Markets for Road Use

Eliminating Congestion through Scheduling, Routing, and Real-Time Road Pricing

Peter Cramton, R. Richard Geddes, and Axel Ockenfels[1]

15 April 2019 [latest draft, presentation, *Nature*]

[1] Peter Cramton is Professor of Economics at the University of Cologne and University of Maryland; since 1983, he has conducted research on market design; he has applied that research to design auction-based markets of radio spectrum, electricity, financial securities, runway access, and other products. Rick Geddes is Professor in the Department of Policy Analysis and Management and Director of the Program in Infrastructure Policy at Cornell University; his research studies private infrastructure investment through public-private partnerships, corporate governance, and regulation. Axel Ockenfels is Professor of Economics at the University of Cologne and Speaker of Cologne Excellence Center for Social and Economic Behavior; his research focuses on behavioral economics and market design. We have benefitted from helpful discussions with Ken Fields, Declan Ganley, Hector Lopez, Bob Poole, Yip Chun Seng, and Leong Wai Yan. Ockenfels thanks the German Research Foundation (DFG) for financial support through the Research Unit "Design & Behavior" (FOR 1371).

Contents

Contents 2
Introduction 4
The time is right 4
The benefits extend beyond eliminating congestion 5
The tyranny of the status quo 6
The challenges to change can be overcome 7
Why non-price approaches to cure congestion fail 8
Dynamic pricing and traffic congestion 9
Current use of congestion pricing 10
New technologies and dynamic road pricing 12
Designing markets for road use 14
Market objectives 14
Evolution of market designs 15
Congestion pricing of road use 16
How today's apps would need to change 18
A wholesale market for transport 20
Governance 27
Promoting acceptance 29
Trade-offs in transport 29
Equity 31
Privacy 33
Behavioral responses, test-bedding and opportunities for early implementation 34
An incremental approach to road pricing in Singapore 36
Step 0: Identify the issues with the current system 36
Step 1: Introduce GPS devices, measure congestion and simulate road pricing 36
Step 2: Introduce initial time-and-location prices 36
Step 3: Re-estimate prices and introduce refined time-and-location prices 37
Step 4: Repeat step 3, periodically refining prices and learning behavioral response 37
Step 5: Introduce real-time pricing 37
Step 6: Introduce forward purchase 37
Step 7: Study long-term impact 38
Summary of implementation plan 38
Conclusion 38
References 39

Summary

Traffic congestion is a global problem with annual costs approaching $1 trillion. The cost of traffic congestion across the combined British, French, German and American economies was estimated at $200 billion, or about 0.8 percent of GDP, in 2013. In Los Angeles alone, traffic jams cost $23 billion each year. The health and environmental costs are severe in urban centers worldwide.

With the right policies those high social costs can be avoided. Advances in mobile communications and computer technology now make it possible to efficiently schedule, route, and price the use of roads. Efficient real-time pricing of road use can eliminate traffic congestion, enhance safety, improve the environment, and increase vehicle throughput. It also raises reliable, much-needed revenue to modernize decaying infrastructure while improving long-term investment in transport. We describe the design of a market for road use that is based on efficient scheduling, routing, and pricing. Under our design, road use is priced dynamically by marginal demand during constrained times and locations. In unconstrained times and locations, a nominal fee is paid for road use to recover costs, as in other utilities. Transport is scheduled based on forward prices and then routed in real time based on real-time road-use prices.

Efficient pricing of network capacity is not new. Indeed, wholesale electricity markets have been dynamically priced for over a decade. Communications markets are adopting dynamic pricing today. Efficient pricing of road use, however, has only recently become feasible. Advances in mobile communications make it possible to identify and communicate the location of a vehicle to within one cubic meter—allowing precise measurement of road use. User preferences can be communicated both in advance to determine scheduled transport and in real time to optimize routes based on the latest information. Computer advances also facilitate efficient scheduling and pricing of road use. Consumer apps help road users translate detailed price information into preferred transport plans. Computers also allow an independent system operator to better model demand and adjust prices to eliminate congestion and maximize the total value of road infrastructure. An independent market monitor, distinct from the operator, observes the market, identifies problems, and suggests solutions. A board governs the market subject to regulatory oversight.

The market objective is to maximize the value of road infrastructure via scheduling, routing, and real-time pricing of its use. The optimization of road use eliminates congestion, making our roads safer, faster, cleaner and more enjoyable to use. The road-use market maximizes the value of our existing transport infrastructure while providing essential funding for the roads network as well as valuable price information to evaluate road enhancements. The market is complementary with and indeed promotes rapid innovation in the transport sector.

Introduction

Traffic congestion is a pervasive and growing worldwide problem. Global congestion costs are estimated at about $1 trillion. The cost of traffic congestion across the combined British, French, German and American economies was estimated at $200 billion in 2013, about 0.8 percent of GDP, and these figures are expected to strongly rise in the next years (*The Economist* 2014). In Los Angeles alone, traffic jams cost $23 billion each year. Two-thirds of those costs stemmed from wasted time and fuel, with the remainder due to increased costs to businesses; they do not include the price of needless carbon-dioxide emissions and local pollution. For the United States, traffic congestion wasted 6.9 billion hours of motorists' travel time and almost 3.1 billion gallons of fuel in 2014 (Schrank et al. 2015). Moreover, congestion's social costs are growing over time. The U.S. congestion "invoice" for added costs in terms of fuel and time grew from $42 billion in 1982 to about $160 billion in 2014 (in 2014 dollars)—almost a three-fold increase—in the 471 urban areas studied by the Texas Transportation Institute (Schrank et al. 2015). Pishue (2017) provides more recent data on U.S. congestion. Similar increases are occurring worldwide as car ownership rises rapidly with development. Roads are thus failing to perform their core task of safely facilitating vehicle passage with minimal environmental impact, and the problem is getting worse.

The time is right

Fortunately, a major improvement in transportation performance will occur over the next twenty years, if not sooner. Complementary technological advances—especially in computers and communications—have the potential to greatly reduce congestion, thus improving mobility. We explain how to harness existing technological developments which, when combined with advanced markets in road use, can *fully*

eradicate traffic congestion along with its attendant social harms. In our terminology, a road network operating to maximize throughput is congestion free. Of course, in this max-flow state, there is congestion in the sense that each additional car slows down the others, but the congestion is not destructive—it does not reduce throughput. Adding traffic beyond the max-flow capacity level, however, is destructive.

These improvements in transportation performance can occur now. Today's technology allows for monitoring a vehicle's road use and charging directly for that use. Real-time charging for road use allows efficient pricing of congestion. With the right prices to guide behavior, traffic congestion can be eliminated. Efficient road-use pricing not only ensures that the right vehicles are using the roads at the right times, but dramatically increases the number of vehicles that can flow through the available lanes (i.e. throughput) during peak periods. Traffic jams are avoided.[2]

The benefits extend beyond eliminating congestion

Efficient road-use pricing generates many additional benefits, however. It improves environmental quality since vehicles operate more efficiently, spend less time idling, while pollution costs are recognized explicitly via road prices. Safety improves due to more consistent, predictable traffic flow. Perhaps as importantly, real-time road prices provide essential information to identify those investments most needed while generating the funds that underlie that investment.

Wholesale electricity markets offer a useful analogy.[3] They have operated on a similar basis for over a decade (Cramton 2017). Electricity resources are optimally scheduled and priced one day ahead, after which a real-time market allows for necessary adjustments throughout the day. In transport, crude versions of the pricing we envision are observed in certain express lanes, in Uber's surge pricing (Castillo et al. 2017), in train ticket pricing, and in airlines' pricing of seats, among other examples. Our approach goes beyond those early versions, however, building on the proven success of electricity markets. In our framework, transport is scheduled and priced in advance with real-time routing and pricing to reflect inevitable changes in demand.

In addition to mitigating traffic congestion—and unlike most taxes—dynamic road-use pricing creates a non-distortionary revenue source. The resulting congestion prices also provide essential information to direct scarce investment resources toward projects where those dollars are most highly valued by road users and away from lower-valued projects. Dynamic pricing thus creates the crucial link between customers' value of a facility and investment flows. Stated differently, congestion pricing provides

[2] Vickrey (1963) is the seminal theoretical study on the efficiency of road pricing, but the time was not right then for implementation. As Harstad (2008) described it, Vickrey "was appalled at the notion of adding to traffic congestion to collect tolls, and railed against tollbooths, urging the development of a system where small radio transmitters would transmit vehicle or driver IDs over a distance of a few feet, and a computerized system connected to roadbed receivers would calculate liabilities and bill drivers periodically. A few years afterward, Vickrey was challenged that the system he proposed was infeasible. He responded in typical fashion: in the mid-1960s, he first built a rudimentary computer in his home and connected it to a radio receiver, then limited himself to a $3 budget for parts with which he built a small radio transmitter placed under the hood of his car. He could then show anyone who asked a printout of the times his own car went up or down his driveway."

[3] A second analogy is in wireless communications (Cramton and Doyle 2017), but time-and-locational markets in communications are only beginning to emerge.

objective market signals regarding where additional investment should be directed that are based on motorists' willingness to pay for that investment. Our proposal thus also addresses one of the most challenging problems facing transportation policy today: the perceived misdirection of scarce public dollars caused by politicization of spending.

Price information alone is inadequate to ensure that the best investment choices are made, however. Good governance and oversight of the planning and investment process is also critical. Yet price information remains essential for expert planners guided by an objective of maximizing social welfare. This approach has worked well in restructured electricity markets, especially the successful market design used in most of North America (Cramton 2017). That system prices energy every 5 minutes at every location. In electricity, regulatory capture by special interests is reduced or avoided entirely via good governance and a data-driven analytical planning process.

Although these are the core benefits, congestion pricing is appealing for several additional reasons. Those include: (i) divorcing charges for the use of road space from fuel type used, which makes road charges independent of rapidly evolving engine technology; (ii) adopting the basic horizontal-fairness principle that the motorists using a road should pay for them, which enhances social equity; (iii) allowing scarce road space to be allocated to motorists who value it most highly at that specified time of day; (iv) incentivizing technological innovations that reduce the cars' demand on scarce road capacity; and (v) encouraging commuters to explore the travel alternatives of their choice during peak times by providing current road-use prices. Although partial, this list suggests that the social benefits created by dynamic road-use pricing are likely to be substantial. Moreover, such benefits are apt to be shared broadly. All citizens share in the health and environmental benefits of eliminating congestion, and nearly all benefit from improved throughput.

The tyranny of the status quo

Commentators have been aware of road pricing's many benefits for decades. Writing in the early 1950s, Milton Friedman and Daniel J. Boorstin state that:

> At first glance, it seems hardly possible that this apparently trivial problem of how to charge people for the highway services they use is a key to the whole problem of how to plan and pay for better highways; yet it is just that. This fact cannot be too strongly emphasized. It is a key not only for a system that would involve operation of roads by private enterprise but equally for the present system of public operation. Should a particular road be built? How should it be built? How should it be financed? Should an existing road be maintained, improved, or allowed to deteriorate? If we could charge directly for the service of the road, we could answer those questions—whether under private or public ownership—in the same way that we now decide how many automobiles should be manufactured, what kind of automobiles should be manufactured, how their production should be financed, whether a particular model should be discontinued, and so on (Friedman and Boorstin 1996, 223).

At the time Friedman and Boorstin were writing, tolls were paid entirely in cash. Widespread tolling implied stopping to pay cash tolls, thus slowing travel. Massive leaps in communications and computer

technology since then make new charging approaches feasible. Nonetheless, reform is slow. This would come as no surprise to Milton Friedman:

> There is enormous inertia—a tyranny of the status quo—in private and especially governmental arrangements. Only a crisis—actual or perceived—produces real change. When that crisis occurs, the actions that are taken depend on the ideas that are lying around. That, I believe, is our basic function: to develop alternatives to existing policies, to keep them alive and available until the politically impossible becomes politically inevitable. (Friedman 1982, ix)

Despite broad academic agreement, use of direct road-use fees in the United States is limited. Although Oregon debuted a system-wide user-fee program (Morris 2015), road pricing is most often limited to new lanes, such as high-occupancy toll (HOT) lanes or on conversions from high-occupancy vehicle (HOV) to HOT lanes. This has left existing transportation facilities—often older roadways in need of fresh investment—out of the funding streams generated by that pricing. Moreover, only a few small sections of the U.S. highway system are dynamically priced.

Looking abroad, many cities around the world have adopted some form of congestion pricing. Examples include Singapore, London, Oslo, Stockholm, and Trondheim. However, its application is typically limited to a cordon around the city, or to a small set of roads within. Moreover, prices are typically not responsive to demand changes in real time. Our proposal represents a significant leap forward to what is the most efficient end-state for road pricing. We propose a comprehensive system of direct, variable road-use charges. This creates a market for transport that maximizes the value of the transport network through efficient scheduling, routing and pricing of road use.

By scheduling of road usage, we mean an advance plan of when road segments are used and by whom. In some cases, scheduling can occur well in advance of actual usage. In other cases, it may occur one or more hours ahead. Scheduling contrasts with routing in that routing results in a real-time usage plan. Scheduling, however, requires at least some demands to be expressed in advance of real time. Of course, schedules can change in response to unexpected events and new information. Our proposal can handle such changes efficiently.

Under existing policies, traffic is getting worse and will inevitably end in a crisis. Our proposal for a market for road use eliminates the inefficiency, frustration and unfairness that comes with the current system.

The challenges to change can be overcome

It is efficient and fair that those who use roads at peak times and thus place the greatest demand on the road system pay a higher price than those who avoid traveling at peak times and thus generate much smaller social costs. But there is much resistance against road pricing. Analysts have attributed the limited use of road-use fees in the United States to motorists' opposition to new rates and fees. Because motorists often think that roads are "already paid for" via gas taxes, dynamic pricing may initially be viewed as a new tax. This suggests that rebating state fuel taxes paid when motorists pay usage fees (the approached used in Oregon) is critical. Such issues are particularly important for existing transportation facilities, where user-fee revenues are usually utilized to improve road quality rather than to add the additional

capacity that is easily visible to motorists. Many commentators argue that displaying tangible improvements from road pricing enhances motorists' acceptance of user fees (Poole 2014). The improved throughput and absence of congestion are two salient benefits—both seen and felt by drivers. Also, there are various effective policy options to strongly alleviate potential concerns about affordability of congestion pricing, similar to policy options to alleviate affordability and equity concerns when pricing water, electricity and communications.

