NBER WORKING PAPER SERIES

CONGESTION PRICING, CARPOOLING, AND COMMUTER WELFARE

Michael Ostrovsky
Michael Schwarz

Working Paper 34261
http://www.nber.org/papers/w34261

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2025

We are grateful to Amirreza Bagheridelouee and Suraj Malladi for research assistance. Ostrovsky thanks the National Science Foundation (grant #1824317) and the Business, Government & Society Initiative at the Stanford Graduate School of Business for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Michael Ostrovsky and Michael Schwarz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Congestion Pricing, Carpooling, and Commuter Welfare
Michael Ostrovsky and Michael Schwarz
NBER Working Paper No. 34261
September 2025
JEL No. D47, D62, H23, L98, Q58, R41, R48

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.

Michael Ostrovsky
Stanford University
Graduate School of Business
and NBER
ostrovsky@stanford.edu

Michael Schwarz
Microsoft
mschwarz@microsoft.com

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 and heatlh externalities), it also 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 community opposition to building additional housing (Ding and Taylor, 2022).

Congestion pricing has been championed by economists as a potential solution to traffic congestion for more than a century, going back to Pigou (1920) and Vickrey (1969). The typical 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"

[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 major 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.

model of Vickrey (1969), summarized in Sections 2, 3, and 6 below, makes these statements precise). Thus, *none of the commuters are directly better off when congestion pricing is introduced*—for them, the time cost of being stuck in traffic is simply replaced by an equivalent monetary cost. 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.[3] As just one illustration (out of many), the state of New Jersey has sued 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."[4] 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." Kreindler (2024) obtains empirical estimates of the parameters of the Vickrey (1969) based on a careful field experiment in Bangalore and concludes that "[the] results suggest limited commuter welfare benefits from peak-spreading traffic policies."

As our paper shows, carpooling can alleviate this shortcoming of congestion pricing. We augment Vickrey's bottleneck model with the possibility of carpooling, which incurs a small disutility Δ. We then show that with optimal tolls, *all* commuters are *strictly* better off in the world with congestion pricing than in the world without it, even if all revenue from tolls is destroyed. Intuitively, at peak times, instead of serving as a lever to push commuters away from traveling on a highly demanded road ("peak-spreading"), the toll now serves as an incentive to carpool—because by carpooling with one other person, a commuter would only have to pay one half of the toll instead of the entire amount. This incentive increases the throughput of the road (as measured in commuters, not in cars), thus effectively increasing the available supply and thus making the demand side better off under market-clearing prices.[5]

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

[3] Hall (2018) discusses the difficulties of allocating toll revenue in ways that make all drivers better off, and also raises the question of whether tolls can make all drivers better off even if the revenue from tolls is destroyed. He shows that putting appropriate tolls on some of the lanes on a congested highway can lead to a Pareto improvement, by relieving hypercongestion on those lanes and thereby increasing those lanes' throughput—which would in turn lead to fewer cars and a lower congestion level on the remaining free lanes. However, as Hall (2018) notes, "Obtaining a Pareto improvement [...] comes at a cost. By only pricing a portion of the lanes, we leave the other lanes congested, with all the resulting social costs. That said, inasmuch as generating a Pareto improvement makes it politically feasible to implement tolling, then we are trading potential, but unrealized, welfare gains for actual welfare gains." Our results show that with carpooling, one may not need to face this trade-off: it may be possible to both obtain the first-best outcome *and* generate Pareto improvement for all commuters (relative to the congested status quo), even if all revenue from tolls is destroyed.

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

[5] A related issue of HOV lanes and tolling is explored in Huang et al. (2000), Konishi and Mun (2010), and Zhang et al. (2024). See Section 1 of Zhang et al. (2024) for a review of the literature.

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.

For simplicity, our model only considers the case of two potential carpoolers per car, but the insight is more general. Most immediately, there can of course be more than two commuters in every cars, with road capacity expanding even further during the most demanded hours. Moreover, 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 several passengers (typically in a van), 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 may arise endogenously in developed countries as well, substantially raising the throughput of the roads (measured, again, in people rather than cars).

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 the upcoming 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 world with autonomous transportation.

