Congestion pricing, carpooling, and commuter welfare

Michael Ostrovsky
ostrovsky@stanford.edu

Michael Schwarz*
mschwarz@microsoft.com

September 7, 2023

Preliminary and incomplete. Please do not circulate.

Abstract

Building on the canonical "bottleneck" model of Vickrey (1969), we show that carpooling and road pricing are highly complementary in addressing traffic congestion: they can be much more effective jointly than each one separately, and can improve commuter welfare without having to rely on the redistribution of government revenue. By contrast, technological advances that make time in traffic more comfortable or productive (e.g., self-driving cars), implemented without additional economic incentives, may result in zero improvement in social welfare.

*Ostrovsky: Stanford University and NBER. Schwarz: Microsoft. Ostrovsky thanks the National Science Foundation for financial support (grant #1824317).

1 Introduction

Heavy traffic congestion is a widespread, major, and persistent problem, affecting the richest and the poorest countries alike. In addition to the immediate problems caused by traffic congestion (delays, time in traffic, stress, wasted energy and resources, environmental externalities), it contributes to such seemingly disparate issues as earnings inequality (as women are less likely than men to spend large amounts of time commuting to work (Le Barbanchon et al., 2021)) and the community opposition to building additional housing (Ding and Taylor, 2022).

Congestion pricing has been advocated by economists as a potential solution to traffic congestion for more than a century, going back to Pigou (1920) and Vickrey (1969). The typical main arguments in favor of congestion pricing are expressed, e.g., in an open editorial by Yale economists making the case for congestion pricing in Connecticut: "By charging people higher tolls during rush hour, tolls can operate as a congestion tax that channels people with more flexible schedules to drive at different times" and "The toll can also encourage people to take different modes of transportation such as the train and public transit, thereby creating cleaner air and benefiting the environment"(Ayres et al., 2019). These arguments are, of course, correct. However, in this paper, we argue that the economic analysis of congestion pricing needs to consider an additional major force that has traditionally been overlooked: carpooling.[1]

The most immediate way to see why considering carpooling should be of first-order importance is to note that the majority of road capacity is currently wasted: most commuters drive solo, leaving multiple seats in their cars empty—despite having many potential carpooling partners who travel along similar routes at the same time (which is what creates traffic congestion in the first place). If all (or most) commuters carpooled, road throughput (measured in commuters, not cars) would automatically increase by a substantial amount. But in the absence of congestion pricing, individual commuters do not have an incentive to do so. In a sense, we are in a typical Prisoners' Dilemma: commuters would be better off if everyone carpooled, but regardless of what others do, an individual commuter prefers to drive solo.[2] What congestion pricing can do is "fix" this Prisoners' Dilemma by giving commuters individual incentives to carpool.

Conversely, carpooling can "fix" a major problem that congestion pricing faces if all commuters drive solo. When roads are congested, the time wasted in traffic serves as a non-monetary "price" that commuters pay to get to work (for concreteness, following Vickrey (1969), we will focus on the morning rush hour), clearing the underlying demand and supply (i.e., road capacity). When tolls are introduced, to clear the same underlying demand and supply, they need to be of comparable magnitude to the monetary value of time wasted in traffic before tolls (the canonical "bottleneck" model of Vickrey (1969), summarized in Sections 2, 3, and 6, makes these statements precise). Thus,

[1] E.g., in the 900-page Handbook of Transportation Economics (de Palma et al., 2011), "carpooling" does not appear in the index (by contrast, there are 33 entries in the index for the word "parking" and its variations: "parking economics", "parking policy," and so on). Likewise, "carpooling" does not appear in the recent textbook on the mathematics and economics of road pricing (Yang and Huang, 2005), even though it contains 18 pages of references.

[2] Of course, this description is an exaggeration, but not a big one: during a typical morning or evening rush hour, the fraction of commuters carpooling with others, relative to the total number of commuters, is vanishingly small.

none of the commuters are directly better off—for them, the time cost of being stuck in traffic is simply replaced by an equivalent monetary cost. Even this outcome is in a sense the best case scenario, not taking into account such additional complications as heterogeneity among commuters in terms of their flexibility and value of time and the ability of policymakers to calculate appropriate tolls in the face of uncertainty. Of course, with tolls, the government does collect revenue, and so in principle redistributing this revenue should make it possible to make the commuters better off, but in practice these types of redistribution policies are often very challenging politically, both because it is often unclear how to target these redistribution policies (without distorting incentives) and because they require a high degree of trust from the public. As just one illustration (out of many), the state of New Jersey is currently suing the U.S. Department of Transportation and the Federal Highway Administration to block "the ill-conceived congestion pricing plan put forward by the Metropolitan Transportation Authority (MTA) and New York City and State agencies."[3] As the press release states, "Since New York and the MTA first revealed their proposed congestion pricing scheme, [New Jersey] Governor Murphy has remained a staunch advocate for the New Jersey commuters, transportation agencies, businesses, and residents who would suffer as a result of this unreasonable and unprecedented proposal."