A key challenge in implementing a real-time market for road use is coordinating the multiple jurisdictions (e.g., city, county, state, and federal) that own various adjoining facilities. Although demanding, the large gains from coordination that our approach offers can motivate jurisdictions to work together, especially when regulators properly encourage coordination. Again, this has worked well in U.S. electricity markets where federal and state regulators have played an essential role in facilitating coordination among market participants.

Direct road-use pricing may also raise privacy concerns. Privacy is however readily addressed through a policy limiting who sees what information about a vehicle's road use. Those same issues are successfully addressed in other industries, such as communications. Who and when calls are made are recorded and priced, but access to this information is strictly limited. Further, many users are today willing to provide locational information to service providers, such as Waze, Google and Uber, in exchange for improved services.

Enforcement is an added concern. A mechanism to ensure that users pay for road use is critical. There are also standard ways of addressing enforcement. For example, license plates could be randomly photographed, and fines applied whenever the vehicle's road-use device is not operational. As an aid to the user, the vehicle could provide a warning whenever the road-use device is not engaged.

Why non-price approaches to cure congestion fail

Numerous studies have shown that an increase in road capacity does not relieve congestion. Whenever a new road is built, or a new lane is added, current residents drive more, there is more commercial traffic, and more in-migration. This is called induced demand: building more roads will bring even more cars on the roads. This results in "The Fundamental Law of Road Congestion" (Duranton and Turner 2011), which states that, after a certain time, if new unpriced capacity is added, the road will become just as congested as it was prior to the capacity addition. It is related to the simple observation that, without congestion pricing, there is always more demand than road capacity at peak times. Even multibillion investments in new roads in cities like Los Angeles and Houston did little to reduce commuting times. Despite the "Fundamental Law" metro areas that invest more on freeway capacity have less congestion than those that have invested less, and pricing promotes more targeted investment to better address congestion (Poole 2018, text box 6).

Similarly, there is a widespread belief that an increase in public transit, in ride-hailing services like Uber and Lyft, or in self-driving cars, will provide a cure to traffic congestion. There is little support for that view. Providing public transportation frees up road capacity by taking drivers off the roads. The fundamental law of road congestion predicts that this will have little effect on congestion without variable

road pricing. Indeed, that is what the data suggest (Duranton and Turner 2011). Ride-hailing services appears to increase traffic (Schaller Consulting 2017), which probably contributes to why Uber supports road pricing as "the most effective way to manage vehicles on the road" (Salzberg 2017), and why Lyft suggests that "congestion pricing … has not caught on in a big enough way" (Zimmer and Green 2017). Although self-driving cars use roads more efficiently, it is doubtful that this will substantially relieve traffic congestion (Fulton et al. 2017, Henzelmann 2017, SBD 2016). The one approach known to solve the gridlock is dynamic pricing.

Dynamic pricing and traffic congestion

Moving to direct road pricing is an important step in improving transportation policy. However, its effect on traffic congestion will be limited unless the price charged varies with the cost that one motorist imposes on others when utilizing a facility at a specified time of day. Like water flowing through a pipe, or a wire carrying electrical waves, a road, bridge, highway, or tunnel has a physical capacity limit. A highway lane can only transport a certain maximum number of vehicles per hour when traffic is flowing freely. When too many motorists try to use a facility at the same time, the facility becomes congested, like a clogged water pipe. Traffic flow collapses, and the bridge, highway, or tunnel is unable to handle as many vehicles as its physical capacity allows.

This fact—that congestion reduces throughput—is one of the central results of transportation science:

> The number of vehicles that get through per hour can drop by as much as 50 percent when severe congestion sets in. At high-traffic levels, the freeway is kept in this condition of "collapse" for several hours after the rush of commuters has stopped. (FHWA 2009)

Drivers generate various costs when they use their vehicles. Those include fuel cost, depreciation of their vehicles, time spent driving, and wear and tear on transportation facilities, as well as the crowding out of other motorists on the road at the same time. Drivers bear many of those costs directly and will take them into account in their decisions about when, how often, and how far to drive. They will, for example, directly bear fuel costs, vehicle depreciation, as well as time costs.

They will not, however, directly bear the costs associated with facility wear and tear or the crowding out of other motorists who also want to drive at that time. Drivers will thus not take them into account in their driving decisions. Their decisions will be dis-aligned with the true overall cost of using roads at a specified time of day. Facility wear-and-tear costs can easily be addressed through a nominal fee for road use.

That second component, however, is more complex: motorist crowding has highly non-linear effects on travel time. That is, as a road gets close to its physical limits, even a small increase in the number of vehicles leads to a large drop in average speed for all vehicles. If one additional motorist uses the road at those peak times, she imposes a large cost on all other motorists through slower travel times. For example, the motorist who decides to use a highway at 3 a.m. imposes no crowding or congestion costs on other motorists, since there is usually excess road space. One who instead uses the same highway at 8 a.m., when many other motorists also want to use it, imposes substantial congestion costs.

Under the current approach, the individual motorist does not consider those large social-crowding costs. The second key aspect of road pricing is therefore a variable charge—or congestion price—that reflects the costs of social crowding. The phenomenon described above suggests the effectiveness of such charges. If even a small number of motorists can, via congestion prices, be incentivized to drive at other times, to use alternative modes of travel, to carpool, to telecommute, or to adjust in any number of other ways, then traffic flow will rise disproportionately for the remaining motorists who choose to use the facility at that time. Reducing the number of drivers by as little as 5 percent at peak times may enable traffic to flow smoothly, allowing the same road to handle many more vehicles. Variable road prices therefore have the effect of allocating scarce road space at peak times to those who value it most highly. The highest-valuing motorists will choose to use it at those times and pay the associated higher fee.

Available transportation capacity is currently instead allocated by queuing. Queuing is especially wasteful in transport because the queue degrades throughput, much like overfishing degrades the fishing resource. Queuing ignores the cost that one motorist imposes on others in trying to use a facility at peak times, as well as the differing values motorists place on the road's use at specified times.

Because congestion prices keep traffic moving smoothly, travel times also become more predictable when road space is efficiently priced. This is critical to parents, for example, who need to pick their children up (or drop them off) from school or day care at specific times, as well as to companies relying on just-in-time inventory techniques. It also reduces the time wasted in planning for possible congestion, or in leaving a time cushion to allow for travel uncertainty. Furthermore, congestion prices help ensure that roads are used more evenly throughout the day.

In general, calculating the marginal social cost of a vehicle at a time and road segment is difficult. However, under our approach, where road use is limited to max-flow capacity, the marginal social cost is simply the marginal vehicle's value of using the road segment at that time. Ostrovsky and Schwarz (2018) also make this simplifying assumption. Further, this value is readily modeled and confirmed by the choices drivers make, a process known as *revealed preference* in economics. Thus, it is not necessary for users to explicitly express a value for road use. The relevant information is contained in the decisions that users make. This makes the system operator's task of optimizing the transport network relatively straightforward, as we describe below.

Current use of congestion pricing

Congestion pricing is not a novel concept. As noted, similar variable charges have been successfully utilized in many other industries. For example, airfares, cell phone rates, electricity rates, room rates at hotels and resorts, train fares, and some local transit systems use variable pricing.

Congestion pricing has also been used successfully on a small number of U.S. roads. Real-time congestion pricing is currently used in Minneapolis on the I-394 MnPass Express Lanes and in Washington DC on the I-495 Express Lanes. It is used on the I-15 FasTrak Lanes in San Diego, where prices are updated every six minutes, and on the I-25 express lanes in Denver. It is also used on the SR-91 Express Lanes in Orange County, California, where the price varies between $1.15 and $9.25 per trip and is posted prior to entry so motorists can choose between priced and non-priced lanes. Priced lanes are popular because they save

substantial amounts of time. The Oregon pilot program mentioned earlier suggests that variable pricing could be incorporated into an overall pricing approach.

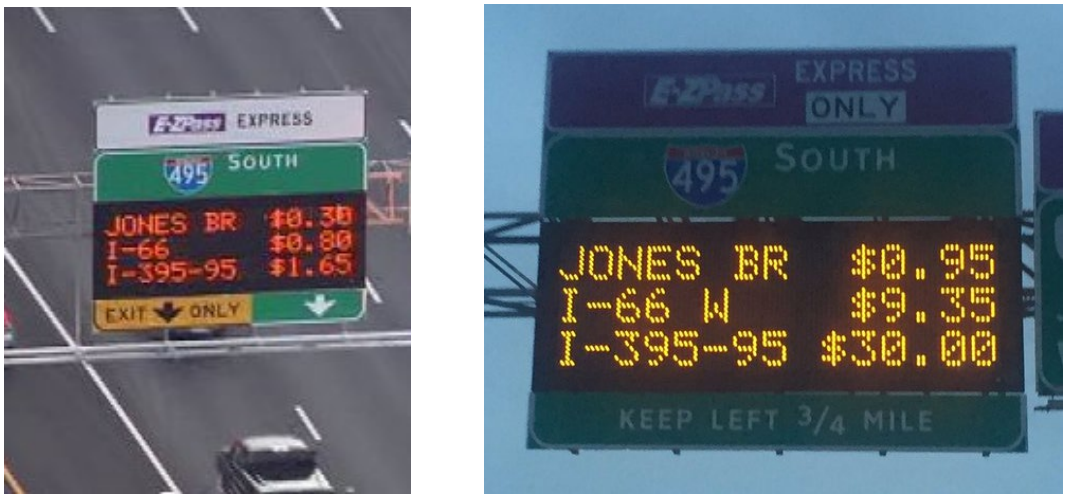

Real-time pricing on Express Lanes in Washington DC

Many international examples of successful congestion pricing are also available. In 1975, Singapore became the first city to implement congestion pricing for urban traffic. Under the original Area Licensing Scheme, cars were charged an additional fee to enter the central business district between 7:30 a.m. and 9:30 a.m. This form of congestion pricing is known as cordon or central area pricing. It was strikingly successful, resulting in a 73 percent decrease in the use of private cars, a doubling of bus usage, and a 30 percent increase in carpooling. In 1995, congestion pricing was extended to three of Singapore's major freeways. On one freeway, average speed during the morning peak increased from 31 to 67 kilometers per hour.

In addition to the demand-side benefit of helping to manage traffic flow, variable pricing creates another benefit on the supply-side: it provides information on how much motorists value the use of roads and thus reveals the most valuable use of the marginal investment dollar. Prices reveal value, and the congestion price required to smooth traffic flows reflects how much value motorists place on a road at that time. High prices suggest that motorists place a high value on using that facility during peak times. The social returns to expanding the road, bridge, or tunnel will thus also be high, and added investment should be directed there. By providing an observable, objective indication of where system expansion should take place, congestion prices also help depoliticize transportation investment, including the improvement of substitute facilities such as public transit. A consensus has largely emerged that tolling (and, importantly, public-private partnerships) can provide critical information on where investment should take place and thus reduce political influences in transportation spending (Geddes 2011).

If road users are prepared to pay a price for road use greater than the costs of providing additional road space (including all costs, externalities etc.) then the additional road space should be built. As in any other economic activity, charges for the use of the new facility should be sufficient to finance its cost. In short, road pricing would yield important benefits on both the demand and supply sides of the transport sector including both roads and public transit.

New technologies and dynamic road pricing

Technological developments in electronics and communications facilitate highly accurate pricing of road use. Currently available technologies allow a road segment to be priced to the sub-meter level. Moreover, road use can be priced in real time based on current conditions of scarcity. Data on the current price of road capacity as established by user demands and measured externalities can be given to a vehicle's on-board computer, directing the driver in real time on what route to take. For autonomous vehicles, real-time road price data can guide the vehicle's route decisions without driver intervention. Real-time adjustment by vehicles in response to highly granular road pricing data has the potential to eliminate all traffic congestion.

Autonomous taxi operating in Singapore beginning in August 2016

Such approaches have been successful in other contexts. The closest analogy is the real-time pricing of wholesale electricity. In restructured markets in the United States, electricity prices at each time and location reflect the efficient congestion prices. Electricity supply and demand is scheduled to maximize gains from trade subject to all physical constraints. The electricity is produced at least cost and consumed by those valuing it the most. The analogous outcome for road usage is efficiently scheduled transportation. Roads would be utilized to provide maximum throughput, priced high enough to reduce demand to the max-flow capacity.

