

A Game-theoretic Approach to Leasing Agreements can Reduce Congestion

Reshef Meir
Hebrew University
Jerusalem, Israel
reshef24@cs.huji.ac.il

Jeffrey S. Rosenschein
Hebrew University
Jerusalem, Israel
jeff@cs.huji.ac.il

ABSTRACT

Leased cars with pre-paid fuel are a significant part of traffic today in many countries. Incentivized to drive as much as possible, their users contribute to pollution, congestion, and other negative societal effects. Calls for change of these leasing arrangements, by environmental organizations and others, are often rejected due to alleged economic rationales. We analyze from a game-theoretic perspective an alternative leasing model, where each driver pays for her own fuel. We show the emergence of a unique equilibrium in which everybody gains: the drivers, their employers who are paying for fuel, and, of course, the environment.

1. INTRODUCTION

There are over 28 million private cars today in the UK alone, driving over 400 *billion* km each year.¹ Approximately one fifth of this vast amount of traffic is attributed to daily commuters, many of them driving leased cars owned by the company that employs them (60% of the cars in the UK are company-owned). The rapidly increasing number of cars on the roads overloads existing infrastructure, and causes a set of environmental and economical problems. These include air, ground, and river pollution, an increase in accidents, the depletion of global oil reserves, and long-term atmospheric impact [8]. In the US, the time overhead due to rush-hour congestion is estimated at 1.2 minutes per kilometer, which adds up to billions of dollars of direct economic losses [2, 23]. In addition, congestion is responsible for increased fuel consumption, and thus further aggravates environmental effects.

Given the large scale of the problem, any change in local or global policies that could result in less traffic should be welcome. Yet many companies provide a strong incentive to their employees to drive *more*, in the form of a leased car with pre-paid fuel.²

We compare two leasing policies and how they affect em-

¹Statistics are taken from the UK Department of Transport [22].

²In Israel, where pre-paid fuel is the standard, leased cars are responsible for 5% of the *total annual mileage* (by combining data from [24] and the Central Bureau of Statistics [18]).

Cite as: A Game-theoretic Approach to Leasing Agreements can Reduce Congestion, Reshef Meir and Jeffrey S. Rosenschein, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX. Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

ployees (i.e., drivers), their companies, and the environment. Under the Common Policy (CP), the company pays for the employees' fuel, whereas in our suggested Alternative Policy (AP), fuel is not included in the leasing agreement.

While it seems intuitive that the Alternative Policy would be in the best interests of the environment, employees often prefer to get pre-paid fuel, which they see as consistent with their own interests. Our main claim against such an attitude on the part of employees is not that it is selfish, but that it is *wrong*. That is, self-interested employees should *prefer* a policy where the fuel is not pre-paid, as should their companies. What we intend to show is that the Alternative Policy induces a “win-win-win” situation: there are fewer cars on the road (thus the environment benefits); companies spend less money; and employees are happier. Moreover, this utopia not only exists—it is also an equilibrium state, meaning that the parties have no incentive to diverge from their behavior.

In the interests of the company and its employees we only consider monetary payments and car usage. That is, the environment is not a player in the game, and we *do not* take into account the effects of driving on the environment in the interests of employees/companies. We note however, that adding such considerations only strengthen our conclusion that the Alternative Policy should be preferred.

The intuition behind our argument is as follows. Under CP, since no additional cost is incurred by driving additional kilometers, employees will use their car at every opportunity, even when there is an alternative. Possible alternatives include cheap substitutions such as public transportation, car-pooling, or simply canceling unnecessary trips, but may also be more expensive, for example, using a taxi or a plane. While we do not explicitly study alternative means of transport, we make the plausible assumption that every car ride has some measurable money-equivalent *utility* to the driver.³ Some rides are more urgent, important, or harder to replace than others, and thus the utility of a 100km monthly quota is the cost of replacing *the most expensive* 100km with the cheapest available alternative. Naturally, the utility of driving 200km a month is higher than that of 100km, but not necessarily double, since the quota now includes less expensive rides. Similarly, we would probably be willing to pay even less in order to increase the quota by another 100km, and so on.