As our paper shows, carpooling can fix this problem. We augment Vickrey's bottleneck model with the possibility of carpooling, which incurs disutility Δ. We then show that with optimal tolls, commuters are better off in the world with congestion pricing than in the world without it. Intuitively, at peak times, instead of serving as a lever to push commuters away from traveling on a highly demanded road, the toll now serves as an incentive to carpool—because by carpooling with another person, a commuter would only have to pay a half of the toll. This incentive increases the throughput of the road (as measured in commuters, not in cars), thus effectively increasing supply and making the demand side better off under market clearing prices.

In effect, taken together, tolls and carpooling operate as a "shock absorber" for times of peak travel demand: on average, there will be more riders in each car during times of peak demand than during off-peak hours. Note that both technologies are needed for this shock absorber to work well, and they reinforce each other. The presence of tolls makes carpooling more attractive, since it allows riders to share the costs of those tolls and thus pay less and travel closer to their preferred times—in contrast to the situation with congestion delays and without tolls, in which carpooling does not reduce any rider's cost of waiting in traffic. A cost paid in the form of dollars can be shared among the carpooling riders and thus substantially reduced for each one of them individually, while a cost paid in the form of time spent in traffic cannot: each of the carpooling riders "pays" the full amount. Conversely, carpooling makes tolls more attractive politically, because it gives price-sensitive riders a feasible way to commute during their preferred times, and can make them substantially better off even if the revenue from tolls does not benefit them directly.

While for simplicity, our model only considers the case of two potential carpoolers per car, the insight is more general. Most immediately, there can of course be more than two commuters in

[3] `https://www.nj.gov/governor/news/news/562023/approved/20230721a.shtml`.

every cars, with road capacity expanding even further during the most demanded hours. Perhaps more surprisingly, suitable congestion pricing with high tolls during peak times can lead to the rise of "intermediate public transport"—the mode of transportation under which a for-hire worker drives a handful of passengers (typically a car or a minivan), either on a fixed or a flexible route. This mode of transportation is common in low and middle-income countries (see, e.g., Jaiswal et al. (2022) for a survey), where many commuters cannot afford a car or a single-passenger taxi and may not have convenient options on (large-scale, fixed-route) public transportation. At the same time, it is rare in rich countries, especially for commuting to and from work, because most workers can afford their own cars and, in the absence of congestion pricing, prefer to drive solo. With appropriate congestion pricing, such modes of transportation will arise in developed countries as well, substantially raising the throughput of the roads (measured, crucially, in people rather than car, which is the appropriate metric).

Before proceeding to the main results of the paper, in Section 4 we address an argument that is sometimes made, that perhaps we do not need to worry about traffic congestion too much due to improving autonomous driving technology. With the rise of self-driving cars, the disutility of being stuck in traffic will be substantially reduced (and commuters will be able to sleep in such cars, work, etc.), and so the argument is then made that therefore the problem of traffic jams will be lessened. It's OK to be stuck in traffic if you can sleep in the car. In Section 4, we show that this argument is fundamentally flawed. The reason for that is that as the *per-minute* disutility of being stuck in traffic drops, commuters are more willing to be stuck in traffic (and less willing to pay to avoid it, e.g., in the form of arriving at work earlier or later), and thus the overall amount of traffic increases, causing the corresponding increases in delays. As a result, while the per-minute disutility goes down, the *overall* disutility of being stuck in traffic remains the same. Endogenous equilibrium behavior by the drivers completely undoes the benefits of the technological improvements, at least as it relates to the issue of welfare losses from traffic jams. This observation further emphasizes the need for intelligent congestion pricing and transportation policies even in the futuristic world with autonomous transportation.

1.1 Related Literature

TBA - INCOMPLETE

Empirical evidence:

- Holland and Watson (1977): Singapore experience after introducing cordon pricing in the 1970s. "The number of cars entering the Restricted Zone between 7:30 and 10:15 fell by 73 per cent ... The volume of cars entering during the half hour before 7:30 a.m. rose by 23 per cent ... The absolute number of car pools entering during the 7:30 - 10:15 period increased by about 60 per cent... The proportion of car pools thus jumped from less than 7 to 37 per cent for those hours."

- Willson and Shoup (1990): Removing parking subsidies by an employer in Los Angeles has led to a large increase in carpooling.