Figure 1: The day-ahead market (orange) enables scheduling and avoids real-time risk (blue)

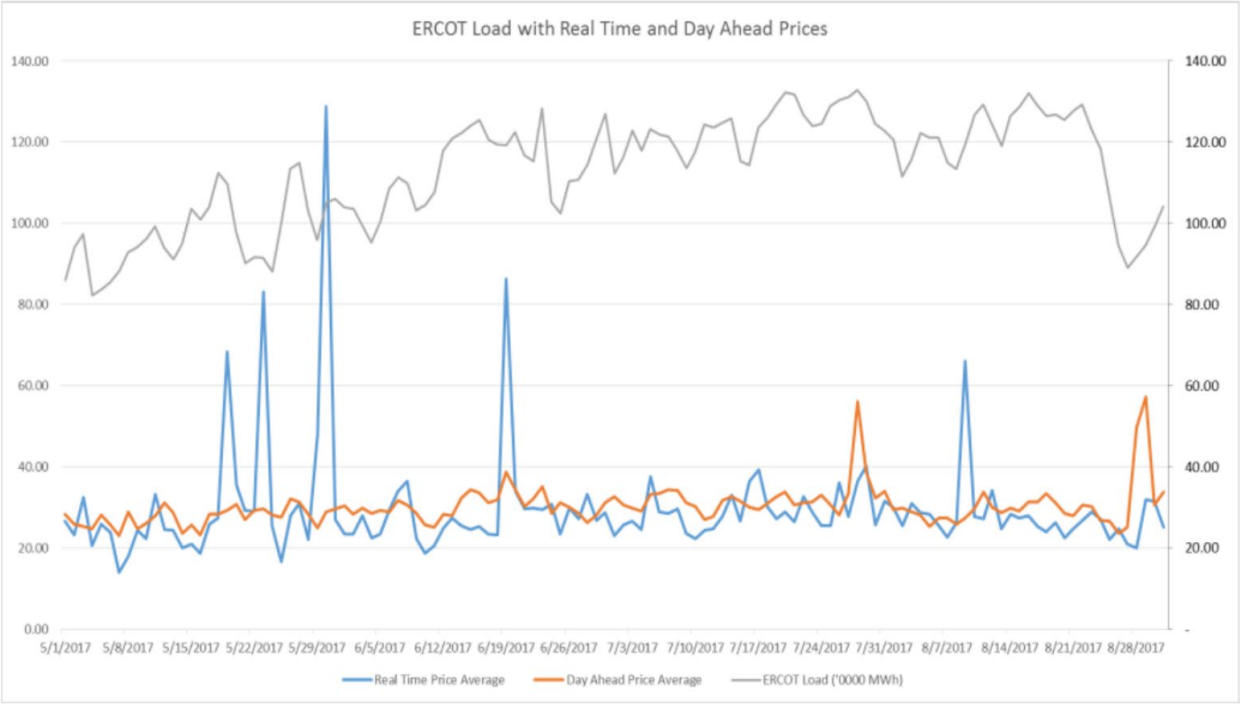

Since real-time congestion prices tend to be volatile, electricity market participants have a desire to manage risks. Forward auctions, conducted in advance of real time, allow participants to make plans and lock in prices consistent with their anticipated needs. Figure 1 shows the real-time and day-ahead electricity prices in Texas in summer 2017. The real-time prices in blue are much more volatile than the day-ahead prices in orange. This is seen in all time and location markets because, as we get closer to real time, participants have fewer options to address shocks to supply or demand. The supply and demand curves that determine price are steeper—quantity varies less in response to price. In electricity markets, most energy is traded in advance of real time. The real-time market, however, plays an essential role in efficiently pricing deviations from day-ahead schedules.

Markets are best tested when they are under extreme stress. Such a period came in Texas at the end of August 2017. Hurricane Harvey over five days in southeast Texas brought 52 inches of rainfall, more than 42,000 lightning strikes, and a record number of tornado warnings. Many substations were underwater, and a significant portion of the transmission grid was destroyed. Remarkably, the Texas electricity system withstood the huge shocks to supply and demand that Harvey produced. Market pricing and system stability were maintained throughout the multiday event. As you can see from the final days shown in Figure 1, day-ahead prices rose sharply during Harvey as participants anticipated the potential for real-time scarcity, but this scarcity did not materialize in real time.

Markets are best tested when under stress—Texas after Hurricane Harvey, August 2017

Such markets also allow traders to arbitrage across related products—such as yearly, monthly, and spot products in the same area—to improve price signals and resource allocation. In modern electricity markets, most energy trades in advance of the real-time market. The forward markets enable planning and hedging of risks, while the real-time market sends the precise price signal to efficiently manage congestion. With road use, more volume is apt to trade in real-time, since individual demands are difficult to predict. Forward markets, however, still will play a useful role for those with more predictable demands.

Designing markets for road use

Market objectives

In developing a straw-man market design for road use, it is best to start with the objectives of the market. Four objectives are salient: efficiency, transparency, simplicity, and fairness. We discuss each in turn.

Efficiency. Short-run efficiency refers to making the best use of the existing road network. Given the size and properties of the road network, this is a complex task. Yet with proper modeling, good data, and state-of-the art economics and computational methods, short-run efficiency is achieved through direct optimization. Provided the demands of participants are truthfully bid and absent other distortions, then the centrally optimized day-ahead market and real-time routing and pricing will achieve an efficient welfare-maximizing outcome.

A second consideration is long-run efficiency. This refers to the market providing the proper incentives for efficient long-run investment for market participants, such as vehicle purchase, and for regulators planning long-term investments in transport infrastructure. This has proven to be an impossible objective

absent pricing. By contrast, an efficient spot market for road use sends the right price signals to motivate efficient long-run investment.

Transparency. The market for road use should be highly transparent. Market rules, including their development and review, should be publicly available. Market data in aggregate form should be available in real time and periodically reviewed monthly, quarterly, and annually. The planning process should also be highly transparent. Transparency helps identify and address problems. It also supports efficient operation and investment.

Simplicity. Like electricity markets, markets for road use are necessarily complex. This follows from the complexity of the engineering and economic problems that must be solved. Designers should nonetheless strive to keep the design as simple as possible, especially for consumers. Complicating features should only be added if they are necessary and consistent with market principles.

Fairness. Key elements of fairness include equal treatment and open access to the market and the road network. Fairness is encouraged through the independence of the system operator and a governance structure that includes representation of all stakeholders.

Our proposed design should draw on best-practice from existing time and location markets. Modern electricity markets, especially the integrated markets operating in most of the United States provide the best example. We will draw on lessons learned in electricity markets, while recognizing the differences between the two settings.

A key market principle in both electricity and transport markets is open access. The transmission grid (electricity) and the road network (transport) are open to all on nondiscriminatory terms. Network capacity cannot be withheld. The congestion price is zero at times and locations where demand at a zero price is less than max-flow capacity. However, on constrained lines the price balances supply and demand and assures that all network constraints are satisfied. When the network is constrained, the price is set at the marginal value of demand. This is called “locational marginal pricing” in the electricity sector. Open access is the basis for restructured electricity markets in the United States, Europe, and other countries. It will be the basis for modern road markets as well.

Evolution of market designs

Congestion pricing for road use is apt to be introduced gradually as the technology evolves and opportunities for implementation are identified. In the simplest application, congestion pricing may be limited to key bottlenecks and new express lanes. Further, pricing may be limited to such a narrow range of facilities and with prices adjusting too slowly to fully relieve congestion. Nonetheless, such steps will reduce congestion and are likely to build support for more significant congestion management.

Additional steps would expand the number of roads with congestion management and relax the constraints on pricing. That will generate further improvements in congestion pricing and reduced congestion. However, the system would remain one in which a system operator is adjusting real-time prices to limit congestion, rather than a full optimization and scheduling of transport.

The final step forward involves both full optimization and scheduling of transport. The system operator would receive preferences from all vehicles for road use at alternative times, such as $10 at the most preferred time to travel from A to B and lower prices at less desirable times. The system operator then optimizes the expressed demands with the available supply to identify the assignment of road use that maximizes total value. The optimization also identifies congestion prices that support the efficient assignment. Although this is a complex optimization, the problem is made easier by the limited number of bottlenecks and the large number of vehicles. The latter reduces the importance of integer constraints in the optimization problem making it more likely that efficient congestion prices exist. Another helpful factor is that aggregate traffic patterns tend to follow a predictable cycle—this greatly facilitates the system operator's modeling of demand.

One important detail is the timing of the optimization. The state of the system, including levels of use, availability, and user preferences, are constantly changing. The initial optimization needs to be done when preferences are relatively complete, not too far from real time, and then re-optimized as circumstances change, say in response to a lane loss. In practice, the system operator strives to sell supply over time in a manner that is consistent with the needs of market participants.

The evolution of the road-use market from primitive to advanced is both desirable and inevitable. This evolution has been observed in other industries. The electricity market may again be the best example. Early markets of about twenty years ago did a good job of facilitating trade across locations, but initially ignored important factors such as congestion pricing. This shortcoming was resolved with the introduction of locational marginal pricing. Many other shortcomings have been addressed both in response to better market rules and improved technology. The result has been a steady and significant improvement in electricity markets over the last twenty years.

The electricity market experience offers key lessons. It is important to adopt an excellent market design early on, but to also ensure that governance and management are structured to generate strong incentives for constant market improvement over time. In electricity markets, scheduling of generating resources plays an important role: many resources are large and limited in how quickly they can respond—both time to start/stop and ability to ramp up/down. As a result, the optimal scheduling of generating resources was introduced in the market early.

Efficient scheduling may be less important in road use because of the large number of vehicles and the speed with which marginal vehicles can respond to price information. This suggests that a useful market simplification, at least initially, is to focus entirely on efficient congestion prices and to initially ignore transport scheduling. We do that in the initial market design discussed below. We then extend the market design to include forward purchase and scheduling. This allows users to then advance purchase the right to use road segments at future times.

Congestion pricing of road use

We next present a market model for road use based entirely on efficient congestion pricing. Operation of the transport system is centralized, although users' decisions regarding road use are decentralized.

Transport in this simplified market is not scheduled. Vehicles and drivers simply respond to congestion prices that are set to efficiently manage congestion.

The market is conducted by an independent system operator (ISO). The ISO's mission is to maximize the value of scarce existing road resources. The key instrument available to the ISO is the ability to set efficient congestion prices. The ISO models demand for road use and establishes prices at each congested segment and at each time to maximize throughput. Usage is monitored and charged to each user based on the marginal cost of segment use including congestion costs.

Demand modeling is critical to establishing efficient prices. Users do not directly express how they substitute across alternative travel times. They instead simply select their most preferred option given the current prices and expectations of future prices. This complicates the ISO's modeling, but the ISO will have abundant data on transport choices with which to model demand and adjust prices. When prices are set too low, congestion forms. When prices are set too high, the segment is underutilized. System uncertainties mean that the ISO cannot operate the road at 100 percent of capacity. Some capacity is reserved to handle momentary surges in demand or drops in supply, as is done in electricity markets. Supply uncertainty arises from road construction, accidents, and weather; demand uncertainty arises from weather and special events, such as the end of a football game, as well as other shocks.

System modeling and management is a complex engineering problem. Nonetheless, the ISO can effectively address this challenge. The problem is well-studied in both economics and engineering when there have been large computational and algorithmic advances in recent years. Most importantly, existing LTE mobile communications allows a vehicle's road use to be precisely monitored, even down to the lane of travel.

A second major challenge is consumer acceptance. Consumers must be able to easily use and understand the system and to realize large benefits. The status quo approach creates much frustration and delay. The reduction or elimination of delay and frustration is thus likely to be the source of consumer gains that will foster acceptance. But ease of use is also critical to consumer acceptance. Sophisticated apps are needed to translate a consumer's preferences and the price information into a recommended transport strategy or menu of choices.

All of this is possible with existing technology. Indeed, most users already have experience with sophisticated apps to facilitate transport. Google Maps, Waze, Uber, and Lyft are examples. Adding road-use prices to those apps is straightforward. Such apps already offer users a menu of choices. The menu will be expanded when road-use prices are included. We can expect rapid innovation in the apps that help users make better and easier transport decisions. Service providers are eager to build apps that create value so that users will use the apps.

A market design focusing solely on congestion pricing is a good starting point. The market may eventually shift to a more complex system that both *prices and schedules* transport system-wide. However, congestion pricing alone is much simpler and likely to go a long way toward maximizing the value of road use.

Such a simple market design would be a good starting point in Singapore. The market could be implemented quickly and would entail little risk. Singapore has three chief advantages as the first mover to implement such a market: (i) Singapore as a city state has a single regulator, the Land and Transport Authority (LTA), that is already committed to introduce the next generation of road pricing; (ii) Singapore already has procured suitable devices for each vehicle to monitor and enforce road pricing; and (iii) Singapore has a long history of innovation in congestion management (LTA 2016).

Figure 2: Integrated market model

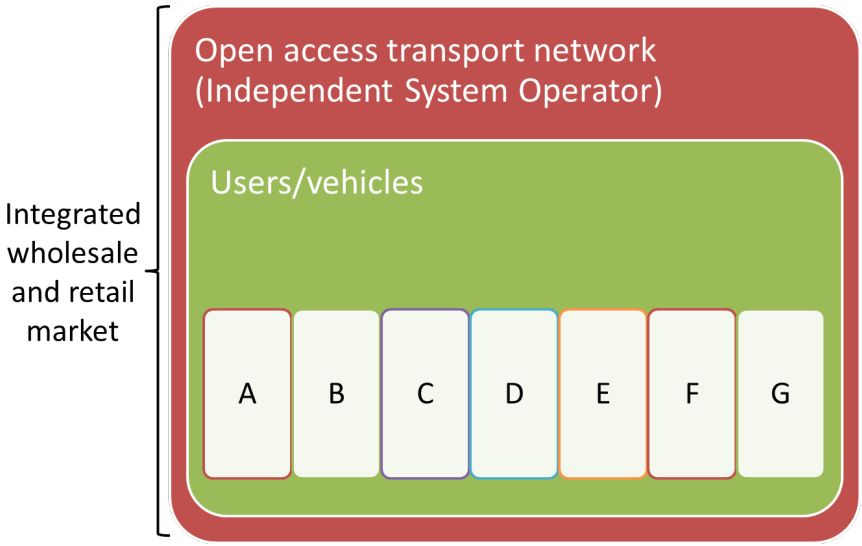

The market in Singapore could be further simplified by integrating the wholesale and retail markets (Figure 2). As a regulator, the LTA could create an independent system operator to run the market. The system operator would model demand and determine real-time road pricing to maintain max flow throughout the road network. Vehicle owners would set up accounts with the system operator for monthly settlement. The system operator would provide real-time road pricing and other relevant data to the public. Service providers, such as Google, Apple, Microsoft, Uber, Lyft, etc., would access this data for use in transport apps, such as Google Maps.

How today's apps would need to change

Remarkably little would change from a user's perspective. Figure 3 shows the current Google Maps phone app, augmented to handle road pricing. Suppose Peter wants to go from home to work. He opens Google Maps and types "W" in the location field and Google puts his Work address at the top of a list of possible destinations. Peter selects it and then Google shows that it will take Peter 32 minutes to drive to work and provides the current estimated price, $3.52 (in red). The red arrow indicates that the price is increasing, so it may be best to leave now.

