



Autonomous cars and dynamic bottleneck congestion: The effects on capacity, value of time and preference heterogeneity



Vincent A.C. van den Berg^{a,b,*}, Erik T. Verhoef^{a,b}

^a Department of Spatial Economics, VU Amsterdam, De Boelelaan 1105, 1081 HV, Amsterdam, The Netherlands

^b Tinbergen Institute, Gustav Mahlerplein 117, 1082 MS, Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Received 15 September 2015

Revised 26 August 2016

Accepted 30 August 2016

Available online 19 September 2016

JEL codes:

D42

D62

H23

R41

R48

Keywords:

Autonomous cars

Heterogeneity

Bottleneck model

Self-driving cars

Road capacity

ABSTRACT

'Autonomous cars' are cars that can drive themselves without human control. Autonomous cars can safely drive closer together than cars driven by humans, thereby possibly increasing road capacity. By allowing drivers to perform other activities in the vehicle, they may reduce the value of travel time losses (VOT). We investigate the effects of autonomous cars using a dynamic equilibrium model of congestion that captures three main elements: the resulting increase in capacity, the decrease in the VOT for those who acquire one and the implications of the resulting changes in the heterogeneity of VOTs. We do so for three market organizations: private monopoly, perfect competition and public supply. Even though an increased share of autonomous cars raises average capacity, it may hurt existing autonomous car users as those who switch to an autonomous car will impose increased congestion externalities due to their altered departure time behaviour. Depending on which effect dominates, switching to an autonomous vehicle may impose a net negative or positive externality. Often public supply leads to 100% autonomous cars, but it may be optimal to have a mix of car types, especially when there is a net negative externality. With a positive (negative) externality, perfect competition leads to an undersupply (oversupply) of autonomous cars, and a public supplier needs to subsidise (tax) autonomous cars to maximise welfare. A monopolist supplier ignores the capacity effect and adds a mark-up to its price.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

'Autonomous cars' are cars that drive themselves and have an automated speed choice. They can drive closer together and at more uniform speeds than 'normal' human driven cars. All else being equal, a group of autonomous cars can move at a greater density for any given speed than 'normal cars', thereby increasing the capacity of roads (Chang and Lai, 1997). Besides this favourable capacity effect, people adopting an autonomous car instead of a normal car may gain a decrease in their value of travel time (VOT) as time in the car can be spent on other activities besides driving. This makes travel time more useful and lowers its costliness. As a result, the VOT may become more heterogeneous for a mix of drivers using normal and autonomous cars. Such heterogeneity may alter the effects of introducing autonomous cars and of policies such

* Corresponding author.

E-mail addresses: v.a.c.vanden.berg@vu.nl (V.A.C. van den Berg), e.t.verhoef@vu.nl (E.T. Verhoef).

Nomenclature

α	Value of time (VOT) for normal car users: the cost of an hour of travel time.
β	Value of schedule delay early: the cost of an hour earlier arrival than the preferred arrival time t^* .
γ	Value of schedule delay late: the cost of an hour later arrival than the preferred arrival time t^* .
δ	Compound preference parameter: $\delta \equiv \beta\gamma/(\beta+\gamma)$.
θ	VOT reduction parameter for an autonomous car: the VOT for an autonomous car is $\theta \cdot \alpha$ with $\beta/\alpha < \theta < 1$.
ΔC	Change in travel cost when switching from a normal car to an autonomous car: $\Delta C = C_a - C_n$.
C_i	Travel cost for car type $i = \{a, n\}$ with a indicating an autonomous car and n a normal car. Travel cost is the sum of free-flow travel time cost and bottleneck cost, which equals the queuing time cost plus the schedule delay cost.
CB	Bottleneck cost equals the queuing time cost plus the schedule delay cost.
f	Fraction of users with an autonomous car.
MC_a	Marginal automobile cost of the autonomous car. The corresponding cost for a normal car is set to zero. MC_a gives how much higher or lower the sum of the fuel cost and per trip marginal production cost is for an autonomous car.
MEB	Marginal external benefit of marginally increasing the number of the users with an autonomous car; it equals the change in travel cost minus the change in average social cost: $MEB = \Delta C - (\partial TTC / \partial f) / N$.
MU_a	Per trip mark-up on the autonomous car. The mark-up on the normal car is normalised to zero.
N	Total number of users.
P	(Generalised) price equals the travel cost plus the per trip marginal automobile cost and the mark-up.
$r[f]$	Function determining the effective bottleneck capacity for autonomous cars of $s/r[f]$.
s	Bottleneck capacity for a normal car.
t	Arrival time.
t^*	Preferred arrival time.
t_e	Moment that the last car arrives and thus the peak ends.
t_s	Moment that the first car arrives and thus the peak starts.
TT	Travel time, which equals free-flow travel time plus queueing time.
TT_{ff}	Free-flow travel time.
TTC	Total travel cost of $N \cdot (f \cdot C_a + (1-f) C_n)$.
TC	Total cost of $TTC + f \cdot N \cdot MC_a$.

as congestion pricing and capacity expansion. It also enforces the relevance of considering distributional effects of policies (see, e.g., [Arnott et al. \(1988\)](#), [Lindsey \(2004\)](#) and [Van den Berg and Verhoef \(2011a, b\)](#)).

Autonomous cars may also lead to fewer accidents, to a smaller or larger total car fleet, to fewer parking spots, to higher speed limits, and to reduced fuel use.¹ There are also potential problems, including the question of liability for accidents, reliability, loss of privacy, and risks of hacking of autonomous cars. For more detailed overviews see [Anderson et al. \(2014\)](#) and [Fagnant and Kockelman \(2015\)](#).

Autonomous cars can be expected to have a considerable influence on urban transport and in the long run on the layout of our cities. This explains the strong interest of policymakers, the media and the general public in autonomous cars as a solution to our transport problems. This paper investigates the welfare impacts of the uptake of autonomous cars. It concentrates on how autonomous cars affect congestion via three channels: the resulting increase in capacity, the decrease in the VOT and the implications of the resulting changes in the heterogeneity of the VOT. We do not consider heterogeneity other than that due to the decision to obtain an autonomous car or not. The VOT is the same for everyone who has the same car. [Section 6](#) gives an exploratory discussion of some model extensions, including pre-existing heterogeneity. Our focus on the congestion effects of autonomous cars reflects the emphasis this has received in policy debates and the media in a congested country like the Netherlands. We could add differences in environmental or safety externalities between the two types of vehicles. However, from an analytical perspective, this would be rather straightforward and would divert attention from the behavioural and congestion impacts in which we are interested.

We are the first to consider the effects of autonomous cars via the VOT, and to consider autonomous cars with endogenous departure time and car type decisions. To achieve this objective, we use the bottleneck model, as do [Lamotte et al. \(2016\)](#). They also have an endogenous choice of car type. However, they have separate roads for normal and autonomous cars and assume that autonomous cars cooperate by having a departure rate equal to the capacity of their road. In their setting, autonomous cars do not affect preferences. [Levin and Boyles \(2016\)](#) use a cell transmission model to study the route choice equilibrium with autonomous and normal cars, but the choices of departure time and car type are exogenously fixed.

¹ If only autonomous cars use a road, it may become possible to redesign it: e.g., to reduce lane widths and turn a two-lane highway into a three-lane one.

With a heterogeneous VOT, users self-select over time. With queuing, the users with the highest VOTs travel in the edges of the peak when queuing delays are minimal; the users with the lowest VOTs travel in the middle of peak and accept long travel times in exchange for a convenient arrival time (Arnott et al., 1988). If users' VOTs fall after buying an autonomous car, they will impose increased external costs on users who arrive closer to the preferred arrival time than they do. The reason is that they now care less about travel time losses and therefore 'need' and create a steeper travel time development over time to be in dynamic user equilibrium (Lindsey, 2004). A user's marginal external cost decreases with its VOT (Van den Berg and Verhoef, 2011a), and when the VOT decreases due to acquiring an autonomous car, this raises this user's marginal external cost.

This 'heterogeneity effect' via the marginal external cost means that, ignoring the capacity effect, an increase in the share of autonomous cars raises the costs for existing autonomous car users. If the capacity effect is sufficiently strong, the total cost is minimised when everyone has an autonomous car; but if not, it is socially optimal when not everyone has one.²

Total cost not only includes travel costs, but also automobile costs due to depreciation and fuel use. Autonomous cars are likely to be more expensive to produce than normal cars, implying higher depreciation costs. However, autonomous cars may also have a higher residual value on resale, implying a lower depreciation cost. At the same time, autonomous cars may have lower insurance premiums and may use less fuel as they drive at more uniform speeds.

We examine three provision schemes of autonomous cars: socially optimal 'public provision', the perfectly competitive case of 'marginal cost provision', and provision by a profit-maximising 'monopolist'. The effects of the provision regimes of autonomous cars have not previously been studied. Private supply would approach the marginal cost case as the provision becomes more competitive. The other extreme of private provision is the private monopoly. When the capacity effect is stronger than the heterogeneity effect, there is a net positive externality and autonomous cars need to be priced below the marginal production cost when normal cars are priced at marginal cost and roads are priced suboptimally. When the heterogeneity effect dominates, marginal cost provision leads to over-consumption of autonomous cars. Provision by a monopolist typically leads to a substantially lower supply of autonomous cars than marginal cost provision as this supplier uses its market power to maximise profits.

We only study settings with unpriced congestion. With optimal congestion pricing and charges that differ by car type, the pricing of cars at marginal production cost is optimal and the issue of potential under- or oversupply with a competitive autonomous car market will disappear.

For our results, the strength of the capacity effect is an important determinant. The literature overview in Section 2 will show a large variation in predictions. The average prediction is that capacity would double with a switch from normal cars only to autonomous cars only. However, predictions vary from a rise of only 1% to a rise of over 400%. These results refer to switching from 100% normal cars to 100% autonomous cars. With a mix of car types, capacity gains may remain modest: the marginal capacity effect probably increases with the level of penetration because this allows autonomous cars to cooperate more efficiently (Tientrakool et al., 2011). We incorporate this by making the capacity for autonomous cars a function of the fraction of autonomous vehicles in the total fleet.