Warner Center (suburban)	Parking Cost/Mo.	\$0	\$30	+\$30
Los Angeles, CA	Solo Driver	90%	46%	−49%
Large private firm	Carpool/Vanpool	6%	48%	+700%
	Transit	0%	0%	0%

2 Model

Our model is essentially identical to the canonical bottleneck model of Vickrey (1969), with three modifications. Two of the modifications are minor. First, to adjust for inflation, we multiply all the monetary amounts in the model by a factor of 10.[4] Second, while Vickrey fixes the number of commuters but allows the capacity of the road to vary, we fix both numbers, which is sufficient for our purposes. The results and derivations in Sections 3 and 6 below are taken directly from Vickrey (1969), subject to these two modifications. The third modification, which we will discuss later on, is substantive: we will allow for the possibility of carpooling, and consequently will introduce a parameter denoting the commuters' disutility from carpooling relative to driving solo. We maintain Vickrey's notation whenever possible.

As in Vickrey (1969), there are $N = 7200$ travelers crossing a particular bottleneck during their morning work commute. Their ideal times for passing the bottleneck are distributed uniformly between $t_a =$ 8:00 AM and $t_b =$ 9:00 AM. Thus, if the bottleneck did not have capacity constraints, $v_m = 7200/60 = 120$ commuters would pass through the bottleneck every minute during the time interval $[t_a, t_b]$. However, there *is* a capacity constraint: only $v = 60$ cars per minute can pass the bottleneck.

Hence (in the absence of carpooling, as in Vickrey's model), it is impossible for all commuters to arrive at work at their ideal arrival times: at least some will have to arrive early or late. Also, a queue of cars waiting to pass the bottleneck will form. To determine the equilibrium behavior of the commuters and the overall pattern of traffic, we need to make assumptions on commuters' preferences, which, like the rest of the model, will be kept very simple.

Each commuter values time spent at home at $w_h = 20$ cents per minute, and values time spent waiting in the queue at $w_q = 0$. The value of time in the office is high after the ideal arrival time, when it is equal to $w_j = 40$ cents per minute, and relatively low before the ideal arrival time, when it is equal to $w_p = 10$ cents per minute. For simplicity, we assume that it does not take any time for commuters to get from their homes to the bottleneck, and likewise it does not take any time for them to get from the bottleneck to their jobs.

[4] According the Bureau of Labor Statistics, the exact CPI adjustment from January 1969 until July 2023 is 8.59 (`https://data.bls.gov/cgi-bin/cpicalc.pl?cost1=1.00&year1=196901&year2=202307`).

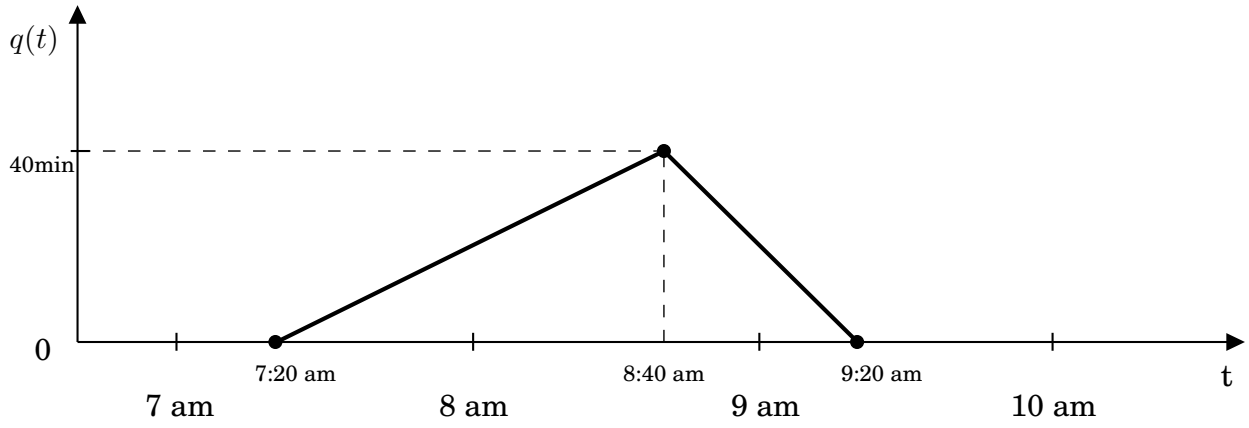


Figure 1: Pattern of traffic with no carpooling or congestion pricing

3 Equilibrium with no carpooling and no congestion pricing

We can now describe the equilibrium of the basic Vickrey model with no carpooling and no road pricing given in Section 2. To do so, it is convenient to consider a function, $q(t)$, that denotes the amount of time (in minutes) that a commuter will need to spend waiting in a queue if he wants to *leave* the bottleneck at time t. (Thus, he will need to arrive there and start queuing at time $t - q(t)$.)