This assumption is known as *decreasing marginal utilities*, and is a standard assumption in economic situations. It

³We also note that cheap alternatives are in general more environment-friendly.

is also clear that at some point, the marginal utility from increasing the quota becomes zero. Otherwise, drivers with pre-paid fuel (that have an unlimited quota) would drive indefinitely.⁴

In exchange for a leased car and pre-paid fuel, under the CP policy a fixed amount is deducted from an employee’s salary. This amount depends on the contract between the company and its employees. In our proposed alternative policy, each employee pays for her fuel, and receives in turn a fixed salary increase (or a smaller salary deduction). We show that in each policy there is a unique equilibrium between the company and the employee, and that the equilibrium attained in the Alternative Policy is better for both sides. This occurs because the employee (i.e., the driver) stops using her vehicle for unnecessary trips, whose utility is lower than the cost of fuel.

While we make some general assumptions on the behavior of involved parties, we do not assume any specific values for the parameters of the leasing agreement (such as the price of fuel, or the distance to work). Our analysis is thus not restricted to a specific company or country.

1.1 Related Work

The multiagent approach has been applied to the traffic and transportation domain in two main ways. The first addresses various organizational problems by modeling the involved parties as self-interested agents, for example, to improve logistics within a freight fleet [13], to increase coordination between transportation companies [14], or to upgrade the service to clients of a public transportation system [12]. We take a similar approach in modeling the parties involved in a leasing arrangement. However, in our model we provide a formal analytic treatment of agent behavior, whereas the complexity of the aforementioned systems typically requires simulations (whose results are dependent on the specific configurations of the systems).

A different route in traffic research focuses directly on the problem of congestion. The agents in this case are typically the drivers that are making routing decisions (such as using the main road / side road) and timing decisions (such as leaving 15 minutes before their preferred time of arrival). Congestion is aggravated by the fact that if all drivers are making decisions that are *privately optimal*, then sometimes the *global* outcome is far from optimal. A famous example is the Braess paradox [9], which is a traffic-related instantiation of the problem known in economics and game theory as the *tragedy of the commons* [17]. Several attempts to alleviate this problem use some manipulation of the information that is available to the drivers [6, 4, 25], whereas others apply direct intervention of drivers’ incentives via external payments/tolls. Methods are evaluated mainly by simulations [21, 5, 10], but also using a formal analysis of the equilibrium, where possible (e.g., [1]). Another paper by Balbo and Pinson [3] combined the two research routes by modeling both the vehicles and the control components as agents, with the goal of regulating traffic. The proposed system (SATIR) have also been implemented and tested on data gathered in Brussels.

Crucially, all aforementioned congestion-related work makes the assumption that the total amount of traffic should be considered as *fixed*, and concentrates on preventing a “bad”

⁴And there is, in any case, a limited number of hours per month available for driving.

scenario, in which there are too many cars at the same place at the same time. This assumption was even made explicit by Tumer et al. [21], who stated that “no individual action is intrinsically bad, but that combinations of actions among agents lead to undesirable outcomes”.

The grave implications of excessive traffic, described in the previous section, make us question this assumption, as we believe that driving a car can be *intrinsically worse* than using an alternative. The goal of this paper is not to disperse traffic in time and space, but rather to reduce the total amount of traffic, thus alleviating congestion, but also all other negative consequences of traffic.

Complex models that take into account changes in traffic volume (due to tolls and due to the congestion itself) have also been proposed [19]. These models still assume that drivers are very sophisticated and that they always find the exact equilibrium of the network (see [15, 7] for critique on this assumption). In contrast, our model suggests a simple policy transition, leaving the players with an obvious optimal strategy that does not depend on the topology of the network.

1.2 Paper Structure

We first clarify some game-theoretic concepts, and formalize our intuition from the first section by considering a simplified game where only the employee acts strategically. We then add the company as a strategic player, and analyze the equilibrium in the induced game. In the remaining sections we show how our results extend to more realistic situations, where there are multiple employees with different preferences, and where employees also get the option of dropping the leasing contract. In the last section, we discuss some implications of our results, and compare them with the situation in practice.

2. PRELIMINARIES

Game.