Another important parameter is the effect of using an autonomous car on the VOT, the size of which is an even more open question than the size of the capacity effect. Finally, there is little to no guidance from the literature on how much higher or lower the costs of purchasing and operating autonomous cars will be, when compared to normal cars. The effects of autonomous cars and their provision regimes are sensitive to these parameters. Hence, it is vital to do extensive sensitivity analyses, and we will do so in Section 5.

For ease of reference, the nomenclature box below summarises the notation. The notation will also be introduced in the text.

2. Overview of the research on the effect of autonomous cars on capacity

It is difficult to assess the possible effect of autonomous cars on the capacity of our roads. There is a large literature that considers this question. But it remains an open and debated issue. Table 1 gives an overview of some estimates. Zwaneveld and Van Arem (1997) review the early literature and argue that a doubling or tripling of capacity is likely. In more recent literature, the expected increase in capacity can be as low as 1% for autonomous cars that do not cooperate (Shladover et al., 2012) and as high as 414% with very efficient cooperation (Fernandes and Nunes, 2012). The predictions thus vary from almost no effect to a quintupling of capacity.

All results refer to switching from 100% normal cars to 100% autonomous cars. If car types travel mixed, the effects could be much less beneficial as the strongest gains are likely to be realized when going from many to only autonomous cars (e.g. Tientrakool et al., 2011). This would make the capacity a highly convex function of the share of autonomous cars. Van Arem et al. (2006) even find that introducing autonomous cars may reduce capacity when their share is low. Levin and Boyles (2016) find that the fraction needs to exceed 75% for capacity not to decrease due to inefficiencies with mixed traffic at intersections. In the dynamic equilibrium of our model, autonomous and normal cars will travel perfectly separated in time. Nevertheless, we assume that the capacity convexly increases with the penetration rate via a non-linear capacity effect.

² Gubins and Verhoef (2011) find a similar result on the effects of teleworking. Lamotte et al. (2016) also find a similar result when studying cooperation between vehicles.

Table 1
Increase in capacity from switching from only normal cars to only autonomous cars.

Study	Uncooperative autonomous cars	Cooperative autonomous cars
Chang and Lai (1997)	33%	X
Fernandes and Nunes (2012)	x	84–230% at 36 km/h, 186–414% at 72 km/h
Ni et al. (2010)	x	20–50%
Shladover (2011)	x	80%
VanderWerf et al. (2002)	7%	Around 220%
Shladover et al. (2012)	1–4%	97%
Tientrakool et al. (2011)	90% at 50 km/h, 40% at 100 km/h	200% at 50 km/h, 270% at 100 km/h

This may reflect that in reality cars may interact on the network elsewhere, but also that the separation over time of car types may vanish if there were pre-existing heterogeneity.

3. The model

3.1. Set-up

We use the Vickrey (1969) bottleneck model to capture the effects of autonomous cars in a stylised setting that allows for closed-form analytical solutions while capturing the important behavioural aspects discussed above. We focus on how introducing autonomous cars changes the equilibrium and will be brief in our description of the standard model. For complete reviews see Arnott et al. (1993, 1998) or Small (2015).

We assume that demand is fixed. We thus ignore that, if autonomous cars lower costs, this will increase demand and congestion. We also ignore other transport modes such as public transport and taxis (see International Transport Forum (2015)).

The travel time cost equals travel time multiplied by the value of time (VOT). The VOT for a normal car is α ; for an autonomous car, it is $\theta \cdot \alpha$, where $\theta < 1$ is the VOT reduction parameter. Travel time, TT , is the sum of free-flow travel time, TT_{ff} , and the delay from queueing at the bottleneck. The queueing delay equals the number of cars in the queue at the instant of joining it divided by the capacity of the bottleneck during the queueing time. When normal cars pass the bottleneck, the capacity is s . When autonomous cars pass the bottleneck, the capacity is $s/r[f]$ with $0 < r[f] \leq 1$. The lower $r[f]$ is, the higher the capacity for autonomous cars.

A person's bottleneck cost, CB , equals the queueing time cost plus the schedule delay cost. The schedule delay cost is the cost due to arriving at a different time than the most preferred arrival time, t^* . We assume that t^* is the same for all and normalise it to 0. We follow Small (1982) and use schedule delay costs that are linear in the time difference between t^* and the actual arrival time, t . The shadow cost per hour for arrivals earlier than t^* is β ; for hours late it is γ . The schedule delay cost thus equals $Max[-\beta \cdot t, \gamma \cdot t]$.

The travel cost equals the free-flow travel time cost plus the bottleneck cost, CB :

$$C_n[t] = CB_n[t] + \alpha \cdot TT_{ff} = Max(-\beta \cdot t, \gamma \cdot t) + \alpha \cdot TT[t], \quad \text{with a normal car,} \quad (1)$$

$$C_a[t] = CB_a[t] + \alpha \cdot \theta \cdot TT_{ff} = Max(-\beta \cdot t, \gamma \cdot t) + \alpha \cdot \theta \cdot TT[t], \quad \text{with an autonomous car,} \quad (2)$$

Subscript n indicates a normal car and a indicates an autonomous car.

As Arnott et al. (1988) show, users with a lower VOT self-select themselves into the centre of the peak, when travel times are long and schedule delays are small. These users care less about the travel time and can thus enjoy the small schedule delays in the centre peak. Likewise, drivers with a larger VOT self-select into the early and late shoulders of the peak. Consequently, normal and autonomous cars will travel separated over time, and the bottleneck capacity is either s or $s/r[f]$.

The generalised price per trip (or price for the sake of brevity) is the sum of the travel cost, the automobile cost and the mark-up on the car. The mark-up is determined by the producer of the car and is expressed in a per-trip equivalent. The generalised cost (or cost for brevity) is the price minus the per-trip mark-up. The automobile cost is assumed to be a fixed amount per trip. This cost includes the fuel cost and the depreciation per trip, capturing the marginal production cost of the car. Autonomous cars probably have a lower fuel cost, but will probably be more expensive to produce and purchase, raising the depreciation cost. The depreciation cost should also account for the residual value when selling the car. Autonomous cars may have a higher or lower residual value than normal cars, which also depends on how fast their technology becomes obsolete.³ It seems likely that autonomous cars will have a higher net automobile cost. Nonetheless, we do not impose this, and our numerical analyses also study autonomous cars with lower automobile costs. The mark-up on the car is also reworked into a per-trip equivalent. The per-trip profit for a car manufacturer equals the number of cars times this mark-up.

³ We ignore the risk of accidents to keep the model tractable. Autonomous cars may have lower accident risks, if only due to the absence of drunk driving. This may lower insurance premiums for autonomous cars, thus lowering the per trip cost.

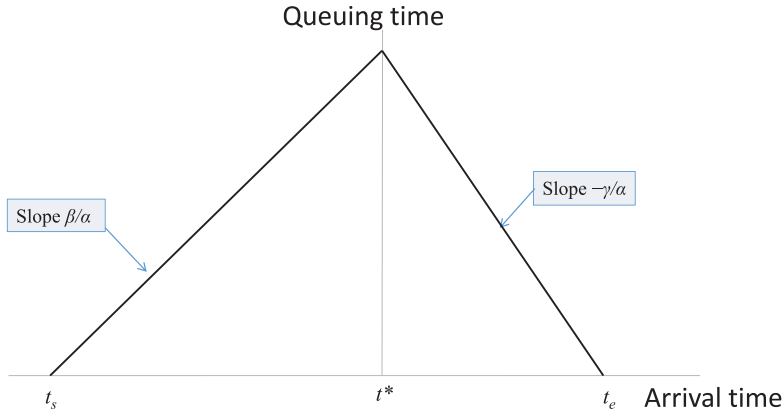


Fig. 1. Equilibrium queuing times without autonomous cars.

MU_a gives the mark-up per trip on autonomous cars. This mark-up may be negative, zero or positive. For normal cars, we normalise the automobile cost and the mark-up to zero. The automobile cost for an autonomous car, MC_a , therefore reflects by how much its marginal production cost exceeds that for normal cars.

In the dynamic equilibrium, for both types of car, the travel costs need to be constant over the arrival times used by their users and can be no lower at other times. Therefore, we can omit the arrival time indicator and write the equilibrium prices as:

$$P_n = C_n, \tag{3}$$

$$P_a = C_a + MC_a + MU_a. \tag{4}$$

3.2. Standard bottleneck model without autonomous cars

We now briefly reintroduce the bottleneck model without autonomous cars and thus with homogeneous users. The three subsections hereinafter look at three cases of interest, namely where autonomous cars only affect capacity, only the VOT, and finally both in the full model. This order of presentation helps identify the different effects of autonomous cars.

In the dynamic equilibrium with homogeneous users, the travel cost should be constant over time as long as arrivals occur. Fig. 1 illustrates this equilibrium. It shows the travel time pattern over arrival time, t , which leads to constant costs. The curve can be interpreted as an iso-cost function since costs are constant along it. Shifting the curve upwards implies a higher cost level. This iso-cost interpretation will prove useful later on. The first driver arrives at the starting time, t_s , of the peak and faces a zero queuing time. The last driver arrives at t_e and also incurs no queuing.

To keep the travel cost constant over time before t^* , the travel time must grow linearly over arrival time at a rate β/α , so that the decrease in schedule delay cost over arrival time is matched by the increase in the queueing time cost. Similarly, after t^* , the travel time must shrink at a rate γ/α over arrival time.

The equilibrium travel cost is:

$$C = \delta \cdot \frac{N}{s} + \alpha \cdot TT_{ff}, \quad \text{with} \quad \delta = \frac{\beta \cdot \gamma}{\beta + \gamma}. \tag{5}$$

N is the number of users and δ is a composite scheduling preference parameter. The term $\delta \cdot N/s$ gives the bottleneck cost, which consists of the schedule delay cost and the queuing time cost. The term $\alpha \cdot TT_{ff}$ gives the free-flow travel time cost.