Consider now a commuter who arrives at work later than his optimal time. If he shifts his arrival time by a small amount ε (which could be positive, denoting a later arrival time, or negative, denoting an earlier arrival time), his value of time at work is reduced by εw_j, his value of time in the queue is increased by $\varepsilon q'(t) w_q + o(\varepsilon)$, and his value of time at home is increased by $\varepsilon(1 - q'(t))w_h + o(\varepsilon)$. In equilibrium, every commuter is optimizing, and so for a commuter who arrives at work late and leaves the bottleneck at time t, it has to be the case that $-w_j + q'(t)w_q + (1 - q'(t))w_h = 0$. Thus, in this case, we have

$$q'(t) = \frac{w_h - w_j}{w_h - w_q} = -1. \tag{1}$$

Similarly, for a commuter who arrives at work earlier than her optimal time and passes the bottleneck at time t, we have

$$q'(t) = \frac{w_h - w_p}{w_h - w_q} = 0.5. \tag{2}$$

These equations allow us to characterize the overall shape of function $q(t)$, illustrated in Figure 1, and the corresponding pattern of traffic. Traffic starts building up at 7:20 AM, reaches its peak for commuters leaving the bottleneck at 8:40 AM, and then starts going down, until 9:20 AM, when it ends. Commuters leaving the bottleneck at 8:40 AM get to work at their ideal time, but also spend the longest in traffic: q(8:40 AM) = 40 minutes. A third of commuters (those leaving the bottleneck after 8:40 AM) will be late for work, and two thirds (those leaving the bottleneck before 8:40 AM) will be early.

The total amount of time spent in the queue by all commuters is 144,000 minutes (20 minutes

per commuter). The total delay relative to the ideal arrival time (for those who arrive at work late) is 24,000 minutes (10 minutes per commuter who is late). Finally, for the commuters who arrive earlier than the ideal arrival time, the average amount of time by which they are early is 20 minutes, for a total of 96,000 minutes.

In monetary terms, this corresponds to a disutility of \$28,800 from waiting in queue (144,000 minutes at the disutility of $w_h - w_q = 20$ cents) and \$14,400 from displaced arrival times (24,000 minutes at the disutility of $w_j - w_h = 20$ cents and 96,000 minutes at the disutility of $w_h - w_p = 10$ cents). The disutility from the commute is the worst for commuters whose ideal time for passing the bottleneck is 8:40 AM. Their disutility is \$8. The disutility is the lowest for those whose ideal times for passing the bottleneck are the earliest possible or the latest possible (8:00 AM or 9:00 AM). Their disutility is \$4.

4 Reducing the disutility of being stuck in traffic

A major component of the overall social loss discussed above is the disutility of waiting in the queue: the value of time in traffic ($w_q = 0$) is lower than all other potential uses of time (w_h, w_p, and w_j). A natural inclination is then to say that if technological progress could make time in traffic more comfortable or productive (by, e.g., introducing air conditioning and music players as we have done in the past, or developing self-driving technology as we are hoping to do in the future), the social loss from traffic congestion will be reduced.[5] Unfortunately, as we illustrate below with a simple modification of the baseline model, this hope may be misplaced.

Consider a modification of the model of Section 3 with one change: instead of having $w_q = 0$, suppose time in traffic becomes more valuable – say, $w_q = 4$¢.[6] From equations 1 and 2, we then get

$$q'(t) = \frac{20 - 40}{20 - 4} = -\frac{5}{4}$$

for commuters arriving at work late, and

$$q'(t) = \frac{20 - 10}{20 - 4} = \frac{5}{8}$$

for commuters arriving at work early.

The resulting "delay" function $q(t)$ is illustrated in Figure 2. It is qualitatively similar to that of Figure 1 (which is shown in Figure 2 with a dotted line), but with one important quantitative difference. As before, traffic starts building up at 7:20 AM, reaches its peak for commuters leaving

[5] E.g., here is one expression of this common sentiment: "When we reach the point where human intervention behind the wheel is no longer needed, autonomous vehicles will drastically improve the daily commute. Imagine, instead of sitting behind the wheel, you'll be able to stretch out in the back, get ahead on some work, or simply relax and catch up on your latest Netflix obsession. When you think that Americans spend 19 full working days a year stuck in traffic on their commute, that's an awful lot of time commuters will be able to claw back for themselves." (`https://www.forbes.com/sites/bernardmarr/2020/07/17/5-ways-self-driving-cars-could-make-our-world-and-our-lives-better`)

[6] The conclusions would be the same for any value of w_q less than 10¢.