A *game* consists of a set of agents N , a set of strategies for each agent $\{A_i\}_{i \in N}$, and a utility function for each agent $U_i : \times_{j \in N} A_j \rightarrow \mathbb{R}$. The set of strategies A_i does not have to be discrete. For example, the strategy may be to decide on an amount of money to spend. A joint selection of strategies for each agent $\mathbf{a} = \{a_j \in A_j\}_{j \in N}$ is called a *strategy profile*. The profile of all agents *except* i is denoted by $a_{-i} = \{a_j \in A_j\}_{j \neq i}$.

Equilibrium.

We say that the strategy profile \mathbf{a} is an *equilibrium*, if no agent can gain by choosing a different strategy, assuming that all other agents keep theirs. Formally, \mathbf{a} is an equilibrium if for any agent i and any strategy $a'_i \neq a_i$, we have that $U_i(\mathbf{a}) \geq U_i(a_{-i}, a'_i)$. Our definition coincides with the standard definition of a *pure Nash equilibrium*. Since we do not allow agents to randomize between strategies, we only consider pure equilibria. Thus, it is possible that a game does not contain any equilibrium.

Dominant strategies.

$a_i^* \in A_i$ is a *dominant strategy* of i if agent i always prefers a_i^* , regardless of the choices of other agents. Formally, for all \mathbf{a} , $U_i(a_i^*, a_{-i}) \geq U_i(\mathbf{a})$. Note that if some player has a

unique dominant strategy, then all other players can assume that this strategy will be played. This simplifies the game, as the size of the strategy space is significantly reduced. In particular, it is possible that in the new, simplified game, there is an agent $j \neq i$ that has a dominant strategy (under the assumption that i plays a_i^*). It is sometimes possible to continue to remove strategies from the game until there is only one strategy profile left. In this case we say that the game is *iterated dominance solvable*. The outcome \mathbf{a}^* is called the *iterated dominant strategy equilibrium*, and it is also the unique Nash equilibrium of the game.

For a detailed discussion regarding these definitions and for more background in game theory see e.g. [20].

3. INITIAL MODEL

3.1 The Common Policy

In the simplest case, we model the interaction between a single company c and a single employee e . The utility of the employee (denoted by $u = U_e$) is composed of two factors: one factor is her income, which we denote s . The other factor is the number of kilometers she drives in a month (mileage), denoted by x . While the income s is not controlled by the employee, she is free to choose how much to drive; thus her strategy space is $A_e = \mathbb{R}_+$ (and $x \in A_e$).

In the common leasing policy (which we denote by CP), the utility of the employee can be decomposed as

$$u_{CP}(s, x) = s + f(x) \text{ ,}$$

i.e., there is some function f that makes the two factors comparable. As explained in the introduction, we make the following assumptions regarding the utility of the employee:

ASSUMPTION 1. *The employee has a decreasing marginal utility from driving more, and there is a maximal mileage that the employee has no reason to exceed. Formally:*

- a. f is non-decreasing and continuous.
- b. f is concave, i.e., for all $y < z$ and $\epsilon > 0$, $f(y + \epsilon) - f(y) \geq f(z + \epsilon) - f(z)$.
- c. There is some x^* s.t. f has a maximum in $f(x^*)$.

For simplicity, we will make the technical assumption that f is *strictly concave* in the range $[0, x^*]$, although this assumption is not necessary and can be relaxed. Thus, for all $0 \leq y < z \leq x^*$ we have that $f(y + \epsilon) - f(y) > f(z + \epsilon) - f(z)$.

Clearly, in CP the dominant strategy of the agent is to drive x^* , thus maximizing utility. This holds for any fixed income s , and therefore the company has no influence on the strategy of the employees regarding their mileage.

We also compute the utility of the company (denoted by $v = U_c$), although for now we will not treat the company as a player in the game (i.e., we will not consider the rationality of its actions). The fuel cost is linear in the mileage and we denote by k the average cost of fuel per 1 kilometer of driving; thus $v_{CP}(s, x) = -s - x \cdot k$.

If we assume that the employee follows her dominant strategy, we get that in CP, the utility of the employee is $u_{CP}(s, x^*) = s + f(x^*)$, while for the company $v_{CP}(s, x^*) = -s - x^* \cdot k$.