3.3. Special case 1: autonomous cars only increase capacity

The first special case that we consider assumes that autonomous cars only affect capacity and not the VOT. The equality of preference parameters means that drivers will not self-select into separate time windows, but will instead travel at the same moments in a pooled equilibrium. Of the users, a fraction f has an autonomous car and faces a capacity of $s/r[f] \geq s$ when they pass the bottleneck. The others have normal cars and face a capacity of s . The new equivalent capacity is the weighted harmonic mean of the normal car capacity and the autonomous car capacity. For now, f is treated as a given. Section 4 will derive the equilibrium shares. The literature review found that the maximum capacity gain is very uncertain: $r[f=1]$ may be close to 1 or as low as 1/5. Moreover, autonomous cars may increase capacity more when they form a larger fraction of the fleet of cars. Therefore, we assume that $r[f]$ is a concavely decreasing function of f , which we ensure by using $\partial r/\partial f < 0$, $\partial^2 r/\partial f^2 \leq 0$ and $0 < r[f] \leq 1$.

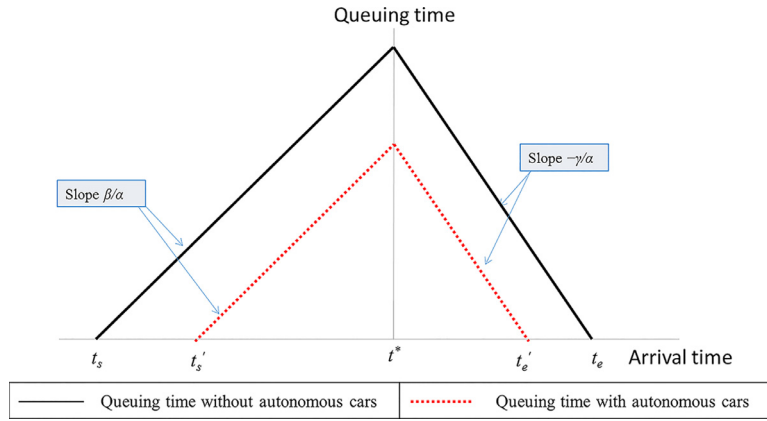


Fig. 2. Equilibrium queuing times when autonomous cars only affect capacity.

Fig. 2 illustrates the equilibrium. The solid line shows the queuing times without autonomous cars. The dotted line applies with a positive number of users with an autonomous car. The increased capacity shifts all users to the lower dotted iso-cost curve. Everyone gains the same cost reduction. The slope of the equilibrium travel times is the same for all drivers because α , β and γ are equal across users. Hence, everyone is on the same iso-cost curve and thus faces the same cost.

Without an effect on the VOT, the equilibrium travel cost is:

$$C_a = C_n = \delta \cdot \frac{N}{s} ((1 - f) + r[f] \cdot f) + a \cdot TT_{ff}. \tag{6}$$

This is the same cost function as for the standard bottleneck model, except that a fraction f of the users face a higher capacity of $s/r[f]$. The travel cost does not differ by car type as everyone has the same preferences.

The total travel cost, TTC , equals:

$$\begin{aligned} TTC &= N \cdot (C_n \cdot (1 - f) + C_a \cdot f) \\ &= \delta \cdot \frac{N^2}{s} ((1 - f) + r[f] \cdot f) + a \cdot TT_{ff} \cdot N. \end{aligned} \tag{7}$$

A higher f increases the total capacity and lowers the total travel cost. In this special case, the social optimum is always at a corner solution: either everyone should have an autonomous car if the extra automobile cost of autonomous cars is not too great or otherwise nobody should have one.

Because normal and autonomous cars have the same travel costs, a marginal user who switches car type will experience a zero change in travel cost, ΔC :

$$\Delta C = C_a - C_n = 0. \tag{8}$$

Yet, the total travel cost decreases in f . Consequently, buying an autonomous car instead of a normal car imposes a positive ‘marginal external benefit’, MEB :

$$MEB = \Delta C - \frac{\partial TTC / \partial f}{N} = \delta \cdot \frac{N}{s} (1 - (r[f] + f \cdot \partial r / \partial f)) \geq 0. \tag{9}$$

The smaller $r[f] + f \cdot \partial r / \partial f$ is, the larger the increase in capacity following a marginal increase in f and the greater the positive externality.

To conclude, when autonomous cars only affect capacity, all users have the same cost and everyone gains by the same amount if the share with an autonomous car increases. Hence, getting an autonomous car causes a *positive* externality because it reduces the negative congestion externality. If cars are priced at marginal production costs, users will ignore the positive externality and there would be an under-consumption of autonomous cars.

3.4. Special case 2: autonomous cars only affect the values of time

Now we turn to autonomous cars only affecting the VOT, while capacity remains unaffected. Users of a normal car have a VOT of α . Users of an autonomous car have a VOT of $\theta \cdot \alpha$. The following constraint must hold: $\beta / \alpha < \theta < 1$.⁴ When some

⁴ The θ has to be above β / α as in dynamic congestion models the VOT must be above the value of schedule delay early (Arnott et al., 1990). The constraint is also intuitive: a VOT below β would imply that people prefer sitting in the car over getting out and reaching the destination to, e.g., work or visit family. Similarly, we assume that an autonomous car will decrease the VOT and thus $\theta < 1$. A θ equal to 1 leads to an uninteresting case with no effect on the VOT. We ignore $\theta > 1$: because an increased VOT would be strange as one could still switch off autonomous driving and because this would complicate the modelling.

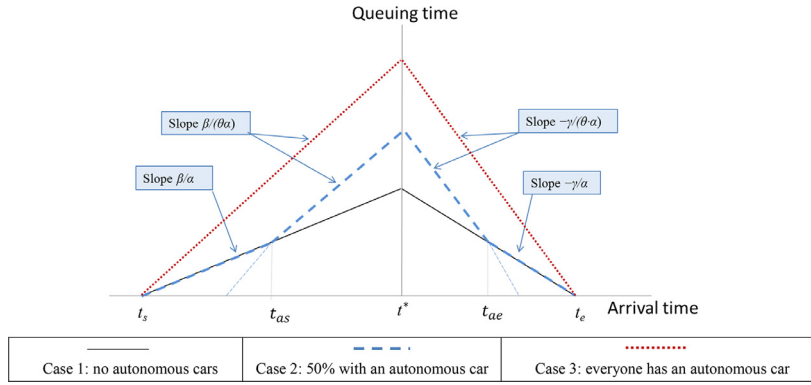


Fig. 3. Equilibrium queuing times when autonomous cars only affect the value of time

Note: For the 50% case, t_{as} is the moment autonomous cars start to pass the bottleneck and thus arrive at the destination; t_{ae} is the moment the last autonomous car arrives. The thin dashed lines are an out-of-equilibrium continuation of the iso-cost function for autonomous cars: autonomous cars do not arrive at these times, but the thin lines indicate what the travel times would need to be, although they are not, for them to have the same cost.

users have an autonomous car and others do not, the VOT is heterogeneous. A user who switches from a normal car to an autonomous car gets a lower VOT, thereby gaining a lower free-flow travel time cost and a lower bottleneck cost. The bottleneck cost equals the queuing time cost plus the schedule delay cost.

The share of drivers with an autonomous car affects external costs and total bottleneck cost by changing the heterogeneity in the VOT. We will refer to this as the ‘heterogeneity effect’, to distinguish it from the ‘capacity effect’ just discussed.

Fig. 3 illustrates this effect by depicting three situations: when nobody has an autonomous car (solid lines), when 50% of users do (dashed lines) and when everybody does (dotted lines). Note again that the queuing time functions shown can be interpreted as iso-cost lines. In the 50% case, autonomous cars arrive in the centre peak between t_{as} and t_{ae} . For this 50% case, the thin dashed lines denote an out-of-equilibrium continuation of the iso-cost function for autonomous cars. They indicate what the travel times would have to be, but are not, for autonomous cars to be willing to travel at those moments outside their equilibrium time window $[t_{as}, t_{ae}]$. Equilibrium travel times, however, are higher due to the flatter iso-cost lines for normal cars. The reverse applies for normal car users: the high equilibrium travel times in the centre peak $[t_{as}, t_{ae}]$ make that arrival time window less attractive than the two shoulder periods. When both car types are present, autonomous car owners self-select into the centre peak, and normal car owners self-select into the shoulders of the peak. The decreased VOT of autonomous car users means that they can enjoy the lower schedule delays in the centre peak as they care less about travel times than users of normal cars do.

Without a capacity effect, normal cars impose lower externalities than autonomous cars. The lowered VOT with an autonomous car means that, for its users to be in user equilibrium, they need a steeper development of travel time over time. Hence, such a user imposes longer travel times on all those who travel closer to the central peak (Lindsey, 2004). The marginal external cost of a user thus decreases with its VOT (Van den Berg and Verhoef, 2011a). When a driver’s VOT decreases due to acquiring an autonomous car, this raises the driver’s marginal external cost. The lower θ is, the more strongly an autonomous car reduces the VOT and therefore the larger the difference in the marginal external cost imposed by a normal and an autonomous car user.

Increasing the share of autonomous cars hurts existing autonomous car users by lowering the share of users who impose a lower queuing externality onto them. During the early arrival window of normal car users, the queuing time increases at a rate of β/α by arrival time; it changes by $-\gamma/\alpha$ for late arrivals. For autonomous car users, these slopes are steeper and equal $\beta/(\theta\alpha)$ before t^* and $-\gamma/(\theta\alpha)$ thereafter. Normal car users impose less steep travel time changes over arrival time than autonomous cars users. Hence, a normal car user imposes a lower queuing time cost on an autonomous car than a fellow autonomous car user does. Consequently, the iso-cost line for autonomous car drivers in Fig. 3 shifts upwards when their fraction increases.

Without a capacity effect, changing the fraction f has no effect on normal car users. They always travel in the shoulders of the peak and therefore will not face an upward pressure of equilibrium queuing times when someone switches car type. Because the duration of the peak is now independent of the share of autonomous cars, a normal car user stays on the same iso-cost line and thus keeps the same cost, no matter what the fraction of autonomous car drivers is.

Following Van den Berg and Verhoef (2011a, b), the travel costs for normal and autonomous cars respectively can be shown to be:

$$C_n = \delta \cdot \frac{N}{S} + \alpha \cdot TT_{ff}, \tag{10}$$

$$C_a = \delta \cdot \frac{N}{S} (\theta(1 - f) + f) + \alpha \cdot \theta \cdot TT_{ff}. \tag{11}$$

Eqs. (10) and (11) confirm that increasing the share f has no effect on the cost for a normal car driver, whereas a higher f raises the cost of existing users of autonomous cars.