Selecting the trip yields the next screen, showing the best route as well as alternative methods of transport. Since traffic conditions may change, Peter is assuming some price risk by driving his own car. Price risk falls with time; over a month it is much reduced, given the largely independent shocks that occur in demand and supply that may cause price variation. That is, Peter's average price per trip over the month

depends on roughly 40 independent realizations of price, reducing the standard deviation of the average price over the month by a factor of $\sqrt{40} = 6.3$. This follows from the law of large numbers.

Figure 3: Google Maps showing trip from home to work

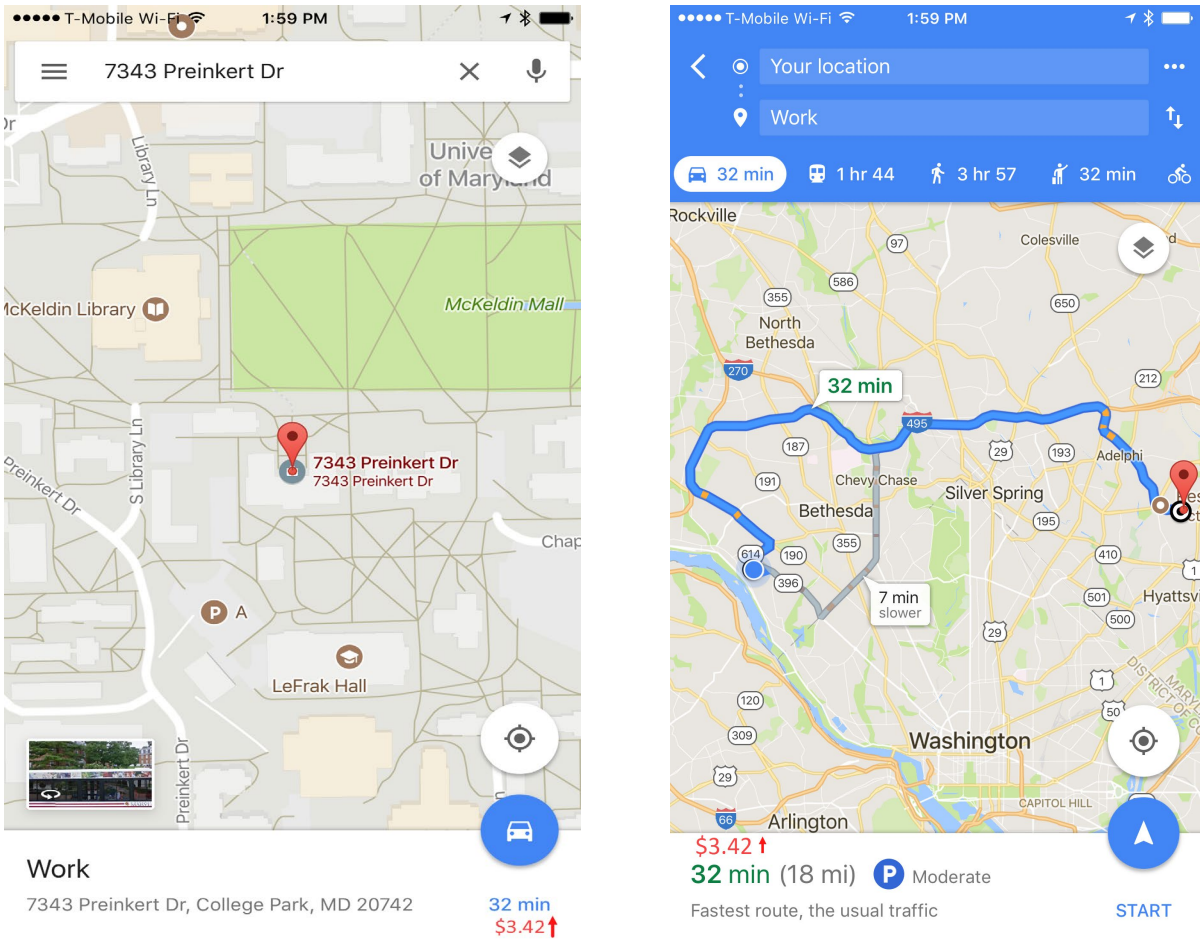

If Peter wants to use Uber or Lyft instead, he simply clicks on the ride hailing icon (Figure 4), which brings up Uber, his default service. From there Peter can see the pricing, select the type of service (UberX), and then book the service. The pricing already includes the road use charge including any additional savings from using POOL—thus it reflects the social cost savings of a shared ride. Remarkably, *nothing needs to change in these screens*. All the work in introducing road pricing is done by the engineers and the computers. The user simply uses the familiar apps in the same way the user does today. Uber in competition with Lyft and other ride-hailing services would likely decide to bear the price risk and offer the customer a fixed price, as they commonly do today. For Uber, the price risk is negligible as it is averaging over thousands or even millions of price realizations.

Figure 4: Google Maps phone app showing Uber request

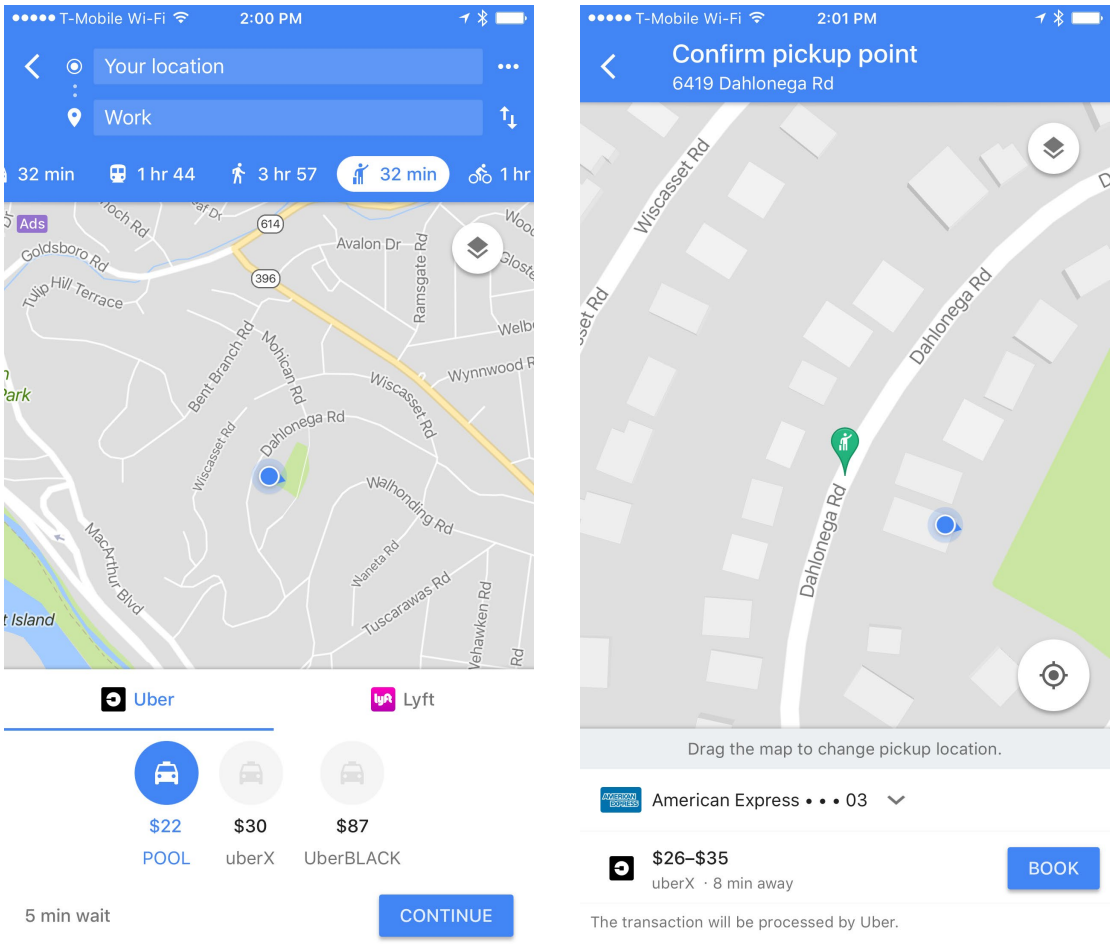

A wholesale market for transport

We next discuss an alternative market design that is built on the wholesale/retail market model. This model has been successfully used in electricity markets for over a decade (Cramton 2017; O'Connor et al. 2015). The independent system operator has the same mission—to maximize the value of road use—but does so by operating a wholesale market in which service providers (e.g., Google, Uber, and Lyft) compete for road use in forward markets as well as in real time, aggregating the demands of individual users. Wholesale pricing is determined in frequent auctions. The market model is shown in Figure 5.

Figure 5: Wholesale/retail market model

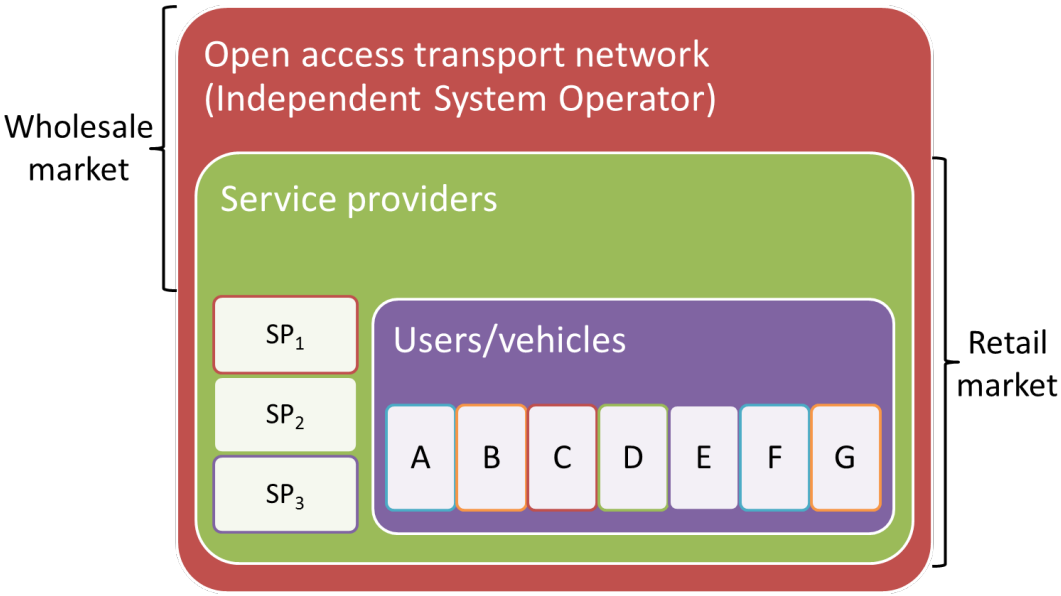

Users provide the fundamental demand for road use. Service providers compete for users in the retail market. Providers that offer more attractive plans are likely to be more successful. Some large companies, such as UPS and FedEx, would participate directly in the wholesale market.

An advantage of the wholesale market model is that entry at the service-provider level is relatively easy. This fosters competition and innovation, which is desirable given the important role played by service providers. In the wholesale market, service providers aggregate user demand. To do this, the service provider must develop a user app that enables users to easily and effectively express demand. The service provider also guides the user, both in scheduling future demand as well as routing during real time. Finally, the service provider establishes user plans and settles payment. We expect a great deal of innovation to occur in service provision.

The wholesale market allows a relatively simple product design. The product is a slot on a congested road segment at a specified time. That is, the "commodity" being traded is the right to use a well-defined section of road for a specific time slot. Time is broken into discrete intervals, such as 10 minutes, to keep the number of products manageable.

There are three promising features of the market. First, the number of congested road segments is limited. The most obvious congested segments are bridges, tunnels, and other bottlenecks. Second, the congested segments are highly predictable. Rush hours are a good example. Third, demand responds to price in close to real time. The demand response takes one of four forms: (i) time shifters, who shift transport to a less congested time; (ii) route shifters, who shift to a less congested route; (iii) mode shifters, who decide to take a train, bus, or bike, rather than drive; and (iv) curtailers, who decide to ride-share, work at home or otherwise reduce demand.

In one sense, responding to prices in transport is nothing new. Today one's "price" is paid in delay cost. Users do respond to this price, but the price is waste and is set incorrectly—the delay cost does not reflect

the negative externality one user imposes on others. In the market model, congestion is eliminated—the real-time price is set at the marginal value of demand at the point of supply and demand balance.

The foundation of the transport market is the real-time market, which prices road use in real time, for example every ten minutes. The real-time market is a physical market, based on actual (i.e. physical) road use. One challenge of the real-time road-use market is that prices can be volatile. It becomes more difficult for demand to respond to events, such as a lane loss, in a short period of time. User apps can redirect traffic given preferences and prices, but the response is limited.

The basic market mechanism is market-clearing. The system operator adjusts the price to clear the market at each time and location. Effectively, there is an auction to balance supply and demand. The auction method is easily visualized when there is a single product. It is a single-price auction, as in Figure 6. Bidders express demand schedules, which indicate the quantity demanded at each price. The ISO forms the aggregate demand curve and crosses it with supply to find the clearing price (P*) and quantity (Q*) where supply and demand balance. All trade takes place at the market-clearing price. Demand (and supply as negative demand) are expressed via a series of price-quantity pairs to form a piecewise-linear decreasing demand curve. Notice that this approach identifies the unique clearing price and quantity. The quantity and its assignment to participants maximizes social value, and the clearing price supports the assignment, since all buyers and sellers are happy with the quantity that they traded given the price.