3.2 The Alternative Policy

Now suppose that in addition to the fixed income, our employee also has to pay for consumed fuel. We define a




Figure 1: Utility of the employee as a function of the mileage in the common leasing policy (top) and in the alternative policy (bottom). x^* and x' are the employee's dominant strategies in both policies, respectively.

new game for the alternative leasing policy (AP), with the same strategies but different utility functions. Fuel cost is linear in the mileage, thus the utility of the employee in AP is:

$$u_{AP}(s, x) = s + f(x) - k \cdot x \text{ .}$$

The relation between the common and alternative policies w.r.t. the utility of the employee is demonstrated in Figure 1. The best strategy for the employee in AP depends on both k and f , and we make the following observations:

- $f(x) - k \cdot x$ is still concave.
- f has a peak in some $x' < x^*$ (see Figure 1).
- Regardless of s , the dominant strategy of the employee in AP is to drive x' .
- The utility for the dominant strategy is $u'_{AP}(s) = s + f(x') - k \cdot x'$.

It is also clear, that if the income of the employee remains the same, then paying for fuel will only decrease her utility, i.e., $u_{AP}(s, x) < u_{CP}(s, x)$. However, we should keep in mind that the company saves money by not paying for the employee's fuel.

The utility of the company becomes even simpler in the Alternative Policy, as the only factor that has an effect is the salary itself, i.e., $v_{AP}(s, x) = -s$. If we assume that the employee follows her dominant strategy, we have that in AP, the company may increase the salary by Δ and still gain (compared to CP), as long as $v_{AP}(s + \Delta, x') > v_{CP}(s, x^*)$. This supplies us with a simple formalization of the intuition given earlier.

PROPOSITION 2. *There is a strategy for the company such that both company and employee gain by switching to the Alternative Policy. Formally, there exists $\Delta > 0$ such that*

1. $u_{AP}(s + \Delta, x') > u_{CP}(s, x^*)$, and
2. $v_{AP}(s + \Delta, x') > v_{CP}(s, x^*)$.

PROOF. The constraint on v is satisfied as long as $\Delta < k \cdot x^*$, since $v_{AP}(s + \Delta, x') - v_{CP}(s, x^*) = s + k \cdot x^* - s - \Delta$.

Also, as long as $\Delta > f(x^*) - f(x') + kx'$, the constraint on u is satisfied, as

$$\begin{aligned} u_{AP}(s + \Delta, x') - u_{CP}(s, x^*) & \\ &= s + \Delta + f(x') - k \cdot x' - (s + f(x^*)) \\ &= \Delta - (f(x^*) - f(x') + kx') > 0 . \end{aligned}$$

It is thus left to show that both constraints can be satisfied at the same time. Recall that x' was defined such that $f(x') - kx' = \max_{x \geq 0}(f(x) - kx)$. In particular,

$$\begin{aligned} f(x') - kx' &> f(x^*) - kx^* && \Rightarrow \\ kx^* &> f(x^*) - f(x') + kx' && \Rightarrow \\ \exists \Delta, & \quad kx^* > \Delta > f(x^*) - f(x') + kx' . && (1) \end{aligned}$$

□

3.3 Weaknesses of the Initial model

Our basic model provides some formal flavor for the intuition that win-win-win situations can be achieved simply by transferring the fuel cost from the company to the employee, with very few additional assumptions. Unfortunately, such a simple analysis still suffers from several weaknesses.

First, although the employee had a dominant strategy in both policies, we did not analyze the actions available to the company from a strategic point of view. Therefore it is possible in principle that the new state is not an equilibrium.

Second, we only modeled a single employee. Although the results hold if we add *identical* employees, we have a problem when employees have different utility functions. For example, if one employee lives closer to the train station, or happens to like bicycling, then the utility of 100km for her might be lower than the utility for her colleague.

Third, we did not consider other actions available to the employee, such as dropping the leasing deal altogether.

In the following sections, we address these issues by refining our model and adding more assumptions where needed.

4. LEASING AS A TWO-PLAYER GAME

In this section we will extend our initial model by formally defining the different factors affecting the company's utility from the leasing interaction, and analyze the stability of the possible outcomes w.r.t. both the company and the employee.