For a user who switches from a normal car to an autonomous car, the travel cost decreases:

$$\Delta C = C_a - C_n = -\delta \frac{N}{s} ((1-f)(1-\theta)) - \alpha \cdot TT_{ff}(1-\theta) < 0. \quad (12)$$

The higher f is, the smaller the cost difference in favour of the autonomous car.

The total travel cost is:

$$\begin{aligned} TTC &= N(C_n(1-f) + C_a \cdot f) \\ &= \delta \frac{N^2}{s} (1+f(1-\theta)) + \alpha \cdot TT_{ff} \cdot N(1-f(1-\theta)). \end{aligned} \quad (13)$$

The derivative of the total travel cost with respect to f is:

$$\frac{\partial TTC}{\partial f} = (1-\theta)N \left(\delta \frac{N}{s} (2 \cdot f - 1) - \alpha \cdot TT_{ff} \right). \quad (14)$$

When $f \leq 1/2$, the total travel cost decreases with f . When f gets sufficiently far above $1/2$, the total cost may increase with f . The derivative $\partial TTC/\partial f$ increases with f : the more people who already have an autonomous car, the less socially beneficial increasing the share becomes. Still, if $\delta \cdot N/s < \alpha \cdot TT_{ff}$, the derivative $\partial TTC/\partial f$ remains negative for the possible range of f between zero and one. This condition implies that the bottleneck cost is relatively unimportant and the free-flow travel time is relatively important, and if this is the case, equipping another driver with an autonomous car always lowers the total travel cost.

Switching to an autonomous car lowers one's own cost, but it also imposes a negative externality by increasing the travel time of other autonomous cars. The $MEB \leq 0$ is:

$$MEB = \Delta C - \frac{\partial TTC/\partial f}{N} = -\delta \frac{N}{s} f(1-\theta) \leq 0. \quad (15)$$

Users will ignore this negative externality. If cars are priced at the marginal production cost, there will be over-consumption of autonomous cars. Conversely, with only a capacity effect, we found a positive externality and marginal cost pricing resulted in under-consumption. There is no a priori reason why either effect would dominate.

3.5. General case: autonomous cars affect capacity and value of time

Now we turn to the full setting, in which autonomous car ownership brings changes in the capacity and the VOT. The effects will prove to be a combination of those in the previous two subsections. Increasing the fraction with an autonomous car, f , hurts existing autonomous car users due to the heterogeneity effect. Yet, increasing this fraction also increases the total capacity and this lowers the travel cost for everyone. Normal car users always gain from increasing f due to the increased capacity.

Travel costs are a combination of those in the previous two subsections:

$$C_n = \delta \cdot \frac{N}{s} ((1-f) + f \cdot r[f]) + \alpha \cdot TT_{ff}, \quad (16)$$

$$C_a = \delta \cdot \frac{N}{s} (\theta(1-f) + f \cdot r[f]) + \alpha \cdot \theta \cdot TT_{ff}. \quad (17)$$

There are two opposing effects of raising the share f on the existing autonomous car users, which are in accordance with the separate effects in the previous two subsections. Even though an increased share raises average capacity, it may hurt existing autonomous car users as those who switch to an autonomous car will impose increased externalities, due to their altered departure time behaviour. Depending on which effect dominates, the travel cost in an autonomous car may increase or decrease with f . When $\theta > r[f] + f \cdot \partial r/\partial f$, raising f decreases the travel cost in an autonomous car. The capacity effect is stronger when $r[f] + f \cdot \partial r/\partial f$ is smaller, making an increased share more beneficial for all users. A smaller θ means that an autonomous car user needs a steeper travel time development over arrival time. Consequently, this user imposes longer travel times on those who arrive closer to t^* than the user. Hence, a smaller θ strengthens the heterogeneity effect, making increasing f less beneficial for current autonomous car users.

The iso-cost curves in Fig. 4 illustrate this. The figure shows equilibria for three penetration levels: 0%, 50% and 100% autonomous cars. It does so for two cases. The left panel has a weak capacity effect relative to the heterogeneity effect. The right panel has a relatively strong capacity effect. In the left panel, going from 50% to 100% autonomous cars *hurts* those who already have an autonomous car by shifting them to a *higher* iso-cost curve; in the right panel, these people *gain* because the relatively strong capacity effect brings them to a *lower* iso-cost curve.

When a user switches from a normal to an autonomous car, the user's travel cost decreases:

$$\Delta C = C_a - C_n = -\delta \frac{N}{s} ((1-f)(1-\theta)) - \alpha \cdot TT_{ff}(1-\theta) < 0. \quad (18)$$

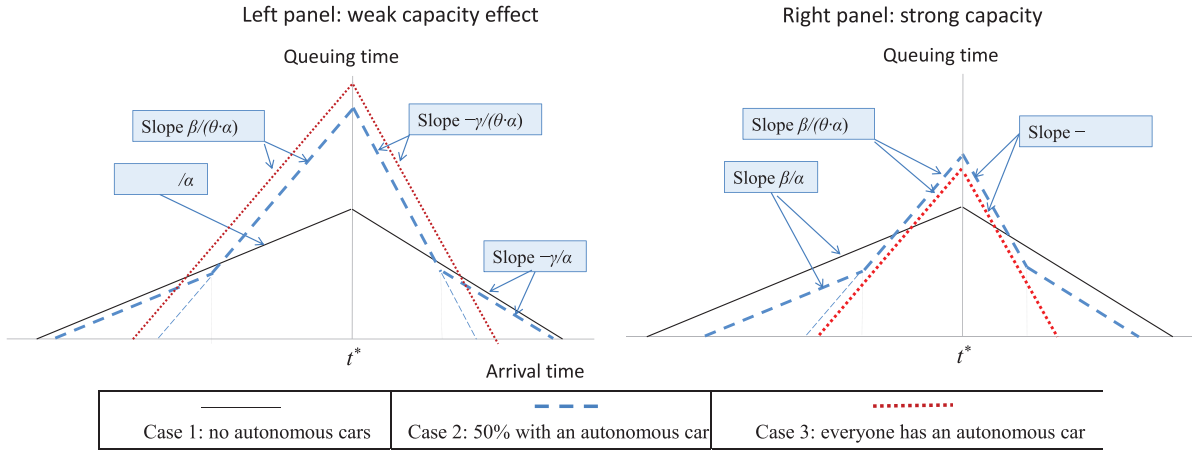


Fig. 4. Equilibrium queuing times when autonomous cars affect values of time and capacity.

Note: In the 50% case, the thin dashed lines are an out-of-equilibrium continuation of the iso-cost function for autonomous cars: autonomous cars do not arrive at these times, but the thin lines indicate what the travel times would need to be, although they are not, for them to have the same cost.

The larger f is, the smaller the difference in bottleneck costs between car types and thus the less switching to an autonomous car lowers one's own travel cost. The individual's cost gain from switching is independent of the capacity effect. The total travel cost is:

$$TTC = \delta \frac{N^2}{s} (1 + f^2(1 - \theta) - f(2 - r[f] - \theta)) + \alpha \cdot TT_{ff} \cdot N(1 - f(1 - \theta)). \tag{19}$$

Just as in the previous section, increasing f may increase or decrease the total cost.

The marginal external benefit due to switching to an autonomous car may be negative or positive:

$$MEB = \Delta C - \frac{\partial TTC / \partial f}{N} = \delta \frac{N}{s} (1 - f(1 - \theta) - (r[f] + f \cdot \partial r / \partial f)). \tag{20}$$

It is negative when the heterogeneity effect dominates the capacity effect, which occurs when $f(1 - \theta) > 1 - (r[f] + f \cdot \partial r / \partial f)$.⁵ The MEB is more likely to be negative when the fraction f is larger as then there are more autonomous car users who are hurt by increasing the share. A smaller θ means that getting an autonomous car decreases the VOT more, which strengthens the heterogeneity effect and thereby lowers the marginal external benefit. When $r[f] + f \cdot \partial r / \partial f$ is smaller, switching to an autonomous car brings a larger capacity gain and thus the MEB is higher. Eq. (20) shows how these different forces jointly determine the overall MEB .

4. Provision regimes in the full model

Now we turn to the equilibrium outcome under three provision regimes of autonomous cars: 1) marginal cost pricing, 2) socially optimal public supply and 3) profit-maximising monopoly provision. The provision at marginal cost leads to over-consumption of autonomous cars if there is a negative externality from their purchase, and to under-consumption under a positive externality. Under private provision, as the market becomes more competitive, the outcome will approach marginal cost provision. The other extreme of private provision is the private monopoly.

The total cost, TC , is the sum of the total travel cost, TTC , and the total automobile cost due to buying autonomous cars instead of normal cars, $MC_a \cdot f \cdot N$.⁶

$$TC = TTC + MC_a \cdot f \cdot N. \tag{21}$$

With an interior equilibrium, the travel cost reduction, ΔC , from switching to an autonomous car is exactly offset by the sum of the per trip extra marginal cost and the mark-up of an autonomous car, such that no one wants to switch car type. For a corner solution without autonomous cars, the price for autonomous cars should be above that of normal cars and vice versa for a corner solution without normal cars. As we will see in the numerical analysis, corner outcomes are likely to occur.

⁵ This condition for a positive MEB is less stringent than the earlier condition for $\partial C_a / \partial f < 0$ as now we also look at the effect on normal car users who always gain from increasing f .

⁶ As the mark-up payments to the firm are a transfer and not a cost, the total cost does not contain MU_a . The marginal production cost of normal car is normalised to zero.