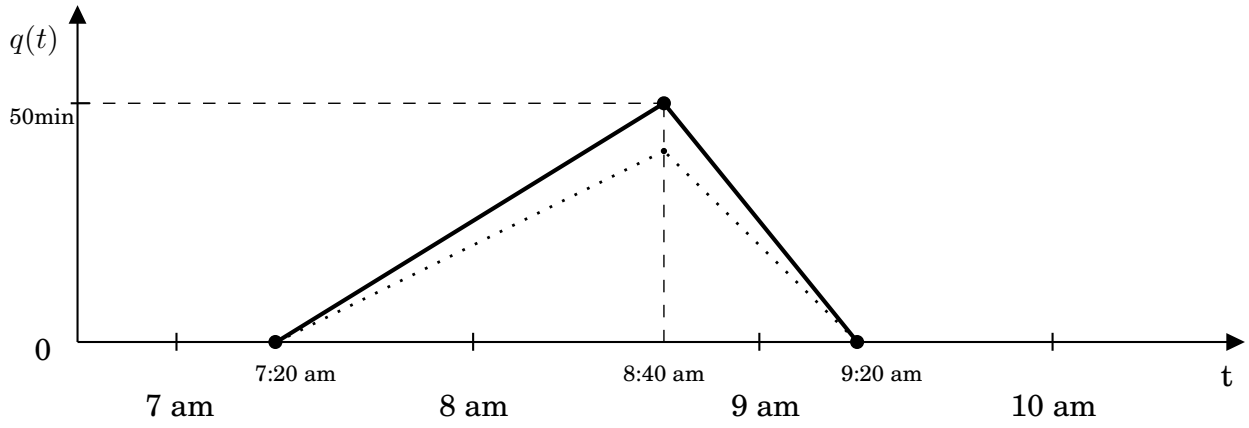


Figure 2: Pattern of traffic with reduced disutility of time in traffic

the bottleneck at 8:40 AM, and then starts going down, until 9:20 AM, when it ends. Commuters leaving the bottleneck at 8:40 AM get to work at their ideal time, but now spend $q(8{:}40\text{ AM}) = 50$ minutes in traffic (instead of 40).

The total amount of time spent in the queue by all commuters is now 180,000 minutes (25 minutes per commuter). The displacement relative to the ideal arrival times is unchanged (both for those who arrive at work late and for those who arrive early).

In monetary terms, the 180,000 minutes corresponds to a disutility of \$28,800 ($w_h - w_q = 16¢$ per minute)—exactly the same as before. In other words, while the disutility per minute of queuing is traffic got reduced by a factor of 5/4, the equilibrium amount of time spent queuing in traffic goes up by the same multiple, thus completely undoing the benefits of higher comfort and productivity.

5 Equilibrium with carpooling and without congestion pricing

We now introduce the possibility of carpooling. In the spirit of keeping the model simple, we will limit carpooling to two people per car, and will not make a distinction between the passenger and the driver. We assume that carpooling imposes on each of the carpooling partners a disutility $\Delta > 0$, which enters their payoff functions additively. This disutility captures such issues as the loss of flexibility relative to driving solo, detour time for the driver, waiting time for the passenger, etc. For simplicity, we assume that the marginal cost of driving solo is also zero, although the results remain unchanged as long as the cost of driving is less than the disutility from carpooling (which of course is largely the case in practice, as the observed prevalence of carpooling is very low).

For now, we continue to assume that roads are free. In this case, the possibility of carpooling does not help: every driver continues to drive solo, and the equilibrium is identical to that of Section 3. Of course, if Δ is small, and everyone carpooled, that would have resulted in a much better outcome for everyone involved—but that is not an equilibrium. An individual commuter's best response is to drive solo, regardless of what others are doing. In a sense, this is a classic case of Prisoners' Dilemma.

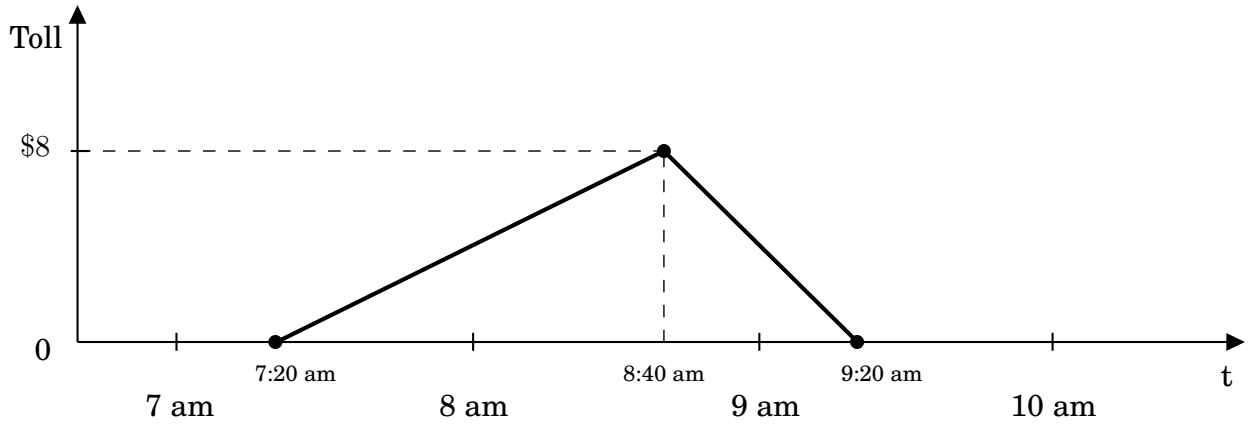


Figure 3: Optimal congestion pricing without carpooling.