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.[6] 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 $\underline{t} = 8{:}00$ AM and $\bar{t} = 9{:}00$ 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 $[\underline{t}, \bar{t}]$. 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.

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,

[6] According to the Bureau of Labor Statistics, the CPI adjustment from January 1969 until July 2025 is 9.07 (`https://data.bls.gov/cgi-bin/cpicalc.pl?cost1=1.00&year1=196901&year2=202507`).

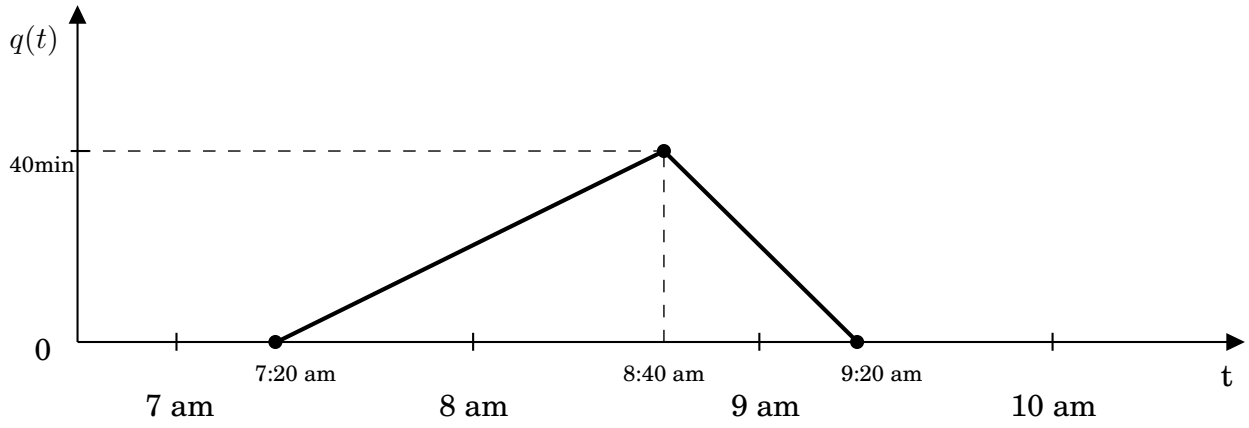


Figure 1: Pattern of traffic with no carpooling or congestion pricing

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 adopting self-driving technology, as we are hoping to do in the future), the social loss from traffic congestion will be reduced.[7] 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, $w_q = 4¢$.[8] 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 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 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

[7] 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`)

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

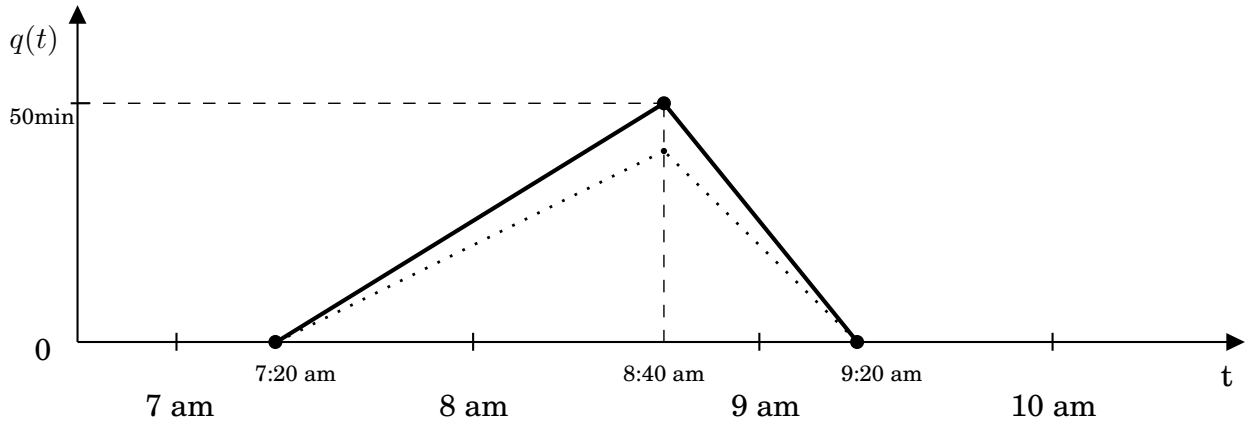


Figure 2: Pattern of traffic with reduced disutility of time in traffic

up by the same multiple, thus completely undoing the benefits of higher comfort and productivity.