Figure 6: Single-price auction model

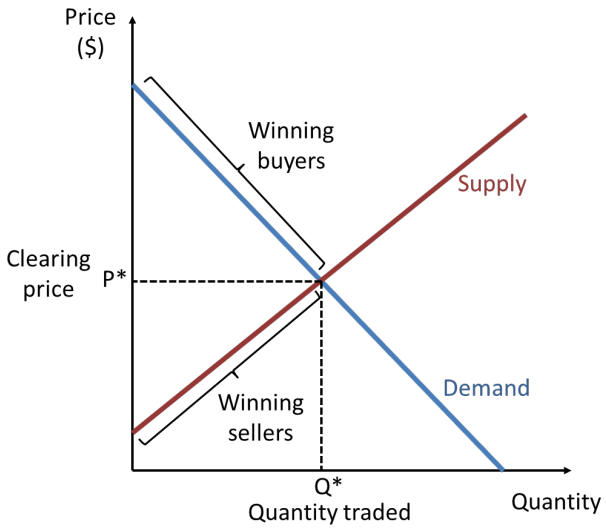

In the actual market, there are many interrelated products—one for each time and location, as well as many capacity constraints to maintain max flow throughout the network. Identifying equilibrium prices and quantities is complicated, but it can be done quickly with modern optimization techniques. Indeed, the problem is easier than the one solved in today's electricity markets, because integer constraints do not play an important role. The number of vehicles that can travel on a road segment in a ten-minute interval is at least several hundred and over one thousand for major highways with three or more lanes. This makes the product effectively a divisible good, and it can be treated as such in forward markets.

Further, with demands expressed as piecewise linear curves, the objective of maximizing gains from trade is quadratic and with linear constraints is readily solved with quadratic programming. Even for large instances with thousands of products and constraints the competitive equilibrium prices and quantities that maximize gains from trade can be computed in seconds using today's computers and algorithms. Maximizing gains from trade maximizes the value of the network subject to max-flow capacity constraints. Feasibility is assured since zero trade satisfies each trader's constraints.

Such an approach can also take advantage of recent advances in cryptography and computation (Parkes et al. 2008, 2009, 2015; Thorpe and Parkes 2007, 2009) to facilitate outcome discovery—both price and quantity discovery—to improve market efficiency with a high level of transparency, while at the same time supporting both trust and privacy. The system operator can thereby ensure that the market rules are being faithfully followed.

Forward auctions enable planning and risk reduction

To mitigate risk and promote efficient scheduling, it is important for the market to provide multiple trading opportunities. We envision four forward (primary) markets: yearly, monthly, weekly, and daily. Secondary markets could operate continuously, but even here we prefer centralized, frequent batch auctions to avoid the arms race for speed seen in continuous markets (Budish et al. 2015). The ISO determines the supply offered in each forward market consistent with customers' and service providers' interest in taking forward positions. The ISO offers more supply at higher prices, creating a rising supply curve. Regarding a schedule of network throughput sales, there are good examples of this from electricity markets, such as the forward sale of congestion revenue rights that allow participants to manage congestion price risk. Figure 7 shows the schedule for offering the congestion revenue rights in Texas for 2021. Each box represents the fraction of rights that are auctioned on the date indicated. The 2021 rights are sold as two products, covering the first half and second half of the year.

Figure 7: Forward sale of congestion revenue rights in Texas electricity system

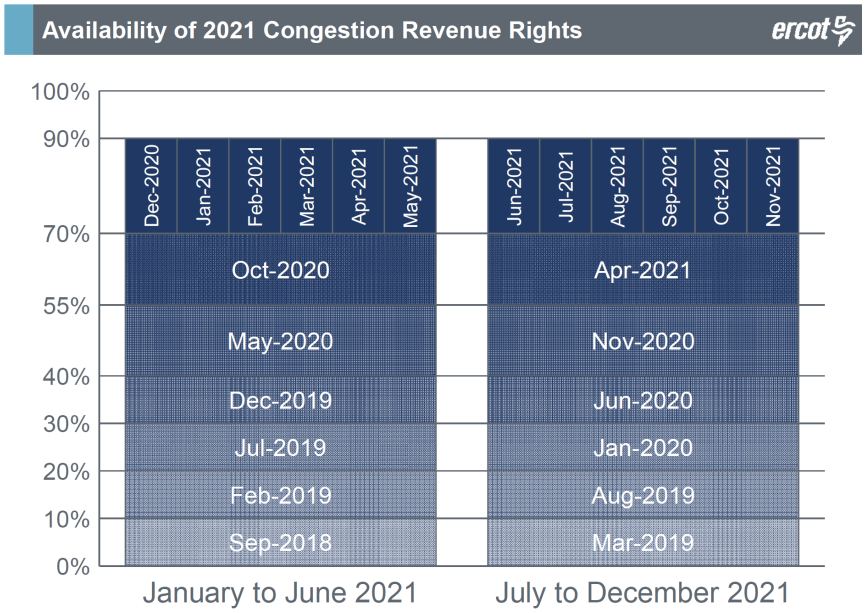

Service providers then bid in the forward and real-time markets to maximize net value to users and manage risk. Typically, this involves purchasing some fraction of user demand in each of the markets and to adjust positions as demand uncertainty is resolved. For the most part, market forces prevent one or more parties from dominating the market. Market rules against manipulation provide further protection against market abuse. Final sources of protection come from two attributes of the real-time operation: (1) real time capacity cannot be withheld and (2) any driver can choose to drive when and where they want.

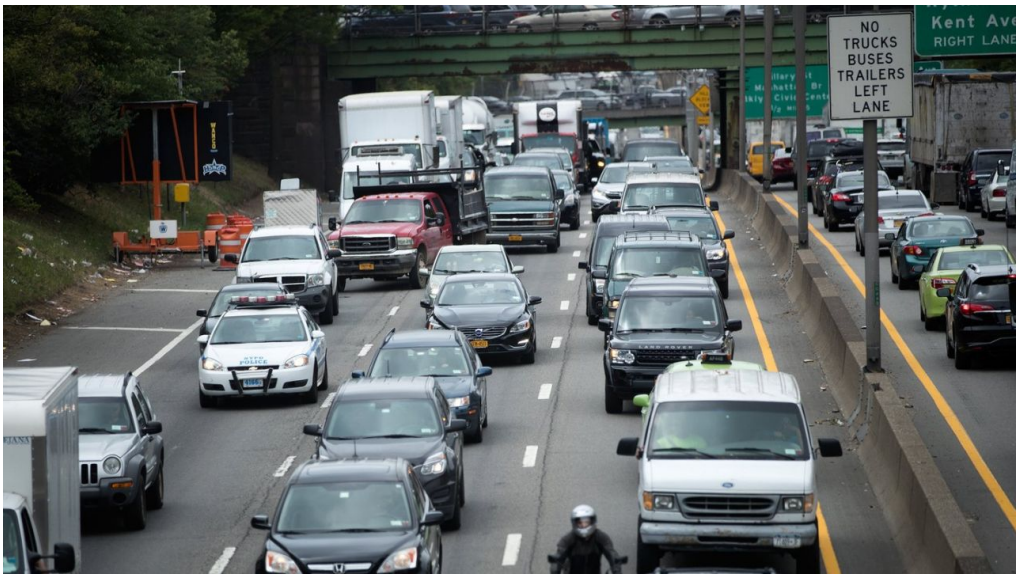

Example of a congested road segment in New York City—the Cross-Bronx Expressway

Figure 8 shows a screen shot from a state-of-the-art auction platform. This is the bidding screen for a monthly auction of the specified road segment, Cross Bronx Expressway from Exit 6A to Exit 2, on a

workday at 8:20 a.m. The bidder's demand in this forward auction is represented with at least two price-quantity pairs. Connecting the pairs forms the strictly decreasing demand curve, shown on the right. The bidder started the auction with a forward position of 1,200 trips per day, which the bidder purchased in the yearly auction. The monthly auction gives the bidder an opportunity to adjust his position—buying some additional trips per day or selling to reduce his holdings. The demand curve expresses a preference to buy more trips if the price is below $22.13 per trip. For prices above $22.13 per trip, the bidder desires to reduce his holdings by selling in this auction. Finally, at prices above $26.86 per trip, the bidder sells more than his entire holding and takes a negative position (i.e. a short sale). Notice how a single curve represents the bidder's preference both for buying and selling.

Figure 8: Sample monthly auction for a congested road segment in New York City

Cross Bronx Expressway from Exit 6A to Exit 2, on a workday at 8:20am

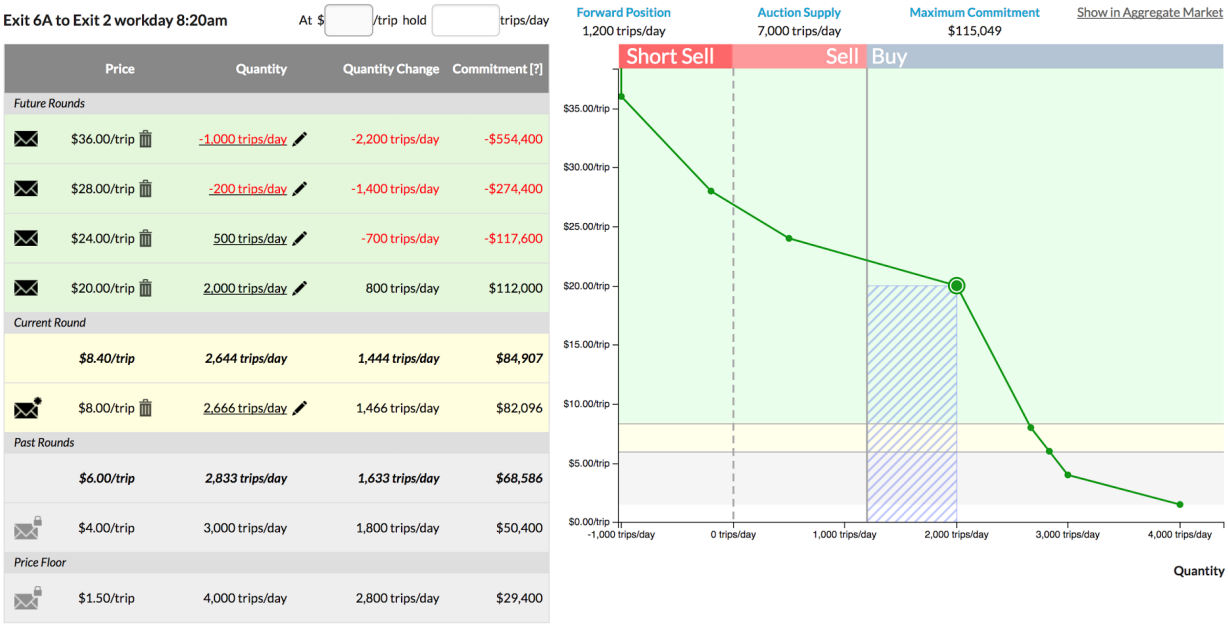

The screen shot in Figure 8 illustrates an ascending clock auction, where price increases continuously like time on a clock. In practice, the auction is conducted over many rounds with the price gradually rising from a low level until excess demand falls to zero. This is a format suitable for settings with relatively few products where price discovery is important. In our setting, where we have many highly substitutable products, such as the trip at 8:10 a.m., 8:20 a.m., or 8:30 a.m., a sealed-bid auction is preferable. This allows each market participant to express preferences for thousands of interrelated products simultaneously. The matching engine then optimizes the entire market at once, determining the assignment and supporting prices that maximize gains from trade. Price discovery in such a smart auction is handled with a sequence of sales, rather than a sequence of rounds.

The forward markets are financial, not physical. Service providers take positions in the forward markets, while subsequent markets allow adjustment of positions as uncertainty is resolved. Speculators also participate in the forward markets, arbitraging across forward markets. Profitable speculators improve price efficiency in forward markets before covering positions in the real-time market.

In the yearly auction, service providers estimate demand for the coming year and bid in the auction based on this demand and expectation of future prices. In the monthly auction, service providers now have better information about demand and can adjust their current positions. In the daily auction, further uncertainty about demand has been resolved. Service providers can further adjust positions. Real-time pricing occurs at the time of physical consumption.

Having multiple opportunities to trade reduces risk to the service provider and facilitates planning. Forward markets also facilitate price discovery through price transparency. Finally, forward trading mitigates market power in the real-time market by putting wholesale customers and service providers in a more balanced position entering the real-time market.

The further forward markets (yearly and monthly) may offer more aggregated products to promote liquidity. However, this may not be needed with the use of a smart market, since market participants can easily express substitution preferences over thousands of alternative products.

With the wholesale market model, service providers compete for retail customers. Offering a service plan that creates the most net value for the consumer is one form of competition. The most sensible and efficient plans will let the consumer express a realistic estimate of usage, informed by past driving behavior, then let the consumer purchase this expected usage on a forward basis to reduce exposure from real-time price volatility. The monthly settlement of such a plan would settle deviations from expected usage with the real-time prices. Thus, the consumer is exposed to the real-time price on the margin—a requirement of efficiency—but almost all the user's expenditure would be at forward prices.

Service providers would also compete by providing useful apps that help customers easily express their preferences. An example retail app is shown in Figure 9. The app lets the user purchase a trip on a forward basis, which may be either a one-time trip or a recurring trip over a specified time-period. In this case, the user is buying a daily roundtrip from home to work for all the weekdays in November. The user enters his preferred departure times and the app calculates a cost per trip. The app also shows how the cost per trip would change if the user departed earlier or later.

Figure 9: Example of a retail app

ForwardTransport, Inc. … *lock in prices and drive with confidence!*

Type	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Daily trip	●	●	●	●	●		

Round trip	One way	Multi-stop	Vehicle
Home		Work	MD 0123

Begin	End
Wed, 1 Nov 2017	Thu, 30 Nov 2017

	Depart	Arrive
Home to Work	8:00am	8:28am
Work to Home	5:00pm	5:28pm

Cost per trip

$ 3.76

Add to Cart

Alternatives and Price Change			
Later		Earlier	
+10 min	$ 0.45	-10 min	$ 0.32
+20 min	$ (0.20)	-20 min	$ (0.33)
+30 min	$ (0.46)	-30 min	$ (0.59)
+40 min	$ (1.11)	-40 min	$ (1.24)
+50 min	$ (1.21)	-50 min	$ (1.34)
+60 min	$ (1.40)	-60 min	$ (1.53)
+70 min	$ (1.56)	-70 min	$ (1.69)
+80 min	$ (1.82)	-80 min	$ (1.95)
+90 min	$ (2.01)	-90 min	$ (2.14)

Importantly, the forward purchase is a means to lock in prices ahead of time. It is a financial transaction, which means that the physical usage of the car is still measured in real time; any deviations from the forward plan are settled at the real-time prices. If the user ends up driving more frequently or at more popular times, then the user will make an additional payment reflecting the extra cost from these deviations. If the user drives less or uses less popular times, then the user would receive a credit reflecting the cost savings from those deviations.