4.1. Marginal cost provision

Perfect competition leads to marginal cost provision and thus to a zero mark-up. Assuming an interior solution, this leads to an equilibrium share of autonomous car users, f^{MC} , of:⁷

$$f^{MC} = 1 - \frac{S}{\delta N(1-\theta)} (MC_a - \alpha(1-\theta) \cdot TT_{ff}). \quad (22)$$

However, often, there is a corner solution. For instance, when $MC_a < \alpha(1-\theta)TT_{ff}$, Eq. (22) would predict $f^{MC} > 1$ and thus we are in the corner outcome $f^{MC} = 1$. Under marginal cost supply, individual users ignore the capacity effect as it does not affect the cost difference between car types.⁸

4.2. Public provision

Public provision minimises the total social cost by finding the optimal fraction, f , under the second-best limitation that there is no direct congestion pricing. Corner solutions are likely. We assume that capacity convexly increases with f . This makes it likely that the total cost is globally concave in f , and if this is the case, the minimum is always at a corner.⁹ But even besides this, it often occurs that the total cost always decreases with the share, implying an optimum at the corner $f^{pub} = 1$.¹⁰ For an interior solution, setting $\partial TTC/\partial f$ to zero implies:^{10F}

$$MU_a = \frac{\partial TTC/\partial f}{N} - \Delta C = -MEB, \quad (23)$$

$$f^{pub} = \frac{1}{2(1-\theta) + \partial r/\partial f} \left((2 - r[f^{pub}] - \theta) - \frac{S}{N\delta} (MC_a - \alpha(1-\theta)TT_{ff}) \right). \quad (24)$$

Therefore, the public operator imposes a negative or positive mark-up to correct for the externality, where a subsidy applies when MEB is positive.¹¹

One may expect f^{pub} to exceed f^{MC} because of a dominating positive externality. Yet this is not always true. The net externality is zero when $f(1-\theta) - r[f] + f \cdot \partial r/\partial f = 0$, and in this case the shares with public and private supply are equal. When this equation is negative, the heterogeneity effect dominates the capacity effect and f^{pub} is below f^{MC} . In our numerical analysis, public provision most often leads to the corner outcome of $f^{pub} = 1$, and marginal cost supply may lead to an undersupply. For interior solutions of our numerical model, however, oversupply with marginal cost pricing is most common and undersupply is rare. Parameterisations that lead (22) and (24) to imply that the public provider's share is above the marginal cost supplier's share also tend to imply that we will be in a corner outcome $f^{pub} = 1$.

The public supplier has a larger share compared to the marginal cost supplier when r is lower. This is because the public supplier then wants to take advantage of the stronger capacity effect, whereas marginal cost provision ignores the capacity effect. When MC_a is smaller or $\alpha(1-\theta)TT_{ff}$ larger, the autonomous car reduces costs more and the share with both regimes increases, but this effect is smaller with a public supplier as it also considers the car purchase externality.

To conclude, public provision often leads to the corner outcome of 100% autonomous cars. However, for interior solutions, marginal cost supply can lead to over- or undersupply.

4.3. Provision by a private monopolist

The monopolist maximises its profit, which equals the mark-up per trip multiplied by the number of autonomous cars:

$$\Pi = MU_a \cdot f \cdot N = (-\Delta C - MC_a) \cdot f \cdot N. \quad (25)$$

The ΔC gives how much higher the travel cost of an autonomous car is than of a normal one and thus $-\Delta C$ measures the willingness to pay for an autonomous car over a normal one.

⁷ As ΔC decreases with f , there is one unique solution for f where $\Delta C + MC_a = 0$.

⁸ Assuming $f^{MC} < 1$, the f^{MC} is higher when: i) $\delta \cdot N/s$ is larger, ii) when $\alpha \cdot TT_{ff}$ is larger, and iii) when θ is smaller. All these changes mean that switching to an autonomous car reduces one's own travel cost more and thus such cars are more attractive to buy. The f^{MC} understandably decreases with the extra cost of autonomous cars, MC_a .

⁹ The first and second order derivatives of total cost with respect to f are:

$$\frac{\partial TC}{\partial f} = \delta \frac{N^2}{s} (2 \cdot f(1-\theta) - 2 + r[f] + f \cdot \partial r/\partial f + \theta) - (1-\theta)\alpha \cdot TT_{ff} \cdot N + MC_a \cdot N, \quad \frac{\partial^2 TC}{\partial f^2} = \frac{\partial^2 TTC}{\partial f^2} = \delta \frac{N^2}{s} (2(1-\theta) + 2 \cdot \partial r/\partial f + f \cdot \partial^2 r/\partial f^2).$$

For an interior optimum to be possible, $1 - \theta + (\partial r/\partial f + \partial^2 r/\partial f^2 \cdot f/2) > 0$ must hold as otherwise the total cost is not strictly convex in f . As $0 < \theta < 1$, this happens when r does not overly depend on f and is not too concave.

¹⁰ When the total cost is strictly convex in f , the other condition that must hold for an interior optimum is:

$$r[f=1] + \partial r/\partial f[f=1] - \theta > \frac{(1-\theta)\alpha \cdot TT_{ff} - MC_a}{\delta N/s} > \theta - 1,$$

as this ensures that the derivative $\partial TC/\partial f$ is negative for $f=0$ and is positive for $f=1$

¹¹ The f^{pub} is higher when: i) the capacity effect is stronger (i.e. smaller r), ii) the capacity function is more curved (i.e. more negative $\partial r/\partial f$), iii) MC_a is lower, iv) $\alpha \cdot TT_{ff}$ is higher and v) $\delta \cdot N/s$ is lower. Unlike with marginal cost pricing, the effect of θ on f is now ambiguous.

For interior solutions, the optimal share for a private monopolist is half of that under marginal cost provision, where the latter already often leads to undersupply:¹²

$$f^{mon} = \frac{1}{2} \left(1 - \frac{s}{\delta N(1-\theta)} (MC_a - \alpha(1-\theta)TT_{ff}) \right) = \frac{1}{2} f^{MC}. \quad (26)$$

The corner solution $f^{mon} = f^{MC} = 0$ occurs in both regimes when the travel cost reduction from owning an autonomous car is so small that nobody wants one.

The monopolist's share may increase or decrease with θ . A lower θ means that switching to an autonomous car lowers one's own travel cost more, thereby raising the willingness to pay for an autonomous car and this makes raising the share more profitable. However, a lower θ also raises the externality that autonomous cars impose on each other, which makes raising the share less profitable. The monopolist therefore internalises the negative heterogeneity externality that autonomous cars impose on each other. Under fixed demand, the monopolist ignores the positive capacity externality as this affects normal and autonomous car users equally and does not change the willingness to pay.

The monopolist almost invariably has a much lower share than the public supplier does. For interior solutions, the monopolist has a lower share unless autonomous cars have no effect on capacity. The corner outcomes without undersupply of $f^{pub} = f^{mon} = 0$ and $f^{pub} = f^{mon} = 1$ are unlikely. The outcome $f^{mon} = 1$ needs a rather negative MC_a and this tends to require that autonomous cars must be cheaper to produce than normal cars. The outcome $f^{pub} = 0$ is also unlikely as this requires autonomous cars to be very expensive or the effects on the VOT and the capacity to be very small. Ignoring implausible corner outcomes, the monopolist supplies too few autonomous cars from the societal perspective.

5. Numerical example

This section presents numerical results for two base calibrations of the model developed above: one for the USA and one for the Netherlands. The latter case may be viewed as representative for European countries. Compared to the USA, the Netherlands has a lower recommended VOT and a lower mileage before a car is sold on the second-hand market. The Netherlands has much higher fuel prices, which are partly offset by a higher fuel efficiency. These differences in parameterisation will turn out to be important. We will find that, for the USA calibration, marginal cost provision and public provision lead to the same outcome, while for the Dutch case, marginal cost pricing leads to too few autonomous cars.

After discussing these base cases, we turn to the sensitivity analyses. The model outcomes are sensitive to the parameterisation and hence it is important to present these results.

5.1. Calibration of the numerical models

We focus on petrol passenger cars. For both countries, the schedule delay parameters will be based on the ratios $\beta/\alpha = 39/64$ and $\gamma/\alpha = 1521/640$ established in Small (1982), as is common in the literature. As in Van den Berg and Verhoef (2011a, b), we use $N=9000$ and $s=3600$. We consider a trip of 20 km, with a free-flow travel time of 20 minutes. To ensure that the regular equilibrium of the bottleneck model holds, the VOT with an autonomous car must be greater than the value of schedule delay early and thus $\theta > \beta/\alpha = 39/64 \approx 0.609$. Hence, θ can be between 0.61 and 1 and the sensitivity analyses will indeed vary θ over this entire range. Our base value is $\theta=0.8$, so that the VOT effect is in the middle of the possible range.

Following the average prediction from the literature, we assume that switching from all normal cars to all autonomous cars doubles capacity. Many studies have found that the capacity effect is non-linear. Following Tientrakool et al. (2011), we assume that going from 90% autonomous cars to 100% leads to half of the total increase in capacity and that capacity monotonically increases in f . Fitting a polynomial leads to $r=1-0.5 \cdot f^{3.85}$. This relationship is illustrated in Fig. 5.

For the USA, we use a VOT of \$18.82/h. This is the recommended value (US Department of Transportation, 2011) updated using 2013 income data (US Census Bureau, 2014). Our base case USA value of MC_a is 1.13. For the USA in 2013, the fuel expenditure per mile was 13¢¹³ \$0.13, implying a fuel efficiency of 26 miles per gallon¹⁴ or 11 km per litre. Let us assume that an autonomous car has a 20% higher fuel efficiency than a normal car. This implies a fuel cost saving for an autonomous car of \$0.31 per trip. We further assume that the extra cost to purchase an autonomous car is \$7500, when also considering the difference in residual value when reselling the car on the second-hand market. We use that in 2013 a car was driven on average 13,476 miles per year.¹⁵ We assume that a new car is sold after 4.5 years. Then, the extra depreciation cost for an autonomous car becomes \$1.54 per trip. This results in an extra automobile cost per trip of \$1.13.

¹² The first-order and second-order conditions of maximisation hold for an interior solution and there is one unique optimum.

¹³ www.rita.dot.gov/bts/sites/rita.dot.gov/bts/files/publications/national_transportation_statistics/html/table_03_17.html (accessed on 24 February 2015).

¹⁴ Using a 2014 average price of \$3.43 per gallon from www.eia.gov/dnav/pet/pet_pri_gnd_dcus_nus_a.htm (accessed on 14 February 2015).

¹⁵ www.fhwa.dot.gov/ohim/onh00/bar8.htm (accessed on 24 February 2015).




Fig. 5. The $r[f]$ function (left) and effective capacity for autonomous cars (right) for the numerical models.