6 Equilibrium with congestion pricing and without carpooling

We now consider the policy response analyzed in Vickrey (1969): road pricing, without considering the possibility of carpooling. The socially efficient outcome in this case is to have all drivers pass through the bottleneck starting at 7:20 AM and ending at 9:20 AM, at a uniform rate of 60 cars per minute, with no queuing. The price schedule that supports this outcome is illustrated in Figure 3. The price for passing the bottleneck starts at 0 at 7:20 AM, rises linearly to $8 at 8:40 AM, and then drops linearly to 0 at 9:20 AM.[7]

This outcome eliminates traffic congestion, but still has serious shortcomings. First, since all commuters arrive at work at the same times as in the original "congested" base case of Section 3, the substantial social costs of suboptimal arrival times are not eliminated. Second, none of the commuters are better off after the introduction of congestion pricing: each of them simply pays in dollars what he or she used to pay in time spent queuing in a traffic jam! As we discuss in the Introduction, this presents a major political problem, leading to a substantial fraction of the population opposing such congestion pricing schemes—and that is precisely what has happened worldwide during the 50 years since Vickrey's paper was published (and the 100 years since Pigou's book), with heavily congested cities vastly outnumbering those rare exceptions that managed to implement time-sensitive congestion pricing.

7 Equilibria with congestion pricing and carpooling

We now get to the main part of the paper—analysis of outcomes under congestion pricing, taking into account the possibility of carpooling. We start out by describing the outcome under

[7]Consider a commuter who arrives at work late. If this commuter chose to arrive one minute earlier, then under this price schedule, he would have to pay 20 cents more in tolls, would spend one minute less at home (a loss of 20 cents), and would spend one minute more at work (a gain of 40 cents), for no net benefit. Alternatively, consider a commuter who arrives at work early. If this commuter chose to arrive one minute later, then she would have to pay 10 cents more in tolls, would spend one minute more at home (a gain of 20 cents), and would spend one minute less at work (a loss of 10 cents), again for no net benefit.

a "simple" policy, when the analysis is particularly straightforward and transparent. We then describe the socially optimal policy and the corresponding equilibrium outcome.

7.1 Simple policy

The simple policy is to charge a constant toll of 2Δ at all times (or a slightly higher toll, if one wants to make individual incentives strict). Then in equilibrium, all commuters will carpool, all will arrive at work at their ideal time, and there will be no traffic congestion. In effect, while the capacity of the road measured in "cars" remains constant at 60 per minute, its capacity measured in "commuters" doubles—without adding more lanes. The social loss relative to the ideal outcome of solo driving and no capacity constraints is 7200Δ—each commuter incurs a disutility Δ from carpooling. Each commuter also pays another Δ as his or her share of the toll.

7.2 Socially optimal policy

The policy described above is simple and transparent—but is not in fact socially optimal. The reason for this is the fact that it forces commuters to carpool "too much."

E.g., in the outcome of the simple policy, consider a small number $\varphi > 0$ of commuters passing the bottleneck right before 9:00 AM. Each one of them incurs a disutility of Δ from carpooling, and so the total disutility they incur is $\varphi\Delta$. Suppose now that instead of having them carpool over the time interval [9:00 AM $- \varphi/120$, 9:00 AM], we have them drive solo, over the time interval [9:00 AM $- \varphi/120$, 9:00 AM $+ \varphi/120$] (with no traffic congestion). On average, each of these commuters is delayed by $\varphi/240$ minutes, at the cost of $w_j - w_h = 20$ cents per minute. Thus, their total cost of delay is $\varphi^2/12$ cents, which is less than or equal to $\varphi\Delta$ for all $\varphi \leq 12\Delta$ (where Δ is measured in cents). The optimal φ maximizes $\varphi\Delta - \varphi^2/12$, and is thus equal to 6Δ.

Consider now a small number $\psi > 0$ of commuters passing the bottleneck right after 8:00 AM under the simple policy. A similar analysis shows that for any $\psi \leq 24\Delta$, it is better to have these commuters drive solo over the time interval [8:00 AM$-\psi/120$, 8:00 AM$+\psi/120$] than to have them carpool over the time interval [8:00 AM, 8:00 AM $+ \psi/120$], and the optimal ψ is equal to 12Δ.