To understand how the equilibrium traffic flows change when the disutility of time in the queue changes, it is helpful to consider function $b(t)$, analogous to the delay function $q(t)$ except that it denotes how much a commuter will have to spend in the queue if she *arrives* at the bottleneck at time t. Thus, she will leave the bottleneck at time $t + b(t)$. There is of course a close relationship between functions $b(t)$ and $q(t)$: specifically, $q(t) = b(t - q(t))$ and $b(t) = q(t + b(t))$.

Functions $b(t)$ for the two levels of per-minute disutility of queuing in traffic are illustrated in Figure 3. In the initial equilibrium (dotted line), cars arrive at the bottleneck at the rate of 120 cars per minute before 8:00am, and at the rate of 30 cars per minute after 8:00am. Recall that the rate of outflow from the bottleneck (i.e., its throughput) is 60 cars per minute. Thus, before 8:00am, the inflow of cars exceeds the outflow, and the queue grows. After 8:00am, the outflow exceeds the inflow, and the queue shrinks. Commuters who arrive at the bottleneck at 8:00am have to wait in the queue the most: 40 minutes. They will then leave the bottleneck at 8:40am.

In the equilibrium with the lower per-minute disutility of being stuck in traffic (solid line), cars arrive at the bottleneck at the rate of 160 cars per minute before 7:50am, and at the rate of $26\frac{2}{3}$ cars per minute after 7:50am. Commuters who arrive at the bottleneck at 7:50am have to wait in the queue the most: 50 minutes. They will then leave the bottleneck at 8:40am as well.

5 Equilibrium with Carpooling and without Congestion Pricing

We now introduce the possibility of carpooling. In the spirit of keeping the model simple, we make several assumptions, although our key insights do not hinge on them. We assume that each car can fit up to two people, and that carpoolers are indifferent between being a passenger and being a 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.

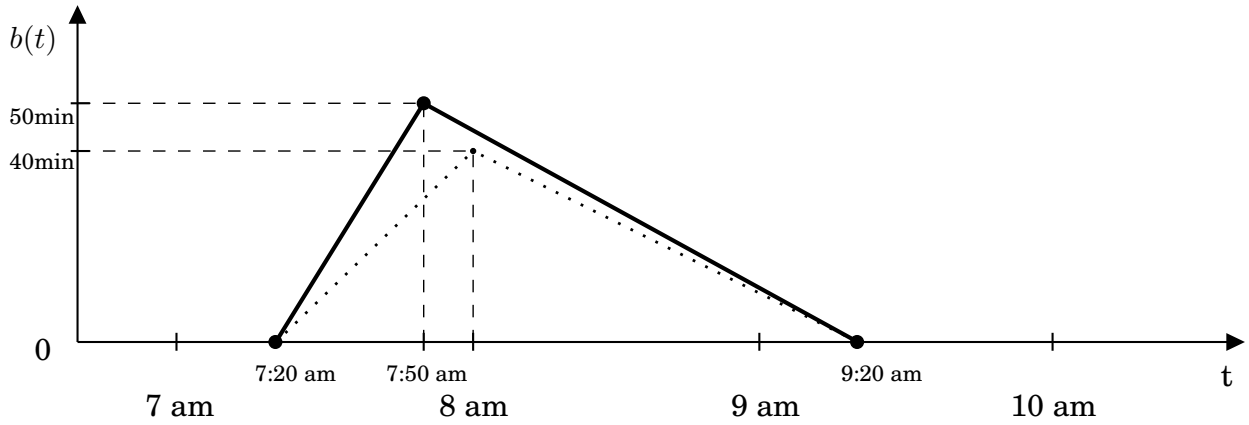


Figure 3: Patterns of waiting times as a function of arrival times

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.

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 4. 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.[9]

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

[9]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.