For the Netherlands, the recommended VOT is €10 (Kouwenhoven et al., 2014), and following a similar calculation as for the USA, an MC_a of €1.51 seems reasonable.^{16,17}

5.2. Base case

Table 2 shows the outcomes¹⁸ for the base calibrations for four cases: no autonomous cars, perfectly competitive provision at marginal production cost, second-best public provision that minimises total social cost, and monopoly provision that maximises the profit of the autonomous car producer. TTC is the total travel cost and TC is the total cost including the automobile cost; $\% \Delta TTC$ and $\% \Delta TC$ are the corresponding percentage changes in totals from the case without autonomous cars. Relative efficiency is the total cost saving of a policy from the case without autonomous cars divided by the cost reduction under socially optimal public provision.

The introduction of autonomous cars lowers costs. Even with a private monopolist, users are slightly better off than without autonomous cars. Of course, users can never be worse off than before as they can always all choose to stick to the normal car. The monopolist asks a substantial mark-up per trip of \$2.34, which implies that the profit is 60% of revenue. The costs are higher for the USA than for the Netherlands because the USA has a higher VOT.

For the USA, marginal cost provision leads to the same corner solution of 100% autonomous cars as public provision. The autonomous car has a lower cost (i.e. travel cost plus automobile cost) than the normal car and hence everybody wants one. For the Netherlands, marginal cost provision results in too few autonomous cars. This MC regime attains only 15% of the societal cost savings from public provision. The private monopoly attains a much lower cost saving in both countries. Partly responsible for these low-cost savings from private supply is the non-linear capacity effect: the smaller market shares with these regimes mean that capacity is much lower. Yet, even under a constant capacity effect, monopoly supply fares much worse than public supply, while also MC supply often does worse than public supply (see Section 5.5 below).

For the USA, a realistic lower bound for MC_a seems to be $-\$0.75$, using a fuel efficiency gain of 50%, an extra purchase cost of \$1000 and a usage of 5.5 years. A reasonable upper bound seems to be $MC_a = \$4$, using a 5% higher fuel efficiency, an extra purchase cost of \$15,000 and 3 years of usage. Under the lower bound case of $MC_a = -\$0.75$, the outcome is similar to that in the base USA case: marginal cost pricing and public provision lead to 100% autonomous cars. Monopoly supply has a similar relative efficiency as in Table 2, although the fraction with an autonomous car is larger. For the upper bound case of $MC_a = \$4$, public supply still leads to 100% autonomous cars, but this requires a large subsidy. Accordingly, autonomous cars need to be very costly for it not to be socially worthwhile for everyone to get one if the capacity effect is as strong as assumed.

¹⁶ The average daily travel distance for people with a car is 41 km (CBS, 2015a). We assume an extra purchase cost (minus residual value) of €5,000 and that the car will be sold after 3.5 years, which is lower than for the USA as a large part of new Dutch cars are leased for a few years only. This results in a depreciation cost per trip of €1.91. Fuel savings would be €0.40 per trip, using an average fuel price for 2014 of €1.695/litre (CBS, 2015b), a 14 km/litre fuel efficiency for normal cars (<http://gemiddeldelgezien.nl/meer-gemiddelden/68-gemiddeld-verbruik-auto> as accessed on 14 February 2015), and a 20% higher fuel efficiency for an autonomous car.

¹⁷ In our welfare evaluation, we ignore that a part of the fuel price is tax. Including this would complicate the modelling but would not change the individual's choice. It may affect the choice of the public supplier, but in our base cases it would not do so as the component for fuel in the automobile cost is small and the public optimums will turn out to be deep in the corner solution of 100% autonomous cars anyway as eq. (24) predicts shares above 150%. Ignoring the tax part does affect the cost savings, but even here the effect is small. For instance, in the Netherlands, taxes form about 64% of the fuel price (www.nl.wikipedia.org/wiki/Benzineaccijns as accessed on 15 April 2015) and removing the reduction in tax payments from the cost saving would reduce the total cost savings in the public case by only 3.7%.

¹⁸ The results were calculated in Mathematica 9. For marginal cost and monopoly supply, we used the analytical conditions and then checked if a corner solution occurred. For public supply, the objective may not be globally convex. Therefore, we used the command 'minimize' to find the global minimum as the total cost is a polynomial function of f .

Table 2
Outcomes under the base calibrations for the USA and the Netherlands.

	USA				The Netherlands			
	Only normal cars	Marginal cost	Public	Monopoly	Only normal cars	Marginal cost	Public	Monopoly
Share with an autonomous car	0	1	1	0.514	0	0.652	1	0.326
Total travel cost (<i>TTC</i>)	261,839	147,857	147,857	241,719	139,128	123,392	78,564	132,136
% Δ TTC	–	–44%	–44%	–8%	–	–11%	–44%	–5%
Total cost (<i>TC</i>)	261,839	158,027	158,027	246,943	139,128	132,256	92,154	136,568
% Δ TC	–	–40%	–40%	–6%	–	–5%	–34%	–2%
Relative efficiency [#]	0	1	1	0.14	0	0.15	1	0.05
Cost per trip for a normal car	–	17.68	17.68	28.64	–	14.70	9.40	15.43
Cost per trip for an autonomous car	29.09	17.56	17.56	28.64	15.46	14.70	9.40	15.43
Mark-up	–	0	0	2.34	–	0	–0.84	0.79
Corner solution	–	yes	yes	no	–	no	yes	no

Note: Relative efficiency is the welfare gain of a policy relative to the case without autonomous cars divided by the gain from public provision.

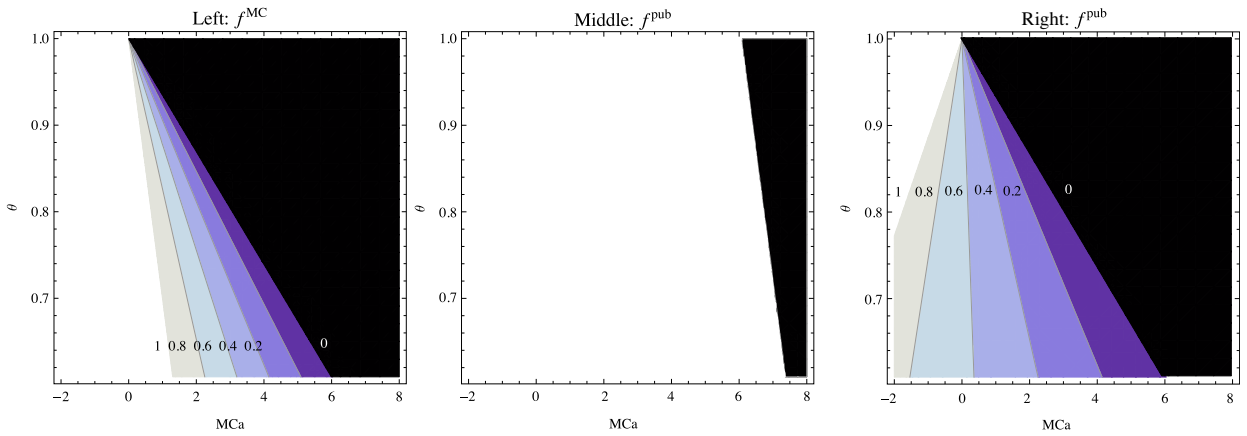


Fig. 6. Effect on share of autonomous cars of MC_a and θ : the left panel is for marginal cost provision, the middle panel for public supply and the right panel for monopoly provision.
 Note: The darker the area of a contour plot is, the lower the share with an autonomous car; a white area is for $f=1$ and a black area for $f=0$.




Fig. 7. Effects of MC_a and θ on the percentage change in total cost with only normal cars.

5.3. Sensitivity with respect to θ and MC_a

There is little to no guidance from the literature on the values of MC_a and θ . There is also a very wide range of predictions for the size and shape of the capacity effect. Therefore, it is vital to do extensive sensitivity analyses. The effects of the other parameters are in line with the theoretical discussion. Moreover, for these parameters we have much more guidance from the literature. Hence, these parameters will not be discussed further here. Section 5.4 will vary the maximum capacity $s/r[f=1]$ for the same general shape of the $r[f]$ function. Section 5.5 will vary this shape. This subsection varies the VOT reduction parameter, θ , from its theoretical minimum of β/α to 1. It also varies the extra marginal automobile cost of an autonomous car, MC_a , from -2 to 8 . As fuel savings can only be so large, a MC_a below -1 tends to imply that autonomous cars are cheaper to produce than normal cars. For The Netherlands, the base case values were $MC_a=1.51$ and $\theta=0.8$.

Fig. 6 shows the equilibrium share with an autonomous car for marginal cost (left panel), public (middle panel) and monopoly provision (right panel). The darker the area of a contour plot is, the lower the share with an autonomous car. A white area is for the corner solution $f=1$ and a black area for $f=0$. The graph is for the Dutch VOT of €10, with the other parameters being the same for both cases. Using the VOT for the USA leads to qualitatively equivalent pictures.

In θ vs MC_a space, public provision only leads to corner outcomes. It leads to 100% autonomous cars unless they are very expensive to make, with $MC_a > 6$. This would imply that autonomous cars cost around €17,500 more than normal cars after considering the difference in residual value as well.

Marginal cost supply also tends to lead to a large share with an autonomous car. If MC_a and θ are sufficiently low, marginal cost and public provision both lead to 100% autonomous cars. Still, for higher values of MC_a , marginal cost provision leads to a below optimal share. Monopoly provision leads to an even larger undersupply and welfare loss. Only when $MC_a < 0$ and when θ is sufficiently large does monopoly provision not lead to too few autonomous cars because this implies that all regimes lead to 100% autonomous cars. For interior solutions, the monopoly leads to half as many autonomous cars as marginal cost pricing. Marginal cost and monopoly provision lead to 0% autonomous cars for exactly the same black parameter range as in this range getting an autonomous car always increases the user's cost. In a part of this range, the public provider still ensures that everyone gets an autonomous car because there is a large positive externality.