These two changes give us the socially optimal commuting pattern:

- solo driving from (8:00 AM $- \Delta/10$) until (8:00 AM $+ \Delta/10$),
- carpooling from (8:00 AM $+ \Delta/10$) until (9:00 AM $- \Delta/20$), and
- solo driving from (9:00 AM $- \Delta/20$) until (9:00 AM $+ \Delta/20$).

The price schedule that implements the socially optimal outcome is shown with solid lines in Figure 4. (The dotted line reproduces Vickrey's price schedule without carpooling, from Figure 3.) In the figure, we set a commuter's disutility of carpooling to $\Delta = \$2$, corresponding to 10 minutes of being stuck in traffic vs. spending that time at home, or 20 minutes of being stuck in traffic vs. spending that time at work before the ideal arrival time. Under the optimal price schedule, the toll rises linearly from zero at 7:40 AM to \$4 at 8:20 AM, stays at \$4 from 8:20 AM until 8:50 AM

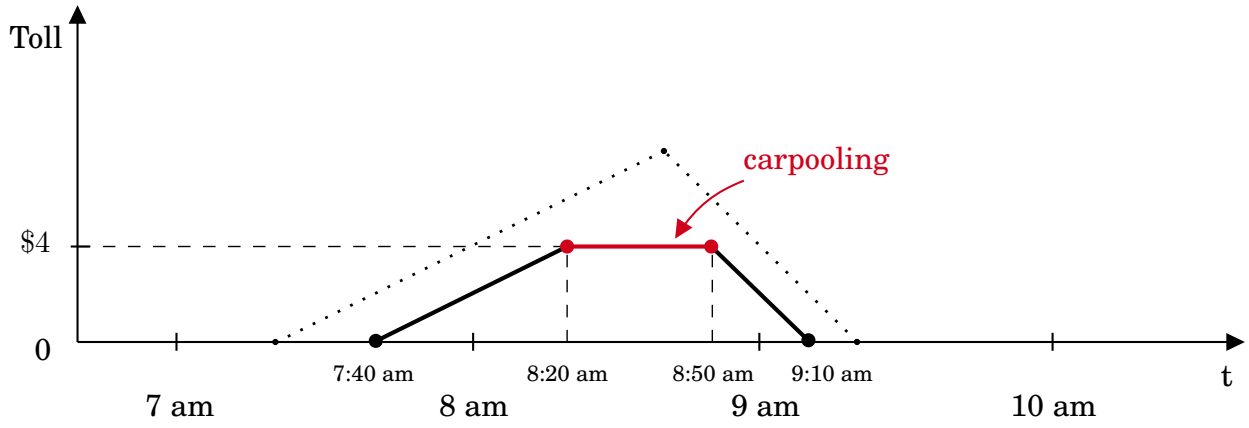


Figure 4: Optimal congestion pricing with carpooling

(which is the time interval during which commuters carpool), and then drops linearly from $4 at 8:50 AM to zero at 9:10 AM. Note that for those who carpool (i.e., those who travel during peak times between 8:20 AM and 8:50AM), the outcome is identical to that under the simple policy: they arrive at work at their ideal time, incur a disutility of $2 from carpooling, and pay $2 as their share of the toll, for the total cost of $4. For those who drive solo, the outcome is strictly better than under the simple policy: their total cost decreases as they get further away from peak times, to $2 (i.e., half as much as under the simple policy) for those whose ideal arrival times are 8:00 AM or 9:00 AM. For all commuters, the outcome is substantially better than that under the Vickrey price schedule without carpooling.

8 Comparisons

We now summarize and compare the total social costs, and the costs facing the commuters (subtracting from the total social costs the tolls the commuters have to pay) under the various policies considered above. As in Vickrey (1969), we compute these costs relative to the benchmark case of unlimited bottleneck capacity, in which case there are no costs of displaced arrival times or waiting in queue, and thus the overall cost is normalized to zero.

The costs are summarized in Table 1. Under the baseline model with no road pricing and no carpooling, there are two sources of disutility: displaced arrival times (some commuters arrive at work before their ideal times, and some arrive after) and waiting in the queue (being stuck in traffic). With 7200 commuters, the average disutility from displaced arrival times is $2 per commuter, and the average disutility from being stuck in traffic is $4 per commuter, for a total per-commuter disutility of $6.

The numbers remain unchanged when the possibility of carpooling is introduced but roads are unpriced, for any $\Delta > 0$. As discussed in Section 5, we are essentially in a Prisoner's Dilemma situation: it would be better for everyone if everyone carpooled, but holding the behavior of others fixed, it is optimal for each individual commuter to drive solo. Hence, nobody carpools.