Governance

An independent system operator (ISO) conducts the market for road use. “Independent” means that the ISO has no ownership interest in the road network and no interest in the congestion revenues collected. The ISO is charged by its board to operate the market to maximize the value of the road network by ensuring max-flow traffic conditions. The ISO’s chief instrument to achieve that efficiency goal is congestion pricing, setting prices for road use that mitigate congestion. An effective governance structure is shown in Figure 10.

Figure 10: Governance of the market

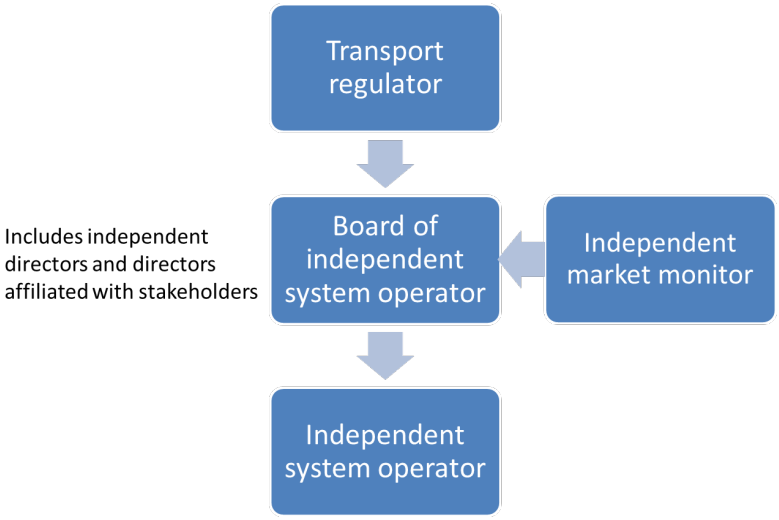

The ISO is a non-profit entity set up with a simple mission: *To serve the public by operating a reliable and efficient market for transport.* The company is analogous to system or grid operators in wholesale electricity markets today. The system operator qualifies market participants and establishes any limits on each participant's bidding activities. The system operator reveals the supply curves for the transport network and conducts all forward auctions. The system operator also conducts the real-time market. Transactions are settled either monthly consistent with the market rules and supply and demand realizations.

The system operator provides information on market performance to market participants and the market monitor. The system operator has a rule-making process to improve the market as problems are identified. In some implementations, the system operator may serve as the default service provider.

An independent market monitor observes the market, identifies problems, and suggests solutions. "Independent" in this case means that the market monitor, in addition to being independent of market participants, is also independent of the ISO. The market monitor brings expert market knowledge. Importantly, the market monitor is not a judge; the market monitor cannot enforce market rules and inflict penalties. The market monitor is instead an observer who writes reports and makes recommendations. In electricity markets, this has allowed the market monitor to quickly identify problems and suggest solutions. The same would be true here. The market monitor reports to the ISO board.

The board oversees the ISO. To ensure that the board includes the knowledge and views of a diversity of market participants, the board has directors who are affiliated with a stakeholder group. The board also has several unaffiliated directors who are independent of the stakeholder groups but bring essential subject-matter expertise. Unaffiliated directors are approved by the regulator.

Promoting acceptance

Before communities take the large step to efficient congestion pricing, researchers must address important questions to promote acceptance: How does pricing alter consumer choice? Can privacy concerns be addressed? Can equity concerns be addressed? Our purpose is not to answer these questions definitively, but rather to provide initial thoughts that satisfactory answers can be found and are apt to support our conclusion that efficient road pricing is good policy.

Trade-offs in transport

Consumer acceptance is essential for road pricing to be successful. An important element of success is providing tangible benefits. The elimination of congestion delays will be the most salient benefit driving consumer acceptance. Anyone living in an urban area—most of the world's population—is well-aware of the frustration of congestion. For those commuting by car, delay is a daily torment of urban life. And it does not go away on the weekends. Weekend congestion is now common in major urban centers.

Congestion delays frustrate drivers

A second salient benefit of road pricing is simplified decision making. Absent road pricing, drivers face uncertainty about how long a trip will take, and those uncertainties depend on departure time. Even adding substantial extra time to account for typical delays may still result in a late arrive, frustrating not only the driver but those who are counting on the vehicle's timely arrival.

Figure 11 illustrates the current decision problem with Google Maps. Peter needs to get from the University of Maryland to a 9 a.m. meeting in Tysons corner, a 22-mile trip mostly on the Beltway. Google provides the best route and shows the anticipated congestion (from orange to red to dark red). Google anticipates the trip will take between 40 and 75 minutes, and recommends that Peter leave at 7:45 a.m. The 7:45 a.m. departure gives Peter a high probability of arriving at the 9 a.m. meeting on time. Peter could also consider leaving earlier. Doing so would increase the probability of a timely arrival and reduce the expected duration of the trip. This may in fact be the preferred choice for Peter, since parking will be

easier, and Peter can read or answer emails while waiting for the meeting to start. Although Peter's objective is clear—arrive at the meeting by 9 a.m.—how to achieve that objective involves many difficult trade-offs in an uncertain setting. It is well-known that decision-making under uncertainty creates stress. The situation is far from ideal.

Figure 11: Today's transport is mostly free, but comes at the cost of uncertain congestion delays

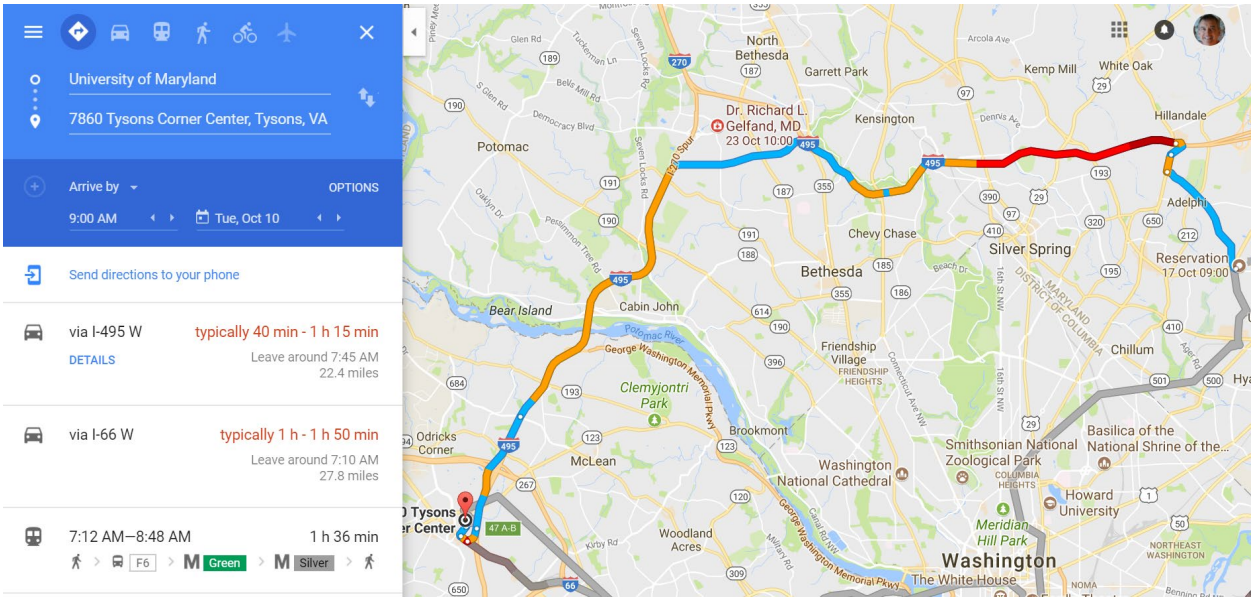

Figure 12: Pricing transport avoids delays and improves throughput

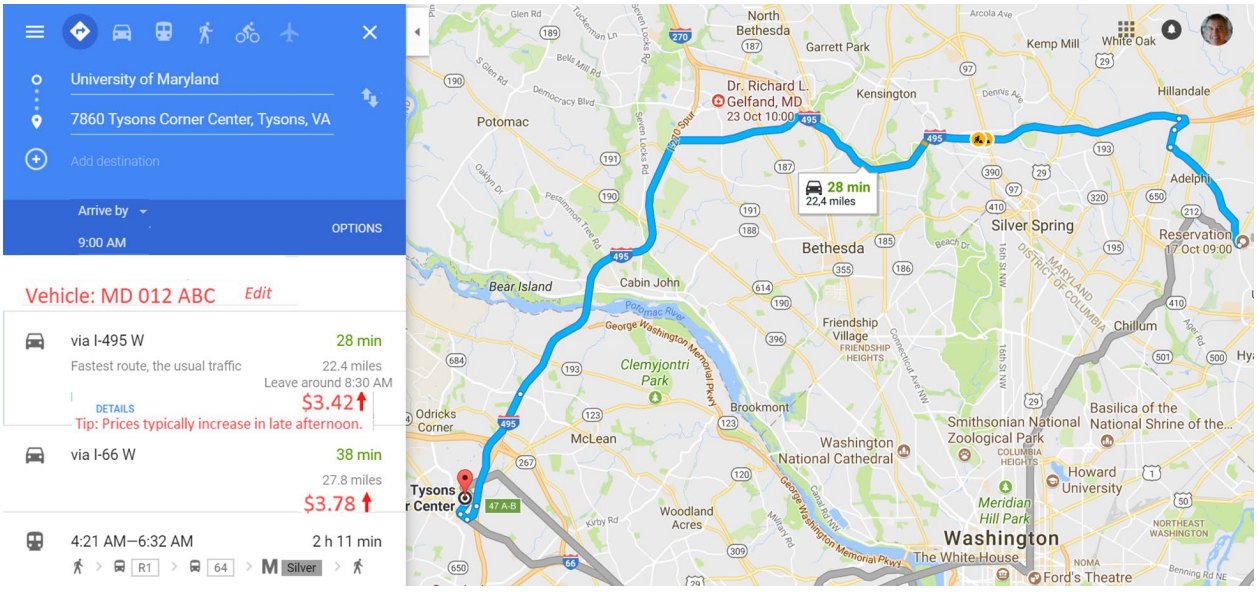

Figure 12 views the same decision problem in a world with road pricing. The first thing to note is how little the decision aid, Google Maps, needs to change to accommodate road pricing. Peter's profile now includes his vehicle's license plate. This allows Google to calculate the vehicle-specific price. The price depends on the characteristics of Peter's vehicle. Peter drives a small gas-powered sedan. Such a vehicle has a lower

price than a large gas-powered SUV, since it uses less road capacity and emits less pollution, but a higher price than a small electric autonomous vehicle, which emits no pollution and uses the least road capacity thanks to optimized driving algorithms and enhanced capabilities, such as platooning. The only change needed to the Google Maps user interface, other than including the vehicle in the user's profile, is the inclusion of price, shown in red in Figure 12. The expected cost for the trip is $3.42 on the fastest route, which takes exactly 28 minutes at max-flow speed. The red up arrow indicates that a departure after 8:30 a.m. would result in a higher price and earlier would result in a lower price.

While the uncertainty about travel time has been eliminated, there now is uncertainty about price, although the volatility of prices is much less than the volatility of travel times on congested roads. Moreover, the harm from price risk is less, because drivers take many trips over a month. What matters for the monthly budget is the volatility of the monthly spend, which is much less, thanks to the law of large numbers. And the price risk can be mitigated further with forward purchase. By contrast, the uncertainty about delay requires drivers who need to arrive on time to take costly actions that exceed the expected cost of delay. For example, without pricing, Peter needs to allow 75 minutes for a trip that is expected to take 55 minutes; with pricing, Peter knows he can complete the trip in 28 minutes.

Both regimes cause some drivers to leave early, late, or not at all, or to switch roads. But the status quo induces an especially costly behavioral response. The congestion causes a loss of throughput of up to 50 percent. And the delay cost borne by drivers is wasteful. Bento et al. (2017) provide an empirical estimate of this delay cost. With pricing, max-flow throughput is maintained, and the road-use cost borne by drivers is retained as revenue, which can be used to build and maintain infrastructure and reduce distortionary taxation.

Equity

A frequent concern about road pricing is equity: that pricing will lead to a road system benefiting the rich at the expense of the poor. This is an essential question. However, our preliminary research suggests that road pricing would have broad benefits across the entire population. There are many reasons for our optimism.

VS

Roads for the rich?

First, road pricing assures max-flow throughput. This yields up to a doubling of capacity on congested road segments at the most popular times. This means that the existing road network can handle many more vehicles—both rich and poor.[4]

Second, all users, rich and poor, benefit from the reduction of uncertainty as well as the benefits from being able to drive when they want to without the fear of congestion delays.

Third, while it is true that poorer users are more apt to shift to less expensive travel times, this implies that low-income households can benefit the most from replacing existing transport taxes with an efficient pricing scheme, and thus can increase consumer surplus (Martin and Thornton 2017; West and Williams 2004; Schweitzer and Taylor 2008).

Fourth, for a similar reason, the poor are likely better able to take advantage of profit opportunities that arise in markets for road use near real time because of supply or demand shocks. For example, a poor commuter may buy a monthly pass to commute at her preferred time, thereby locking in an attractive price per trip of $10. On occasions when there is a lane closure the real-time price may surge to $30. The poor commuter may opt to take public transit on such days and earn $30 per event (her forward purchase would automatically be sold in real time for $30).

Fifth, surveys show that resistance for road pricing projects in operation is particularly low among low-income users. For instance, 70 percent of the lowest income users strongly support San Diego's HOT lanes (FHA 2008).