Fig. 7 compares the changes in total costs that the regimes attain. In this sensitivity analysis, monopoly supply always performs worse than public supply and often leads to a large welfare loss. Marginal cost supply can also lead to substantial

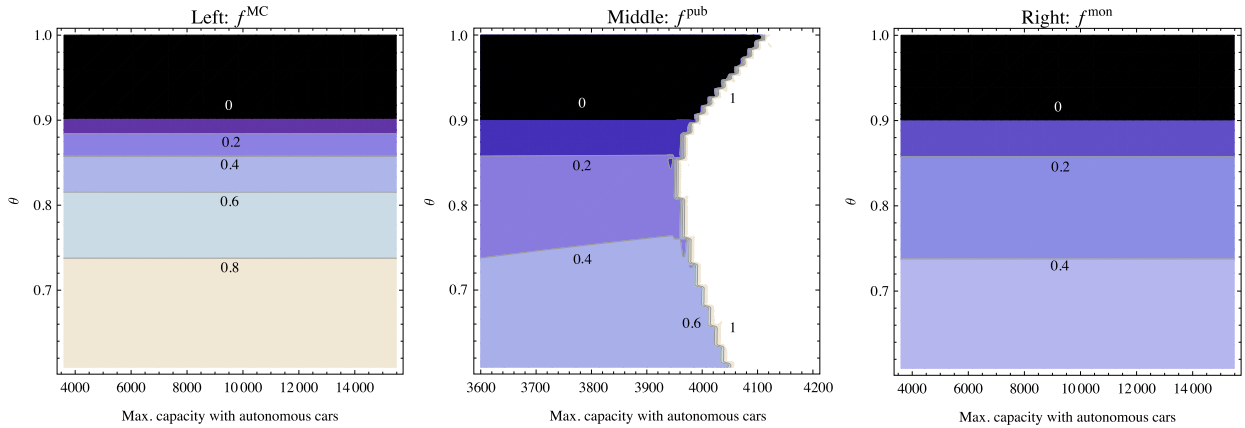


Fig. 8. Effect of maximum capacity and θ on equilibrium shares of autonomous cars: the left panel is for marginal cost provision, the middle panel for public provision and the right panel for monopoly provision.
 Note: The darker the area of a contour plot is, the lower the share with an autonomous car. A white area is for $f=1$ and a black area for $f=0$. For public supply, the range of $s/r[f=1]$ is much smaller than for the other regimes as for $s/r[f=1] > 4120 \approx 1.14 \cdot s$ public provision always leads to $f=1$.

welfare losses, even when the difference in shares is small. For a slightly positive or negative MC_a and thus autonomous cars not being too expensive, marginal cost supply leads to the same corner outcome of $f=1$ as public supply.

5.4. Sensitivity analysis with respect to the maximum increase in capacity

Next, we look at the impact of the maximum increase in capacity due to autonomous cars by varying $s/r[f=1]$. In the base case, the maximum capacity was $2 \cdot s$. We range the maximum capacity from s to $4.4 \cdot s$. We again also vary θ . The power of the polynomial defining r is kept constant at 3.85.

Fig. 8 depicts the shares with an autonomous car. Consistent with the analytics, $s/r[f=1]$ has no effect on the share under marginal cost or monopoly provision as these regimes ignore the capacity effect. The maximum capacity does not affect the shares under these regimes, although it does affect the total costs.

For public provision, we show a much smaller range of the maximum capacity as the only area of interest is for small values. For $s/r[f=1] > 4120 \approx 1.14 \cdot s$, so for a maximum increase of more than 14%, public provision always leads to $f=1$. Only when the capacity effect is very small is it optimal for not everyone to be equipped. Still, such a small increase is within the possible range.¹⁹

Public supply only leads to 0% autonomous cars when the effect on the VOT is limited (i.e. $\theta > 0.9$) and the capacity effect is very small (i.e. $s/r[f=1] < 4120$). Monopoly supply leads to an undersupply unless a corner outcome $f^{pub} = 0$ holds. Still, the difference in shares is minor when the maximum capacity is below 3950. Consistent with the theory, the difference between f^{mon} and f^{pub} goes towards zero as maximum capacity decreases towards $s=3600$. When the VOT effect is not too weak, with $\theta < 0.9$, and the maximum capacity is below 3950, marginal cost pricing leads to an oversupply that can be quite large. Still, for values of the maximum capacity that seem more reasonable—say between $1.25 \cdot s$ and $2.5 \cdot s$ —public supply leads to 100% autonomous cars; marginal cost supply often does the same or leads to a small undersupply. Still, even a small undersupply can lead to a large loss in welfare.²⁰

5.5. Sensitivity analysis with respect to the shape of the capacity function

This final sensitivity analysis looks at the effects of the shape of capacity function. It does so by varying the power, m , of the polynomial:

$$r[f] = 1 - 0.5 \cdot f^m. \tag{27}$$

For a power of $m=0$, we have a constant capacity effect that is independent of f . For $m \geq 1$, the higher m is, the more convex the capacity $s/r[f]$ is in f . Our base case had $m=3.85$ and thus capacity was highly convex. No matter what the

¹⁹ At no point are the parameter areas of the $f^{pub}=0$ and $f^{pub}=1$ corner solutions connected; there is always a (small) area between them where the interior outcome smoothly goes from 0 to 1.

²⁰ We repeated this analysis with a lowered MC_a of -2 and with an increased MC_a of 7 , whereas here MC_a was 1.51 . When $MC_a = -2$, getting an autonomous car always reduces one's own cost. Hence, marginal cost provision always leads to 100% autonomous cars. When $MC_a = 7$, getting an autonomous car always increases one's own cost, and hence both marginal cost pricing and monopoly supply lead to 0% autonomous cars. Still, if the capacity effect is not too weak, there is a positive externality and the socially optimal share may be positive. Besides these changes, the effects of the parameters are qualitatively the same as in the main sensitivity analysis.




Fig. 9. Relative efficiencies of marginal cost and monopoly supply over the convexity of the capacity function as determined by the power of $r[f]$ in Eq. (27). Note: Relative efficiency is the welfare gain of a policy from the case without autonomous cars divided by the gain from public provision. From a power of 1, the higher the power of $r[f]$ is, the more convex the capacity for autonomous cars.

power is, it is socially optimal to have 100% autonomous cars when the MC_a is low enough and the maximum capacity for autonomous cars is high enough. Because the maximum capacity is independent of m , the total cost under public supply is independent of the power m . As Section 4 proved, the shares under marginal cost and monopoly supply are also independent of the capacity function as these providers ignore the capacity effect.

The shape of the capacity function can hence only affect travel costs under marginal cost or monopoly supply. The higher the power in (27) is, the lower the capacity for autonomous cars when $f < 1$, and the lower the welfare gain from a regime that leads to a below socially optimal share. Fig. 9 shows this by giving the relative efficiency of the two private regimes over the power m . It does so for different levels of θ . The relative efficiency of a regime equals its total cost saving from the initial setting with only normal cars relative to the cost savings with public supply. The sensitivity analysis is again for the Dutch case. The pictures for the USA look qualitatively the same.

To conclude, the shape of the capacity function can greatly affect the welfare effects of the regimes. However, the shape never affects the equilibrium fraction of autonomous cars with marginal cost and monopoly provision. For our parameterisations, the shape also has no effect on the fraction with public supply as this regime always leads to the corner outcome of 100% autonomous cars.

6. Model extensions

We have used what is arguably the simplest dynamic model to investigate the effects of autonomous cars via increased capacity, a lower VOT and the resulting VOT heterogeneity. There are a number of logical extensions worthy of mentioning here and it seems worthwhile to explore qualitatively how some of these might affect our results.

First, it seems plausible that the extent to which getting an autonomous car reduces the VOT will differ among people. Some office workers may spend the freed-up time preparing documents or checking email, becoming nearly as productive as in the office. Manual workers may well face some difficulty in using the freed-up time working, but may spend the travel time on leisure activities. A businessperson may already be using the time in a normal car in a productive way and hence buying an autonomous car may bring modest benefits. In any case, the parameter θ is likely to vary over individuals. This would mean that car buyers self-select, with those for whom an autonomous car reduces travel costs more being more likely to get one. This also suggests that increasing the fraction with an autonomous car may be less beneficial. For a high enough share, the user who switches has a relatively small reduction in VOT and thus a small cost saving. A heterogeneous θ gives marginal cost and monopoly supply an extra disadvantage from the efficiency viewpoint: these suppliers only consider the marginal user in deciding their f , whereas the public supplier considers the average autonomous car user who has a lower θ_i than the marginal user. This issue is akin to the classical issue of quality provision by a private firm in Spence (1975).

There may also be pre-existing heterogeneity in the values of time and/or schedule delay. Then, users arrive ordered by the heterogeneous ratios β_i/VOT_i and γ_i/VOT_i , which are now heterogeneous for two reasons: pre-existing heterogeneity and differences in car type. Pre-existing heterogeneity may lessen the heterogeneity effect as autonomous cars may add little in heterogeneity. It may also strengthen the heterogeneity effect as the last users to switch to an autonomous car will tend to have the lowest reduction in travel cost. Pre-existing heterogeneity means that some normal and autonomous car users may have the same ratio β_i/VOT_i or γ_i/VOT_i , implying that they travel mixed and reach the bottleneck at the same time. As car types now no longer travel separated in time, this may reduce the capacity benefit of autonomous cars.

An interesting form of pre-existing heterogeneity is 'proportional heterogeneity', which leads to variation in α_i , β_i and γ_i in fixed proportions and could be argued to be due to differences in income (Vickrey, 1973; Van den Berg and Verhoef, 2011b; Van den Berg 2014; Hall, 2015). Such heterogeneity would not affect the order in which users arrive, but would mean that those with higher incomes may have lower marginal utilities of income and higher travel costs. Hence, people with higher incomes would be more likely to buy autonomous cars and may gain more from doing so. This implies important distributional effects from introducing autonomous cars.

Finally, we assume that total demand is fixed and we ignore other modes. If autonomous cars reduce the price of car travel for a given number of users, this induces more demand and attracts users away from other modes such as public transport. If the travel cost by public transport decreases with the amount of its use, for example, due to economies of scale or frequency benefits, a Downs-Thomson paradox could occur.²¹ If this is the case, a mode shift towards car travel raises the cost of travel by public transport and thus even more public transport users will switch to the car. This then also increases the equilibrium price for car travel and hence everyone is worse off. For related paradoxes, see [Arnott and Small \(1994\)](#).

7. Conclusion

We have investigated the effects of introducing autonomous cars on the costs of travel via a number of distinct channels. These include an increase in the capacity of roads, a decrease in the values of time (VOTs) and the resulting heterogeneity in VOTs. The share of users with an autonomous car was endogenous as we incorporated the equilibrium of the car purchase market. We considered three provision regimes: perfect competition leading to marginal cost pricing, second-best public provision, and monopoly supply.