	Cost ($/morning)						
	Displaced arrival	Waiting in queue	Disutility from carpooling	Total Social Cost	Total Commuter Cost	Social Cost per commuter	Average Commuter Cost
No road pricing, no carpooling	14,400	28,800	-	43,200	43,200	6.000	6.00
No road pricing, carpooling, Delta > 0	14,400	28,800	-	43,200	43,200	6.000	6.00
Road pricing, no carpooling	14,400	-	-	14,400	43,200	2.000	6.00
Simple road pricing, carpooling, Delta = $3	-	-	21,600	21,600	43,200	3.000	6.00
Simple road pricing, carpooling, Delta = $2	-	-	14,400	14,400	28,800	2.000	4.00
Simple road pricing, carpooling, Delta = $1	-	-	7,200	7,200	14,400	1.000	2.00
Optimal road pricing, carpooling, Delta = $3	8,100	-	5,400	13,500	27,000	1.875	3.75
Optimal road pricing, carpooling, Delta = $2	3,600	-	7,200	10,800	21,600	1.500	3.00
Optimal road pricing, carpooling, Delta = $1	900	-	5,400	6,300	12,600	0.875	1.75

Table 1: Social costs

Road pricing (on its own, without carpooling), does of course affect the outcome in a subtle way. Waiting in traffic is eliminated, while the disutility from displaced arrival times remains the same. So the average social cost per commuter is now $2. However, that is not the direct impact on commuters' themselves. For them, the cost of waiting in the queue has been replaced by an equivalent monetary cost of tolls, and their welfare is unchanged: the average welfare loss for the average commuter is still $6. [TBA: add all this crap about politics]

9 Conclusion

In this paper, we show the joint power of road pricing and carpooling. We also show the importance of taking carpooling into account when evaluating the effects of congestion pricing proposals, as well as in determining optimal tolls. The mathematical arguments in our paper are deliberately simple (and are only a small deviation from Vickrey's canonical model), but the resulting economic outcomes, and the forces leading to the first-order improvement, are critically different once carpooling is considered. Intuitively, in Vickrey's model with solo driving, time in traffic serves as a *non-monetary* "price schedule" equilibrating demand and road capacity, and when tolls are introduced, the *monetary* price schedule equilibrating the same demand and the same road capacity ends up being equivalent to the non-monetary one, with the net effect on commuters being zero: for them, the time being stuck in traffic is simply replaced with the equivalent monetary amount in the form of tolls.[8] By contrast, *with* carpooling, the outcomes are fundamentally different from the point of view of commuters. When the "price schedule" is non-monetary and is "charged" in the form of time wasted in traffic, carpooling would not change the per-commuter price - each of the carpooling partners would still have to incur the full price. When, however, the price schedule is monetary and is charged in the form of dollars, carpooling reduces the price paid by each commuter

[8] Of course, the overall social outcome is improved under congestion pricing, but as we discussed in the Introduction, the fact that the direct impact on commuters is zero leads to very serious distributional and political problems severely affecting congestion pricing initiatives.

in half, thus creating incentives to carpool and leading to a first-order improvement to the overall commuter welfare even if the revenue from tolls is destroyed. Carpoolers can split the monetary cost, but not the cost incurred in the form of wasted time.

References

Ayres, I., S. Berry, K. Gillingham, and B. Nalebuff (2019). How we learned to love the toll. *Hartford Courant*.

de Palma, A., R. Lindsey, E. Quinet, and R. Vickerman (Eds.) (2011). *A Handbook of Transportation Economics*. Elsevier.

Ding, H. and B. D. Taylor (2022). Traffic trumps all: Examining the effect of traffic impact analyses on urban housing. *Journal of Planning Literature 37*(1), 3–16.

Holland, E. P. and P. L. Watson (1977). Measuring the impacts of singapore's area license scheme. In *Transport decisions in an age of uncertainty: Proceedings of the third World Conference on Transport Research Rotterdam, 26–28 April, 1977*, pp. 417–422. Springer.

Jaiswal, A., M. Manoj, and G. Tiwari (2022). A review of studies on service quality of intermediate public transport. *IATSS Research 46*(4), 537–546.

Le Barbanchon, T., R. Rathelot, and A. Roulet (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics 136*(1), 381–426.

Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan.

Vickrey, W. S. (1969). Congestion theory and transport investment. *American Economic Review 59*(2), 251–260.

Willson, R. W. and D. C. Shoup (1990). Parking subsidies and travel choices: Assessing the evidence. *Transportation 17*, 141–157.

Yang, H. and H.-J. Huang (2005). *Mathematical and Economic Theory of Road Pricing*. Elsevier.