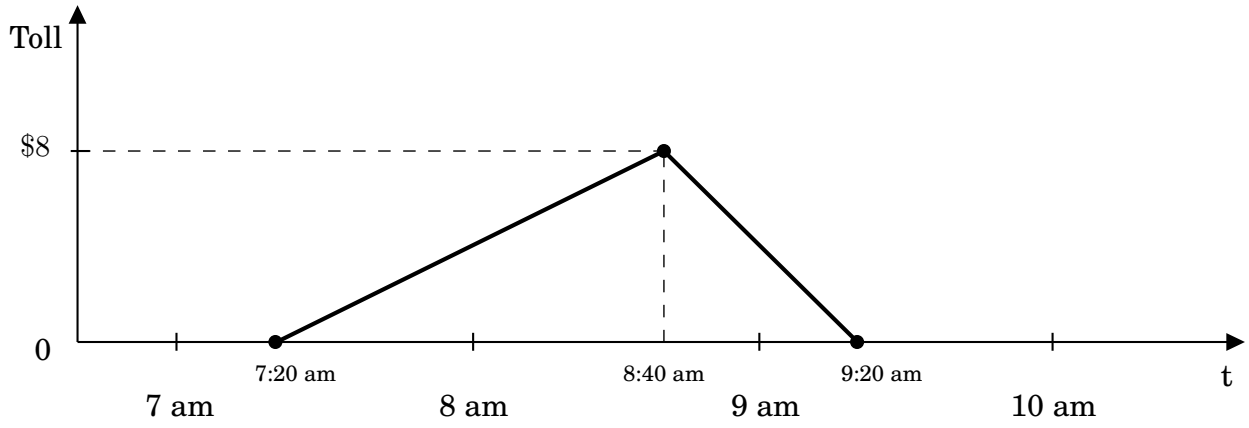


Figure 4: Optimal congestion pricing without carpooling

book), with heavily congested cities vastly outnumbering those rare exceptions that managed to overcome the political challenges and implement congestion pricing.

7 Equilibria with Congestion Pricing and Carpooling

We now consider the outcomes under congestion pricing in the presence of carpooling. We start out by describing the outcome under 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 socially optimal toll level is equal to the simple one around peak rush hour, but is lower before and after the peak.

Specifically, 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

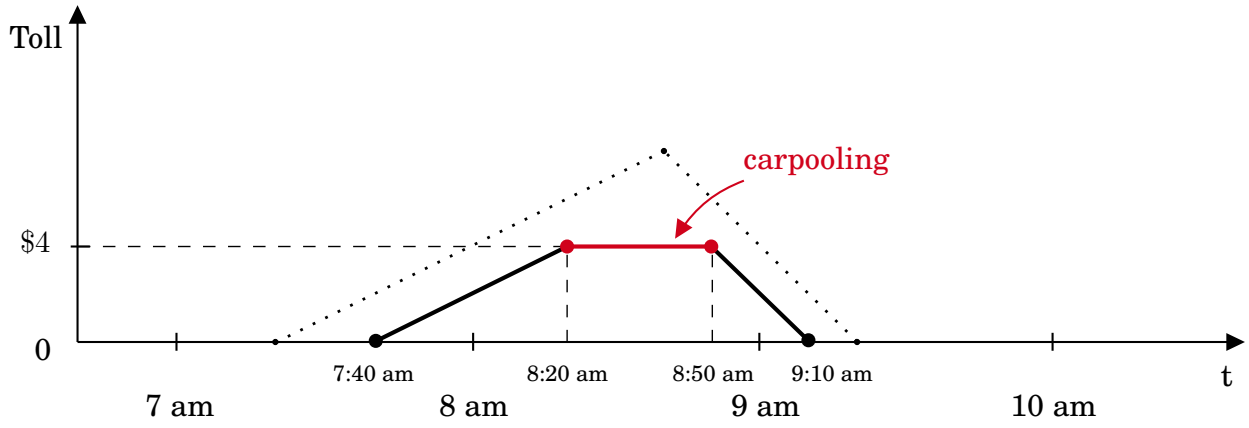


Figure 5: Optimal congestion pricing with carpooling

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 5. (The dotted line reproduces Vickrey's price schedule without carpooling, from Figure 4.) 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 (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.