Buying an autonomous car instead of a normal car raises road capacity and thus imposes a positive externality. However, it also lowers a user's VOT. As under heterogeneity the bottleneck congestion externality decreases with a user's VOT, buying an autonomous car also imposes a negative externality. The net externality may be positive or negative. The numerical analyses suggest that a net positive externality is most likely, especially when the capacity for autonomous cars increases with their market share.

If there is a positive externality due to buying an autonomous car, marginal cost pricing tends to lead to under-consumption of autonomous cars. To prevent this and attain the second-best optimum, the public supplier needs to provide a subsidy. However, if there is a negative externality, a corrective tax is needed to prevent over-consumption. The private monopolist is likely to lead to a large undersupply and welfare loss.

The results are sensitive to the model's crucial parameters. For instance, in the USA base case, public and marginal cost supply both lead to 100% autonomous cars, whereas, for the Dutch base case, marginal cost supply leads to only 76% autonomous cars and public supply to 100%. There is great uncertainty about these parameter values and this makes investigating them an important future research topic. The inclusion of pre-existing heterogeneity in multiple dimensions and the analysis of the resulting distributional effects also seem interesting and important extensions.

We only consider one supplier, but in reality, there will be multiple suppliers of autonomous cars who may also supply normal cars. Will these suppliers face incentives to allow their autonomous cars to cooperate or would it be more profitable to keep them incompatible? And what about supporting infrastructure for autonomous cars along the road? Will each car supplier need to build its own? The introduction of autonomous cars will have great effects on public transport, taxi transport and parking demand. In the long run, it may even affect the structure of the city. These questions deserve attention in future research efforts.

Finally, we have only considered congestion and ignored other externalities, such as pollution and safety. Including these would raise the marginal external benefit of purchasing an autonomous car as they are often argued to be safer and more fuel-efficient. For safety, it may be that this benefit increases with the market share of autonomous cars, just as we also assumed for the capacity benefit. The analytical analysis core is likely to remain unchanged by including these other externalities. However, the parameterisation and the calibration of the safety function would be important extensions to increase the completeness of our work.

Acknowledgements

We would like to thank the editor, Robin Lindsey, and reviewers for their extensive and helpful comments. We thank the participants of the OPTION conference (2015), the ITEA conference (2015) and the Eureka seminar in Amsterdam (2015). We thank Sylvia Bleker for her valuable comments. Financial support from the ERC (AdG Grant #246969 OPTION) is gratefully acknowledged. Any errors are ours.

References

- Anderson, J.M., Kalra, N., Stanley, K.D., Sorensen, P., Samaras, C., Oluwatola, O.A., 2014. Autonomous Vehicle Technology: A Guide for Policymakers. Rand Corporation, Santa Monica, CA, USA Accessed from www.rand.org/pubs/research_reports/RR443-1.html.
- Arnott, R., de Palma, A., Lindsey, R., 1988. Schedule delay and departure time decisions with heterogeneous commuters. *Transp. Res. Rec.* 1197, 56–67.
- Arnott, R., de Palma, A., Lindsey, R., 1990. Economics of a bottleneck. *J. Urban Econ.* 27 (1), 111–130.
- Arnott, R., de Palma, A., Lindsey, R., 1993. A structural model of peak-period congestion: a traffic bottleneck with elastic demand. *Am. Econ. Rev.* 83 (1), 161–179.
- Arnott, R., de Palma, A., Lindsey, R., 1998. Recent developments in the bottleneck model. In: Button, K.J., Verhoef, E.T. (Eds.), *Road Pricing, Traffic Congestion and the Environment: Issues of Efficiency and Social Feasibility*. Edward Elgar, Aldershot, pp. 79–110.
- Arnott, R., Small, K.A., 1994. The economics of traffic congestion. *Am. Sci.* 82 (5), 446–455.
- CBS, 2015a. Mobiliteit in Nederland; persoonskenmerken en motieven, regio's. Accessed from statline.cbs.nl/Statweb/publication/?DM=SLNL&PA=81124NED&D1=a&D2=0&D3=0&D4=1-3&D5=0-8&D6=0&D7=a&VW=T on 21 February 2015.

²¹ As discussed by [Zhang et al. \(2016\)](#), this assumes that total travel demand is fixed, automobile and public transport are perfect substitutes and public transport is not congested or crowded. Relaxing these assumptions often eliminates the paradox.

- CBS, 2015b. Pompprijzen motorbrandstoffen; locatie tankstation, brandstofsoort. Accessed from statline.cbs.nl/StatWeb/publication/?PA=81567NED&LA=nl on 21 February 2015.
- Chang, T.-H., Lai, I., 1997. Analysis of characteristics of mixed traffic flow of autopilot vehicles and manual vehicles. *Transp. Res. Part C* 5 (6), 333–348.
- Fagnant, D.J., Kockelman, K., 2015. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transp. Res. A* 77, 167–181.
- Fernandes, P., Nunes, U., 2012. Platooning with IVC-enabled autonomous vehicles: strategies to mitigate communication delays, improve safety and traffic flow. *IEEE Trans. Intell. Transp. Syst.* 13 (1), 91–106.
- Gubins, S., Verhoef, E.T., 2011. Teleworking and Congestion: A Dynamic Bottleneck Analysis. Tinbergen Institute Discussion Paper, 2011-096/3.
- Hall, J.D., 2015. Pareto improvements from lexus lanes: the case for pricing a portion of the lanes on congested highways. Working Paper from University of Toronto, Department of Economics, nr 548. Accessed from <https://www.economics.utoronto.ca/public/workingPapers/tecipa-548.pdf> on 2 February 2016.
- International Transport Forum, 2015. Urban mobility system upgrade: how shared self-driving cars could change city traffic Corporate Partnership Board Report Accessed from www.internationaltransportforum.org/Pub/pdf/15CPB_Self-drivingcars.pdf on 18 January 2016.
- Kouwenhoven, M., de Jong, G.C., Koster, P., van den Berg, V.A.C., Verhoef, E.T., Bates, J., Warffemius, P.M.J., 2014. New values of time and reliability in passenger transport in the Netherlands. *Res. Transp. Econ.* 47, 37–49.
- Lamotte, R., de Palma, A., Geroliminis, N., 2016. Sharing the road: the economics of autonomous vehicles. HAL working paper, 01281425.
- Levin, M.W., Boyles, S.D., 2016. A multiclass cell transmission model for shared human and autonomous vehicle roads. *Transp. Res. Part C* 62, 103–116.
- Lindsey, R., 2004. Existence, uniqueness, and trip cost function properties of user equilibrium in the bottleneck model with multiple user classes. *Transp. Sci.* 38 (3), 293–314.
- Ni, D., Li, J., Andrews, S., Wang, H., 2010. Preliminary estimate of highway capacity benefit attainable with IntelliDrive technologies. 13th International IEEE Conference on Intelligent Transportation Systems (ITSC).
- Shladover, S., 2011. How vehicular networking can enable automated driving. Presentation for IEEE VNC - Amsterdam November 2011 Retrieved from <http://www.ieee-vnc.org/2011/talks/Shladover.pdf> on 30 October 2014.
- Shladover, S., Su, D., Lu, X.Y., 2012. Impacts of cooperative adaptive cruise control on freeway traffic flow. *Transp. Res. Rec.* 2324, 63–70.
- Small, K.A., 1982. The scheduling of consumer activities: work trips. *Am. Econ. Rev.* 72 (3), 467–479.
- Small, K.A., 2015. The bottleneck model: an assessment and interpretation. *Econ. Transp.* 4 (1–2), 110–117.
- Spence, A.M., 1975. Monopoly, quality, and regulation. *Bell J. Econ.* 6 (2), 417–429.
- Tientrakool, P., Ho, Y.-C., Maxemchuk, N.F., 2011. Highway capacity benefits from using vehicle-to-vehicle communication and sensors for collision avoidance. In: *Vehicular Technology Conference (VTC Fall)*. San Francisco, CA, USA. IEEE, pp. 1–5.
- US Census Bureau, 2014. U.S. Income and poverty in the United States: 2013. Government Printing Office, Washington, DC Current Population Reports, P60-249. U.S. Accessed from www.census.gov/content/dam/Census/library/publications/2014/demo/p60-249.pdf on 24 February 2015.
- US Department of Transportation, 2011. The value of travel time savings: departmental guidance for conducting economic evaluations revision 2. Accessed from www.dot.gov/sites/dot.dev/files/docs/vot_guidance_092811c.pdf on 21 February 2015.
- van Arem, B., van Driel, C.J., Visser, R., 2006. The impact of cooperative adaptive cruise control on traffic-flow characteristics. *IEEE Trans. Intell. Transp. Syst.* 7 (4), 429–436.
- van den Berg, V.A.C., 2014. Coarse tolling with heterogeneous preferences. *Transp. Res. B* 64, 1–23.
- van den Berg, V.A.C., Verhoef, E.T., 2011a. Congestion tolling in the bottleneck model with heterogeneous values of time. *Transp. Res. B* 45 (1), 60–70.
- van den Berg, V.A.C., Verhoef, E.T., 2011b. Winning or losing from dynamic bottleneck congestion pricing? The distributional effects of road pricing with heterogeneity in values of time and schedule delay. *J. Public Econ.* 95 (7–8), 983–992.
- VanderWerf, J., Shladover, S., Miller, M., Kourjanskaia, N., 2002. Effects of adaptive cruise control systems on highway traffic flow capacity. *Transp. Res. Rec.* 1800, 78–84.
- Vickrey, W.S., 1969. Congestion theory and transport investment. *Am. Econ. Rev.* 59 (2), 251–260 (Papers and Proceedings).
- Vickrey, W.S., 1973. Pricing, metering, and efficiently using urban transportation facilities. *Highway Res. Rec.* 476, 36–48.
- Zhang, F., Lindsey, R., Yang, H., 2016. The Downs-Thomson paradox with imperfect mode substitutes and alternative transit administration regimes. *Transp. Res. B* 86, 104–127.
- Zwaneveld, P.J., van Arem, B., 1997. Traffic effects of automated vehicle guidance system: a literature survey. TNO Inro, Delft, the Netherlands, INRO-VVG 1997-17.