Sixth, some revenues could be used to improve public transit, an important substitute to driving, especially for the poor. This is the approach recently adopted in New York City. This improvement of the substitute good makes it more likely that the poor will benefit from congestion pricing.

[4] Hall (2018) gives a simple illustration that road pricing does not need to create losers; a carefully designed, time-varying toll on a portion of the lanes of a highway that can make everybody better off, even before the toll revenue is spent and even with realistic driver heterogeneity. On the other hand, Arnott et al. (1994) and Parry and Bento (2001) show that, under certain conditions, many drivers may not want to support congestion pricing solely based on resulting travel time savings in the short run.

Seventh, various options to alleviate concerns about affordability are available, such as a fixed travel allowance for commuters (Small 1992), tax reduction (Parry and Bento 2001), and allocating monthly budgets to spend on congestion tolls (Kockelmann and Kalmanje 2004). We add that, with a market for road use, one could also 'grandfather' vouchers to poor or needy people, which they could then either use themselves or sell on the market.[5]

While more research is needed about the distributional impact of a market for road use, and about the acceptability of proposals to alleviate equity concerns among users and politicians, equity concern can be resolved. Indeed, the norm of a market-based economy is for consumers to buy goods and services at competitive market prices. This is true even for essential services such water, gas, electricity, and communications. Saying that road use pricing is unfair implies that the status quo is fairer. But if we had started with a market for road use, there would be fairness arguments against switching to free road use (e.g., rich people drive more and so benefit more, and poor people suffer more from the increased externalities). The non-pricing of roads stems from the fact that roads were originally uncongested, and that pricing, collecting, and enforcing payments was too costly. In today's urban areas these reasons for non-pricing no longer exit.

Privacy

A common critique of road-pricing is privacy. Monitoring and enforcement require that the system operator know the location of each vehicle during vehicle use. Technically this is easy. Each vehicle would have a device for this purpose, and there would be stiff fines for vehicle operation with a non-functioning device. But users may be concerned that the location information would be used in an inappropriate way—an invasion of privacy.

[5] Extreme surge prices in exceptional circumstances sometimes seem unacceptable to many (Kahenman et al. 1986, Irwin 2017). To further promote acceptability, the market design could include special rules for well-defined exceptional circumstances, as we also see them in other markets.

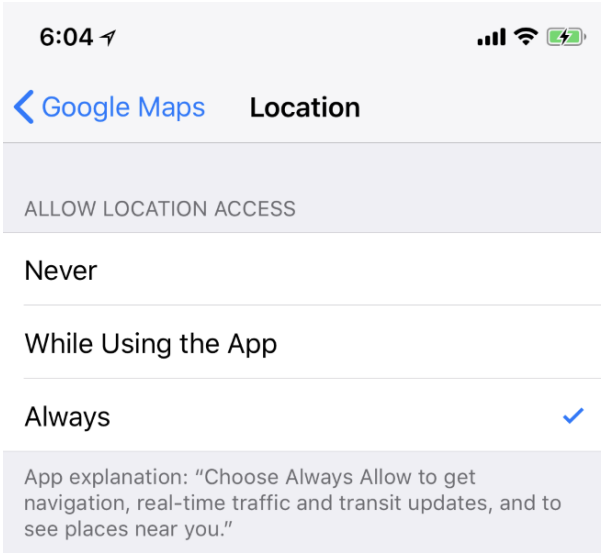

Consumers reveal a preference for improved services over privacy

One response to this concern is that consumers today reveal a preference for letting other parties see location information when by doing so the consumer enjoys improved services. Map apps would be the lead example. The map app becomes much more powerful when it knows the user's location. Nearly everyone with a smart phone uses at least one map app and nearly everyone agrees to give the map app this location information. This consent is given with the belief that the information will be used in a way to help, not hurt, the user. The same could be done for measuring road use. The system operator would have strict rules about how the information is used. The information would only be used to measure road use for purposes of scheduling, routing, and pricing. No individual data would be shared with others. Controls would be in place to assure that the information policy is followed, and the data remain secure. Indeed, modern cryptography makes it possible for the system operator to run the market without any human having access to the individual data (Parkes et al. 2008, 2009, 2015; Thorpe and Parkes 2007, 2009).

Behavioral responses, test-bedding and opportunities for early implementation

How do individuals respond to the incentives set by markets for road use? There is enormous evidence that competitive markets work extremely well if appropriately designed, both in controlled experiments and in the field. The available evidence for markets for road use generally comes to the same conclusion, although more work needs to be done.

Only few highly-controlled laboratory traffic experiments look at behavior along the different relevant dimensions, such as time and space. Rather, they focus on simple, repeated coordination games, sometimes without any pricing (Selten et al. 2007; Chmura and Pitz 2004; Schneider and Weimann 2004; Rapoport et al. 2004), or they include a simple toll (Gabuthy et al. 2006; Hartman 2009). Almost all experiments induce identical driver preferences, inelastic demand and deterministic supply (but see Lopez 2017). That said, most laboratory evidence suggest that people respond to incentives as predicted.

More recently, field experiments started to use GPS data to measure behavioral effects of various congestion pricing schemes. For instance, Martin and Thorne (2017) installed GPS responders in 1,400 vehicles in Melbourne and measured how drivers responded to being charged for road use. The authors found that constant charges do not much affect peak-time behavior, but charges targeted at peak times do reduce congestion (see Kreindler 2017 for a related study based on a large panel data set of GPS travel behavior in Bangalore). Such controlled field studies can strengthen trust in congestion pricing as compared to laboratory experiments, survey studies (Small et al. 2015) and self-reported travel diaries (Karlström and Franklin 2009), but we caution that they may come with their own limitations. For instance, the traffic situation is hardly affected by the experiment treatments, which makes it impossible to measure actual behavioral trade-offs between time, risk and price as would be experienced with optimal congestion pricing. Also, long-term adjustments to congestion pricing, which tend to reduce individual and social costs, are not accounted for in any of those experimental studies.

Road pricing is also found to be effective outside experimental studies. For instance, Foreman (2016) finds a substitution from peak to off-peak times on the San Francisco Oakland Bridge after the introduction of time-varying tolls. Cordon charges have been found to be effective in reducing congestion in London (Leape 2006) and Milan (Gibson and Carnovale 2005).

Overall, the experimental and empirical studies strongly confirm the effectiveness and efficiency of pricing mechanisms to address traffic congestion. Moreover, many studies suggest that drivers substantially differ in their trip scheduling flexibility and in how they value travel time and reliability, which further increases the efficiency of a market-based solution to congestion. That said, more research is needed about details surrounding the engineering of markets for road use. From the behavioral perspective, likely the best approach is to start with simple lab experiments and surveys, and then move to more complex natural settings, refining the hypothesis and designs in the process. This way, experiments can deal with issues such as equity concerns, and cognitive constraints to information processing that require specifically designed user apps and feedback mechanisms to elicit reliable demand information, which in turn will improve price discovery. The simplest laboratory and field experiments start with induced private values and time costs, and a simplified road network; the system operator would be perfectly informed about user preferences. In the more realistic case, the system operator would estimate demands and supplies based on revealed preference over an extended period. The experiments can then gradually improve the external validity by putting subjects under increasingly realistic, data-based stress tests and making the relevant trade-offs in real-time. The next step of the research would use actual supply and demand conditions, and cost estimates, from likely pilot cities, such as Singapore and Phoenix to study the effectiveness of market mechanisms under more 'realistic' conditions, and finally testbed behavioral mechanisms utilizing data from GPS responders in vehicles in the field.

The next step would be to implement the market for road use for field testing. Dynamic pricing of road use can be implemented in any congested area. Both Singapore and the United States are attractive markets for early implementation. Both countries have existing pilots and high adoption rates of the enabling technologies. As we described before, Singapore stands out as the likely first innovator, given its status as a city-state with a long history of innovation in transport and the benefits of a single regulator. In the United States, the Phoenix area stands out as a likely first innovator. Phoenix already has in place

plans for innovation in many areas and has been an early adopter of transport innovations such as autonomous vehicles.

An incremental approach to road pricing in Singapore

Beginning in the early 2020s, Singapore will introduce the next generation of electronic road pricing. The key step in this important transition is the installation of GPS devices in every vehicle that can measure each vehicle's use of the road network, and thereby permit an effective means for time-and-location pricing of road use.

The Land Transport Authority (LTA) is charged with designing and implementing the road pricing plan. Here we propose an approach that ultimately leads to efficient real-time road pricing but does so in steps to mitigate policy risk. In this way, LTA can have a well-defined path to the desired end, progressing in incremental and natural steps. At each step, LTA learns more about the effectiveness of the approach and can make simple refinements that move the system closer to the goal—maximizing the value of the scarce road network. The steps are carefully constructed to be consistent, so that each step is an easy adjustment and improvement from the prior step. System performance is constantly and consistently measured throughout the process, so that LTA and the public can see the benefits of the system and its incremental refinement. The process is dynamic and can be refined as LTA learns more about the benefits of pricing road use.

Step 0: Identify the issues with the current system

This initial step identifies the issues that caused LTA to adopt a GPS-based solution for road pricing. Congestion in Singapore creates delays and uncertainties, and social and environmental costs. The gantry-system measures road use on only some road segments. Pricing is coarse (only at gantries) and unresponsive (updated quarterly). Drivers have an incentive to avoid gantries, causing congestion around the main arteries that are poorly suited for the inflated traffic. Gantries also detract from the beauty of Singapore. Step 0 helps shaping the goals of a new system of road pricing and serves as a baseline to assess performance improvements in subsequent steps.

Step 1: Introduce GPS devices, measure congestion and simulate road pricing

Step 1 introduces GPS devices to measure road use. This will allow analyzing both, individual driver responses to the current road pricing (such as price evasion), and, on the aggregate level, congestion throughout the road network. The GPS data also serves as the input for a simulation of the Singapore road network with the goal to determine steady-state (equilibrium) road prices by time and road segment. These prices are like the prices under the current system but are computed for the entire road network. They are more granular in both time and location than under the current system.

Step 2: Introduce initial time-and-location prices

Step 2 introduces the simulated equilibrium road prices from step 1 to serve as the initial prices for charging for road use. The prices are publicly available and remain fixed for an extended period (one to three months). Providers of transport apps, such as Google and Apple, will be informed in advance, so that their apps can be customized for Singapore. This allows users to easily observe prices for trips, both in real time and in advance. The prices are not true real-time prices, as they do not respond to real-time

events such as lane closures but are fixed throughout the step—just as prices today are fixed throughout the quarter.

The prices are set to maximize the simulated value of the road network. To address uncertainty, the prices are such that road use typically does not exceed a threshold capacity, such as 90% of maximum max-flow capacity. This approach is like the current approach, but with more granular pricing across time and location. We anticipate that Google and Apple and other transport app providers will be eager to update apps for Singapore. They would view this as a major field test of a system that is apt to be adopted in other parts of the world. Further, the integration of pricing is straightforward.

Once implemented, the data provided by the GPS devices before and after the introduction of the more granular prices will make it relatively easy to assess the gains from step 1 (the current gantry-based system), both in terms of the simulated value of road use and measured congestion.

Step 3: Re-estimate prices and introduce refined time-and-location prices

Re-estimate the steady-state time-and-location prices based upon data from the initial trial (step 2). These revised prices are fixed and then used for another extended period (one to three months). LTA continues to measure gains from the policy and any congestion.

Step 4: Repeat step 3, periodically refining prices and learning behavioral response

Repeat step 3 with a one-month duration for fixed prices. Continue until the full gains from fixed, steady-state, time-and-location pricing is measured. During step 4, it may be desirable to make prices more granular in both time and location.

Step 5: Introduce real-time pricing

In step 5, LTA introduces real-time pricing. Prices adjust in real time as congestion begins to emerge. Thus, prices respond to lane closures or other shocks to supply and demand. This does not involve a significant change for those providing transport apps. The providers simply read the real-time prices periodically, such as every ten minutes. Aside from this dynamic updating, and potentially the provision of smart tools to predict and present future prices, the apps are the same. This also makes the transition to real-time pricing straightforward for consumers. Consumers use the same apps as before, but now the prices change in response to events. However, since the prices from earlier steps were steady-state equilibrium prices, the prices are highly consistent with the prior step, except in the face of shocks such as lane closures.

LTA evaluates price risk from real-time pricing in addition to the usual performance measurement of the system.

Step 6: Introduce forward purchase

One disadvantage of real-time pricing is it introduces price risk when necessary to eliminate congestion. In step 6, if deemed useful, LTA introduces a means for consumers to mitigate real-time price risk. One sensible approach is for LTA to establish a default plan based on forward purchase. Users can opt out of forward purchase if they are happy with their experience in Step 5 and not too concerned about price risk.

In the default plan, the consumer buys forward her estimated demand over the entire month. The market operator estimates the consumer's demand based on all available data. Deviations from the forward purchase are priced in real time. Such a plan is a sensible way for consumers to limit price risk. At the same time, it would not diminish the effectiveness and efficiency of road pricing.

Step 7: Study long-term impact

Some benefits of fully comprehensive and fully efficient road pricing will materialize only in the medium- and long-term, when demand for road-use is more elastic, scarcity prices for road-use help to make better investments in vehicles and road infrastructure, and apps that guide drivers become smarter. Evaluating those impacts will provide a complete picture of the benefits of fully efficient road pricing to politicians and the public in Singapore and around the world.

Summary of implementation plan

The key step in introducing the next generation road pricing is the installation of GPS devices in every vehicle that can measure each vehicle's use of the road network (Step 1 above). Once this is done, the next steps incrementally and straightforwardly extend Singapore's current system of road pricing towards the most effective and efficient system to eliminate congestion. This is done by making prices more responsive to time and location of road use. Doing so eliminates the inefficiency and frustration that comes with current road use. Because the transition is done in incremental steps, and the benefits of each single step are constantly measured and monitored, the technological and policy risks of the transition to the most efficient end-state for road pricing are mitigated. Regulators can proceed with confidence.