	Cost ($/morning)						
	Displaced Arrival	Waiting in Queue	Disutility from Carpooling	Total Social Cost	Total Commuter Cost	Social Cost per Commuter	Average Commuter Cost
No tolls, no carpooling	14,400	28,800	-	43,200	43,200	6.00	6.00
No tolls, improved comfort	14,400	28,800	-	43,200	43,200	6.00	6.00
No tolls, carpooling	14,400	28,800	-	43,200	43,200	6.00	6.00
Tolls, no carpooling	14,400	-	-	14,400	43,200	2.00	6.00
Simple tolls, carpooling, $\triangle$ = $3	-	-	21,600	21,600	43,200	3.00	6.00
Simple tolls, carpooling, $\triangle$ = $2	-	-	14,400	14,400	28,800	2.00	4.00
Simple tolls, carpooling, $\triangle$ = $1	-	-	7,200	7,200	14,400	1.00	2.00
Optimal tolls, carpooling, $\triangle$ = $3	8,100	-	5,400	13,500	35,100	1.875	4.875
Optimal tolls, carpooling, $\triangle$ = $2	3,600	-	7,200	10,800	25,200	1.50	3.50
Optimal tolls, carpooling, $\triangle$ = $1	900	-	5,400	6,300	13,500	0.875	1.875

Table 1: Social and commuter costs

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 under various regimes 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.

Technological improvements (in the form of a higher value of time in traffic) do not lead to any changes in payoffs: as discussed in Section 4, the equilibrium response of commuters spending more time in traffic completely offsets the increase in comfort.

Likewise, 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.

Road pricing (on its own, without carpooling) does of course affect the outcome, but as discussed in Section 6, the overall effect is subtle. Waiting in traffic is eliminated, but 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. As we discuss in the Introduction and in Section 6, this is a major political problem, which is likely an important reason behind the lack of widespread adoption of congestion pricing across the world.

The outcomes change substantially once we consider road pricing and carpooling together, though specific numbers depend on the disutility of carpooling. For a moderate amount of disutility (\$2 per carpooler), the simple policy of charging a \$4 toll throughout the day (to induce everyone to carpool) results in no traffic and no disutility from displaced arrival times, though it does incur a disutility from carpooling. On net, the social cost per commuter is the same as under road pricing without carpooling (\$2), but crucially, each commuter is better off—and is better off than under the regime without congestion pricing: the average disutility of a commuter is \$4 (\$2 from the disutility of carpooling and an additional \$2 from the half of the \$4 toll). With optimal road pricing, under which the toll is at the \$4 maximum during the very peak of traffic (and everyone carpools during that time) and is lower during the early and late parts of the commute (and people drive solo during that time, but can be a bit early or a bit late for work), the outcomes are further improved: the average social cost per commuter is now \$1.50 while the average disutility faced by a commuter is \$3.00. The numbers are further improved if the disutility from carpooling can be lowered, e.g., to \$1.

9 Concluding Remarks

In this paper, we show that road pricing and carpooling are complementary. Our model closely follows Vickrey's canonical framework, 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.[10] 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 in half, thus creating incentives to carpool and leading to an increase in the throughput of the road (as measured in people, not in cars) and through that to a first-order improvement in the overall commuter and social welfare—*even if the revenue from tolls is destroyed.*

[10] Of course, optimal congestion pricing improves the overall social outcome, 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 practice, a typical car can seat five people, yet transports only about one and a half.[11] Thus, personal vehicles and public roads are utilized at less than one third of their capacity, at all times. Carpooling can eventually unlock enormous efficiency gains that for now remain unrealized, because with existing technologies and institutions, the individual benefit from carpooling (both relative to driving solo and especially relative to the social benefit that additionally includes the externality on others) is too low. Congestion pricing has a side effect of increasing incentives to carpool. As we show in Section 8, this side effect is of first-order importance for estimating potential welfare gains from optimal tolls. Leaving it out of policy analysis would substantially underestimate the benefits of congestion pricing.