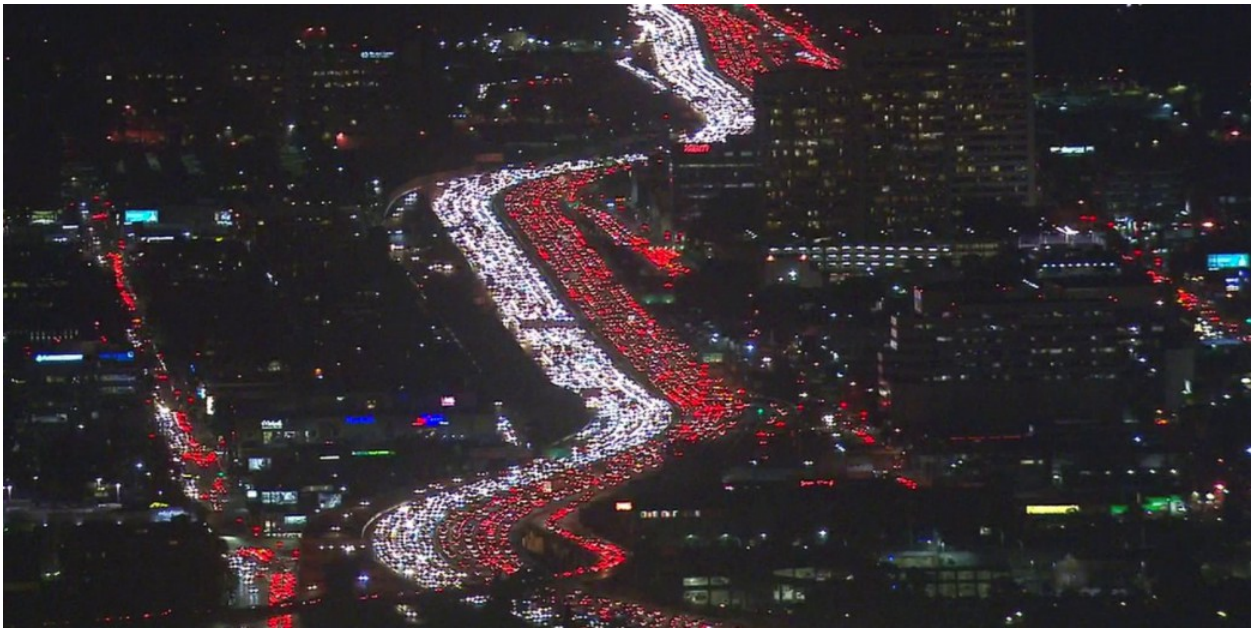

Los Angeles on a holiday

Conclusion

Traffic congestion is one of the biggest challenges of modern societies. It comes at huge individual and social costs and makes the life of many million people miserable. Worldwide congestion costs are

estimated at about $1 trillion per year, and traffic is expected to only become worse over the next years and decades. The "Fundamental Law of Road Congestion" implies that building or widening new roads, or other supply-based policies, will not solve the problem.

Fortunately, technological advances allow us to largely eliminate congestion through efficient congestion pricing. In the simplest approach, an independent system operator models demand and computes real-time prices for road use to maximize the value of road use. Prices on congested road segments are set to induce demand to eliminate the congestion and maximize throughput. User-friendly computer apps armed with this price information then guide consumers in making transport choices consistent with their preferences. Equity and privacy concerns can be effectively dealt with.

A more sophisticated market design is based on a wholesale market model, as we see in electricity markets. The advantage of a wholesale market is that it allows relatively easy entry as a service provider. The ensuing competition among service providers then promotes innovation. That innovation helps service providers to better understand user demands, translate user demands into bids in the wholesale market, and develop forward trading strategies to mitigate risk. A complete system of scheduling, routing, and congestion pricing may seem like a radical idea. It, however, has been successfully applied in electricity markets for over a decade, and is increasingly adopted in communications markets.

Modern mobile communications allow road use to be monitored and charged based on real-time scarcity. Doing so gets the most out of our existing transportation infrastructure and simultaneously provides essential funding of the roads as well as valuable price information to evaluate road enhancements. This is the inevitable future of roads.

References

Arnott, R., A. de Palma, A., and R. Lindsey (1994) "The Welfare Effects of Congestion Tolls with Heterogeneous Commuters," *Journal of Transport Economics and Policy*, 28, 139-161.

Budish, Eric, Peter Cramton, and John Shim (2015) "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130:4, November, 1547–1621.

Bento, Antonio M., Kevin Roth, and Andrew Waxman (2017) "Avoiding Traffic Congestion Externalities: The Value of Urgency," Working Paper, USC.

Castillo, Juan Camilo, Dan Knoepfle, and E. Glen Weyl (2017) "Surge Pricing Solves the Wild Goose Chase," Working Paper, Stanford University, July.

Chmura Thorsten and Thomas Pitz (2004) "Minority Game—Experiments and Simulations of Traffic Scenarios," Bonn Econ Discussion Paper 23/2004, Bonn Graduate School of Economics.

Cramton, Peter (2017) "Electricity Market Design," *Oxford Review of Economic Policy*, 33:4, November, 589-612.

Cramton, Peter and Linda Doyle (2017) "Open Access Wireless Markets," *Telecommunications Policy*, 41:5-6, 379-390, June.

Cramton, Peter, R. Richard Geddes and Axel Ockenfels (2018) "Set Road Charges in Real Time to Ease Traffic," *Nature*, 23-25, 2 August.

Cramton, Peter, R. Richard Geddes and Axel Ockenfels (2019) "Using Technology to Eliminate Traffic Congestion," *Journal of Institutional and Theoretical Economics*, 175:1, 126-139.

Duranton, Gilles and Matthew A. Turner (2011) "The Fundamental Law of Road Congestion: Evidence from US Cities," *American Economic Review*, 101:6, 2616-2652.

Federal Highway Administration (2008) Congestion Pricing – A Primer: Overview, US Department of Transportation, April.

Federal Highway Administration (2009) Transit and Congestion Pricing: A Primer, US Department of Transportation, April.

Foreman, Kate (2016) "Crossing the Bridge: The Effects of Time-varying Tolls on Curbing Congestion," *Transportation Research Part A: Policy and Practice*, 92, 76-94.

Friedman, Milton (1982) *Capitalism and Freedom*, Los Angeles: Phoenix Books.

Friedman, Milton and Daniel J. Boorstin (1996) "How to Plan and Pay for the Safe and Adequate Highways We Need," in Roth, Gabriel, *Roads in a Market Economy*, Aldershot: Ashgate Publishing.

Fulton, Lewis, Jacob Mason, and Dominique Meroux (2017) "Three Revolutions in Urban Transportation," Institute of Transportation Studies, University of California, Davis, Research Report UCD-ITS-RR-17-03.

Gabuthy, Yannick, Matthieu Neveu, and Laurent Denant-Boemont (2006), "Structural Model of Peak-period Congestion: An Experimental Study," *The Review of Network Economics*, 5, 273–298.

Geddes, R. Richard (2011) *The Road to Renewal: Private Investment in U.S. Transportation Infrastructure*, Washington, DC: AEI Press.

Gibson, Matthew, and Maria Carnovale (2015) "The Effects of Road Pricing on Driver Behavior and Air Pollution," *Journal of Urban Economics*, 89, 62-73.

Hall, Jonathan D. (2018) "Pareto Improvements from Lexus Lanes: The Effects of Pricing a Portion of the Lanes on Congested Highways," *Journal of Public Economics*, 158, 113-125.

Handy, Susan (2015) "Increasing Highway Capacity Unlikely to Relieve Traffic Congestion," National Center for Sustainable Transportation.

Harstad, Ronald (2008) "William S. Vickrey," *New Dictionary of Scientific Biography Volume 7* (ed. Noretta Koertge), 149-152, Thomson Gale.

Hartmann, John L. (2012), "Special Issue on Transport Infrastructure: A Route Choice Experiment with an Efficient Toll," *Networks and Spatial Economics*, 12, 205–222.

Henzelmann, Torsten (2017) "Urbane Mobilität 2030: Zwischen Anarchie und Hypereffizienz," Roland Berger Consulting.

Irwin, Neil (2017) "Why Surge Prices Make Us So Mad: What Springsteen, Home Depot and a Nobel Winner Know," New York Times. Oct 14.

Kahneman, Daniel, Jack L. Knetsch, Richard H. Thaler (1986) "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *The American Economic Review*, 76:4, 728-741.

Kalmanje, Sukumar and Kara Kockelman (2004) "Credit-Based Congestion Pricing: Travel, Land Value and Welfare Impacts," *Transportation Research Record*, 1864, 45-53.

Karlström, Anders and Joel P. Franklin (2009) "Behavioral Adjustments and Equity effects of Congestion Pricing: Analysis of Morning Commutes during the Stockholm Trial," *Transportation Research Part A: Policy and Practice*, 43, 283-296.

Kreindler, Gabriel E. (2017) "Road Traffic Congestion Charging: Experimental Evidence from Bangalore," Working Paper, MIT, November.

Leape, Jonathan (2006) "The London Congestion Charge," *The Journal of Economic Perspectives*, 20:4, 157-176.

Linn, Joshua, Zhongmin Wang, and Lunyu Xie (2016) "Who will be Affected by a Congestion Pricing Scheme in Beijing?" *Transport Policy*, 47, 34-40.

Lopez, Hector (2017), "The Power of Weak Incentives," Working Paper, University of Maryland.

Land Transport Authority (2016) "Tender Awarded to Develop Next-Generation Electronic Road Pricing System," Singapore: Land Transport Authority, February 25.

Martin, Leslie A. and Sam Thornton (2017) "Can Road Charges Alleviate Congestion?" Working Paper, University of Melbourne, October.

Morris, David Z. (2015) "Oregon Just Debuted This Radical Alternative to the Gasoline Tax." *Fortune*, July 17.

Ostrovsky, Michael and Michael Schwarz (2018) "Carpooling and the Economics of Self-Driving Cars," Working Paper, Stanford University.

O'Connor, Phillip R. and Erin M. O'Connell-Diaz (2015), "Evolution of the Revolution: The Sustained Success of Retail Electricity Competition," *COMPETE Discussion Paper*.

Parkes, David C., Michael O. Rabin, Stuart M. Shieber, Christopher A. Thorpe (2008) Practical Secrecy-preserving, Verifiably Correct and Trustworthy Auctions, *Electronic Commerce Research and Applications* 7:3, November, 294-312.

Parkes, David C., Michael O. Rabin, and Christopher Thorpe (2009) "Cryptographic Combinatorial Clock-Proxy Auctions." *Financial Cryptography and Data Security*, 305–324.

Parkes, David C. Christopher Thorpe, and Wei Li (2015) "Achieving Trust without Disclosure: Dark Pools and a Role for Secrecy-Preserving Verification," *Proceedings of the Third Conference on Auctions, Market Mechanisms and Their Applications*.

Parry, Ian W.H. and Antonio Bento (2001) "Revenue Recycling and the Welfare Effects of Road Pricing," *Scandinavian Journal of Economics*, 103:4, 645-671.

Pishue, Bob (2017) *US Traffic Hotspots*, INRIX Research, September.

Poole, Robert W., Jr. (2014) "Value-Added Tolling: A Better Deal for American Highway Users," *Reason Foundation Policy Brief* 116.

Poole, Robert W., Jr. (2018) *Rethinking America's Highways: A 21st-Century Vision for Better Infrastructure*, University of Chicago Press.

Rapoport Amnon, William E. Stein, James E. Parco, and Darryl A. Seale (2004) "Equilibrium Play in Single-server Queues with Endogenously Determined Arrival Times," *Journal of Economic Behavior and Organization*, 55, 67–91.

Salzberg, Andrew (2017) "Road Pricing: A Solution to Gridlock," Medium, Uber Under the Hood.

SBD (2016) "How Autonomous Vehicles Could Relieve or Worsen Traffic Congestion," Berlin, Germany: HERE.

Schaller Consulting (2017) "Unsustainable? The Growth of App-Based Ride Services and Traffic, Travel and the Future of New York City," Brooklyn, NY, February.

Schneider Kerstin and Joachim Weimann (2004) "Against all Odds: Nash Equilibria in a Road Pricing Experiment," in Michael Schreckenberg and Reinhard Selten (eds.) *Human Behaviour and Traffic Networks*. Springer Publishing Co, Berlin, Heidelberg, New York.

Schrank, David, Bill Eisele, Tim Lomax, and Jim Bak (2015) *Urban Mobility Scorecard*, Texas Transportation Institute, August.

Schweitzer, L.A.C. and Brian D. Taylor (2008) "Just Pricing: The Distributional Effects of Congestion Pricing and Sales Taxes," *Transportation*, 35:6, 797–812.

Selten Reinhard, Thorsten Chmura, Thomas Pitz, Sebastian Kube, and Michael Schreckenberg (2007) "Commuters Route Choice Behaviour," *Games and Economic Behavior*, 58, 394–406.

Small, Kenneth A. (1992) "Using the Revenues from Congestion Pricing," *Transportation*, 19, 359-381.

Small, Kenneth A., Cliord Winston, and Jia Yan (2005) "Uncovering the Distribution of Motorists' Preferences for Travel time and Reliability," *Econometrica*, 73:4, 1367-1382.

The Economist (2014) "The Cost of Traffic Jams," *The Economist*, November 3.

Thorpe, Christopher and David C. Parkes (2007) "Cryptographic Securities Exchanges," *Proc. 11th International Conference on Financial Cryptography and Data Security*, 163–178.

Thorpe, Christopher and David C. Parkes (2009) "Cryptographic Combinatorial Securities Exchanges," *Financial Cryptography and Data Security*, 285–304.

Vickrey, William S. (1963) "Pricing in Urban and Suburban Transport," *The American Economic Review*, 53:2, 452-465.

West, Sarah E. and Roberton C. Williams (2004) "Estimates from a Consumer Demand System: Implications for the Incidence of Environmental Taxes," *Journal of Environmental Economics and Management*, 47:3, 535-558.

Zimmer, John and Logan Green (2017) "The End of Traffic – Increasing American Prosperity and Quality of Life," Medium.