A natural question is to what extent realistic levels of congestion pricing can lead to substantial increases in carpooling. On this question, there is encouraging evidence, both direct and indirect. Direct evidence comes from the pioneer of congestion pricing—Singapore. When Singapore introduced their congestion pricing scheme in 1975, it exempted cars with four or more occupants ("car pools") from the fee. As Holland and Watson (1977) report, "The exemption [dots] induced a large increase in occupancy during the restricted hours. The absolute number of car pools entering during the 7:30 – 10:15 period increased by about 60 per cent at the same time that the total number of cars including car pools was falling by 73 per cent. The proportion of car pools thus jumped from less than 7 to 37 per cent for those hours." Indirect evidence comes from natural experiments in which the cost of parking changed sharply, since costly parking provides similar incentives to carpool to those provided by congestion tolls. Wilson and Shoup (1990) discuss three instances in which free parking was taken away from employees who drove solo. The change led to a reduction in the number of employees driving to work solo ranging from 20% to 81%. The largest corresponding increase in carpooling was observed at a large private firm located in the Warner Center—a suburban setting without public transportation in the Los Angeles area. The rate of carpooling has increased from just 6% to 48%.

Another takeaway from our results is that reducing the disutility Δ from carpooling is critically important, and complementary to congestion pricing. There are a number of low- or no-cost policy changes that could reduce this disutility. One substantial component of the disutility from carpooling is the time and friction involved in arranging one. Another major component is the loss of flexibility. Unlike carpooling companions, solo drivers do not need to spend time on coordinating with anyone and they to not have to decide in advance when to head out. A well-designed carpooling platform can significantly reduce the disutility from carpooling by making it easy to find carpooling matches with short detours.[12] Likewise, a well-designed platform could help carpoolers seamlessly

[11] `https://www.energy.gov/eere/vehicles/articles/fotw-1333-march-11-2024-2022-average-number-occupants-trip-household`

[12] For example, a highly-successful pre-COVID "Casual Carpool" between Oakland and San Francisco required zero advance planning and minimal detours by the drivers (`https://www.sfmta.com/casual-carpool`). It is now being revived, with the flexibility and ease of coordination again being major benefits ("Longtime participant Greg Riessen created an app that puts a new twist on the program called Rapid Carpool. 'The driver and the rider are connected just moments before the driver is getting off the freeway, and so the driver doesn't even have to exit the freeway and come looking for riders,' Riessen said," `https://abc7news.com/post/casual-carpool-popular-commuter-option-cross-bay-bridge-makes-post-pandemic-return-oakland-stop/17513559/`). While such examples can spring

share costs of driving and other related expenses, such as tolls and parking. However, currently both private and public investment into creating such platforms are relatively low. To our knowledge, none of the major technology companies are currently investing into building such a platform in the United States.[13] One major obstacle is the ambiguity of the current legal framework, which discourages tech companies from investing in carpooling platforms. This may at first seem surprising, since the social desirability of carpooling is well understood by policymakers. E.g., that is why "not-for-profit" carpooling is allowed by standard car insurance policies in all states. Where the legal ambiguity arises is at the point at which a private company starts arranging carpools on a large scale in a centralized fashion. A carpool is considered "not-for-profit" if the driver is reimbursed for her costs and does not make a profit. However, legally acceptable reimbursement rates are open to interpretation. For example, if a driver pays for monthly parking, does carpool remain non-profit if the passenger covers the corresponding daily cost, in addition to the gas and depreciation of the car? Can a driver who carpools in a high-end luxury Lexus SUV legally receive a higher reimbursement from her passenger than a driver who carpools in a bare-bones Toyota sedan, and if so, by how much? If a driver and a passenger were matched by a for-profit platform that handles the payments, does it still constitute a non-profit carpool? Can the platform that matches the driver and the passenger be held liable in the event of an accident? In the Appendix, we discuss a number of these issues and the corresponding steps that policymakers and regulators can undertake to reduce the legal ambiguity surrounding commercially-organized carpooling, thus encouraging private companies to create well-designed platforms that would lower the disutility from carpooling. These steps are particularly important in regions that have introduced (or consider introducing) congestion pricing to alleviate traffic delays.

As we mentioned in the Introduction, while for simplicity the model in our paper involves just two carpoolers, the insights are more general, and apply both to carpools and to various modes of "intermediate public transport," with vans of various sizes that operate with multiple stops and short detours. Such modes of transportation are common in low- and middle-income countries (Jaiswal et al., 2022), but have not gained widespread traction in higher-wage nations, likely because of relatively high driver costs. Efficient congestion pricing may enhance the viability of such types of services in rich nations, by making them relatively more attractive than solo driving (via the main mechanism behind our results—a solo driver pays the entire toll, while a passenger in a 10-seat shared minivan pays only one tenth of that amount). As congestion pricing gradually becomes more widespread, it becomes increasingly important for policymakers to create a favorable legal framework that would make it easier for entrepreneurs to enter the market for intermediate public transport, with dramatic social gains due to the the corresponding increase in road throughput (as measured in people, not in vehicles). Moreover, as autonomous vehicles are finally commercially

up "organically" in some particularly suitable locations, a broader system with a major impact on the overall traffic congestion would require a much more substantial investment.

[13] While there are a number of smaller start-ups that attempt to enter this space, they are not as well positioned to successfully build such platforms on a large scale and to get to a critical mass of good carpool matches, because they lack the location history for hundreds of millions of users that large technology company already have.

available to the public and are spreading rapidly around the world,[14] it is critical to ensure that the society does not end up with most of them carrying just one passenger per vehicle during peak congestion times and destroying the surplus from the incredible technological innovation—a danger we illustrate in Section 4. Instead, with efficient congestion pricing and appropriate regulations, self-driving vehicles have the potential to erase the boundary between private solo driving and public transportation, leading to dramatic increases in welfare for both the commuters and the society as a whole. In Ostrovsky and Schwarz (2018), we present a model of how such a system can be designed and organized, and discuss in depth the interplay among carpooling, congestion pricing, and self-driving vehicles.

[14]While the progress and roll-out of self-driving technology has been gradual over the last decade and a half, we view June 25, 2024 as the specific date when one could say that it was finally available to the public. That was when Waymo launched its commercial driverless service in San Francisco, removing all waiting lists and other barriers and thus making it available to anyone with a cellphone (`https://waymo.com/blog/2024/06/waymo-one-is-now-open-to-everyone-in-san-francisco`). This was followed by a rapid rollout of the service to several other regions (`https://support.google.com/waymo/answer/9059119`, `https://support.google.com/waymo/answer/16011725`). Baidu's Apollo Go commercial service has operated without a safety driver in multiple cities in China since February 2025 (`https://technologymagazine.com/articles/how-baidus-apollo-go-targets-global-robotaxi-expansion`). A number of other technology companies are in advanced stages of testing in the United States, China, and other parts of the world.

Appendix

In this Appendix, we discuss in more detail some of the legal frictions that need to be addressed to facilitate the growth of carpooling and carpooling matching platforms.

Policymakers have long understood the societal benefits of carpooling, which reduces road congestion and air pollution, hence benefiting all citizens, not just carpoolers. Carpooling has a long history in the United States. E.g., during World War II, resource shortages have caused the U.S. government to promote carpooling with a famous poster, ""When You Ride ALONE You Ride with Hitler!"[15] However, the legal infrastructure around carpooling has not been disambiguated for the age of platforms, electronic payments, and GPS-enabled cellphones. Policymakers can encourage entrepreneurs[16] to invest into developing carpooling platforms by removing legal uncertainties about how the laws might be interpreted.

In essence, all standard insurance policies in all states cover "not-for-profit carpooling." Not-for-profit carpool is defined as a carpool in which the driver might obtain some form of reimbursement for driving, but does not profit from it. Stringent commercial transportation regulations do not affect carpooling as long as it remains "not-for-profit." While this legal structure is sensible and seems straightforward, it becomes more ambiguous in the age of platforms. Below is a list of legal and regulatory questions that an entrepreneur considering creating a carpooling platform may face. To reduce barriers to entry and encourage investment into carpooling technology, regulators and policymakers should provide unambiguous favorable answers to these questions.

1. Does carpool remain "not-for-profit" (and hence legal) if it is arranged on a platform that charges a commission for each ride, accepts payments from passengers, and makes reimbursement payments to drivers? A number of for-profit companies have attempted to build carpooling platforms, and to the best of our knowledge, the assertion that carpool arranged on a platform remains "not-for-profit" has never been challenged. But since it hasn't been litigated, the answer remains ambiguous, making it risky for entrepreneurs and their backers to invest in such initiatives. A confirmation from regulators would reassure potential entrepreneurs that carpooling remains "not-for profit" from the point of view of both insurance and transportation regulation, regardless of whether it was arranged among friends or facilitated by a commercial platform.

2. What is the maximum reimbursement rate that preserves the "not-for-profit" status of carpooling? The IRS rules state that the allowed reimbursement rate for charitable driving is 14 cents per mile and for business driving it is 70 cents per mile.[17] Does carpool remain "not-for-profit" if the driver is reimbursed at the 70-cents-per-mile rate? In many congested areas,

[15] `https://digitalcollections.hoover.org/objects/39522/when-you-ride-alone-you-ride-with-hitler-join-a-carsharin`

[16] The term "entrepreneurs" here includes both those starting new organizations to focus on these platforms and innovators within existing organizations starting new initiatives.

[17] `https://www.irs.gov/tax-professionals/standard-mileage-rates`

parking and tolls may cost more than gas and car depreciation. Will carpool remain "not-for-profit" if in addition to the 70 cents per mile reimbursement rate, the driver is also reimbursed for the cost of parking and tolls? If a monthly parking permit costs $300, how much of it is a passenger allowed to reimburse to the driver without violating non-profit status of carpooling? Is it zero, because "not-for-profit" carpool allows reimbursements only for miles driven rather than parking fees? Is it $10 dollars (because parking costs $10 per day), or is it $15 (because parking at work is only used on workdays), or is it $5 (because the driver can carpool both on the way to and from work, thus collecting reimbursement for two rides in a single day)? Or perhaps it is only half of that amount, with the idea that parking fees are equally shared between the driver and the passenger? What if the carpool involves more than two people (which of course would be wonderful from the point of view of reducing congestion)? While these questions are routinely resolved by individual carpoolers amongst themselves on an ad hoc basis, without attracting regulators' attention, for a centralized, potentially large-scale platform that facilitates carpooling (and drives down the disutility Δ by reducing frictions), it is necessary to have clear answers. Clear guidance from regulators and policymakers regarding maximum reimbursement rates would be helpful for anybody considering investing into the development of a carpooling platform. Much like the IRS today provides per-mile reimbursement rates for personal driving, it could provide maximum allowed reimbursement rates for "not-for-profit" carpooling trips, potentially taking into account different costs of tolls and parking across different geographical areas. To prevent abuse, the drivers could be limited to at most two trips per day (presumably to and from work).

3. In the event of an injury accident, a carpool passenger might demand damages from both the carpool driver and the platform that arranged the carpool (beyond the amount covered by the carpool driver's insurance policy and the insurance policies of other drivers potentially involved in the accident). It would be logical to allow the platform to include in its terms of service a legally binding liability waiver that shields it from such demands, and also shields the carpooling drivers from such demands from their passengers (unless, say, the carpool driver is convicted of a felony that caused damages to the passenger). Without such a waiver, a carpool passenger would in effect get greater protection from possible accidents than a commuter driving solo, despite not facing greater risks. Moreover, not being able to rely on such limits to liability may discourage entrepreneurs and their backers from building such platforms, and may discourage drivers from participating in them. A carpool-friendly policy would give the platforms the ability to include such legally binding liability waivers in their terms of service.

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.

Hall, J. 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.

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.

Huang, H.-J., H. Yang, and M. G. H. Bell (2000). The models and economics of carpools. *The Annals of Regional Science 34*(1), 55–68.

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.

Konishi, H. and S. Mun (2010). Carpooling and congestion pricing: HOV and HOT lanes. *Regional Science and Urban Economics 40*(4), 173–186.

Kreindler, G. (2024). Peak-hour road congestion pricing: Experimental evidence and equilibrium implications. *Econometrica 92*(4), 1233–1268.

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.

Ostrovsky, M. and M. Schwarz (2018). Carpooling and the economics of self-driving cars. NBER Working Paper 24349, available at `https://www.nber.org/papers/w24349`.

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.

Wilson, R. W. and D. C. Shoup (1990). Parking subsidies and travel choices: Assessing the evidence. *Transportation 2*(17), 141–157.

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

Zhang, Z., R. Yang, and M. Wu (2024). Designing high-occupancy toll lanes: A game-theoretic analysis.