We first note that our original definition of the company's utility, which only considered expenses, ignored an important factor. The company *gains* something from the employee, otherwise she would not be employed in the first place. As with f , the "gain" function may take different forms, but we can make similar plausible assumptions about it. Formally, we denote the gain function by $g(u)$, where u is the utility of the employee.

ASSUMPTION 3. *The company has a decreasing marginal profit from the utility of the employee. Formally:*

- a. g is non-decreasing and continuous.
- b. g is strictly concave.⁵

⁵As with f , we make this assumption to simplify the analysis, and in practice a weaker restriction on g would suffice.

There are several justifications for the monotonicity assumption. This can be interpreted as "happy employees work harder", but also as "better conditions attract better employees". We do not search for the "correct" interpretation, as the implication is the same. The decreasing marginal profit is a standard economic assumption.

We can now rewrite the utility of the company, considering the productivity of the employee for both policies:

$$\hat{v}_{CP}(s, x) = g(u_{CP}(s, x)) + v_{CP}(s, x) = g(s + f(x)) - s - kx ,$$

and

$$\hat{v}_{AP}(s, x) = g(u_{AP}(s, x)) + v_{AP}(s, x) = g(s + f(x) - kx) - s . \quad (2)$$

Recall that in either policy, the employee has a dominant strategy (either x^* or x'). Assuming that the employee is indeed using her dominant strategy, we expect the company to optimize the salary, so as to maximize its profit. Thus there is an optimal salary s^* that maximizes $v_{CP}(s, x^*)$, and the strategy profile (x^*, s^*) is the iterated dominant strategy equilibrium of CP.

As for the Alternative Policy, recall that in order to make both sides benefit, the company needs to increase the salary of the employee by Δ , under the constraints described earlier (Equation (1)). Let $\hat{\Delta}$ be an arbitrary amount that satisfies the constraints. Note that

$$u_{AP}(s^* + \hat{\Delta}, x') > u_{CP}(s^*, x^*) \quad \Rightarrow \quad \text{(from prop. 2)}$$

$$g(u_{AP}(s^* + \hat{\Delta}, x')) > g(u_{CP}(s^*, x^*)) \quad \Rightarrow \quad \text{(from monotonicity of } g)$$

$$\begin{aligned} \hat{v}_{AP}(s^* + \hat{\Delta}, x') &= g(u_{AP}(s^* + \hat{\Delta}, x')) + v_{AP}(s^* + \Delta, x') \\ &> g(u_{CP}(s^*, x^*)) + v_{AP}(s^* + \Delta, x') \\ &> g(u_{CP}(s^*, x^*)) + v_{CP}(s^*, x^*) \quad \text{(from prop. 2)} \\ &= \hat{v}_{CP}(s^*, x^*) \quad \Rightarrow \end{aligned}$$

$$\hat{v}_{AP}(s^* + \hat{\Delta}, x') > \hat{v}_{CP}(s^*, x^*) , \quad (3)$$

so the company still gains according to \hat{v} . However, this alone does not guarantee stability. Hypothetically, it is possible that in AP there is no equilibrium, or that there is an equilibrium that is worse for either the company or the employee. We will now see that this is not the case.

We denote by $s' = s^* + \Delta'$ the best response of the company to the dominant strategy of the employee (i.e., to x') in AP, thus (s', x') is the (unique) iterated dominant strategy equilibrium of AP.

PROPOSITION 4. *The equilibrium profile (s', x') in the Alternative Policy is preferred by both players to the equilibrium (s^*, x^*) in the Common Policy. Formally,*

1. $u_{AP}(s', x') \geq u_{CP}(s^*, x^*)$, and
2. $\hat{v}_{AP}(s', x') > \hat{v}_{CP}(s^*, x^*)$.

PROOF. Denote by Δ_{inf} the infimum of Δ that obeys the constraints imposed by (1), i.e., $\Delta_{inf} = f(x^*) - f(x') + kx'$. The following lemma shows that in the equilibrium profile, the company must increase the salary by at least Δ_{inf} .

LEMMA 5. *If $\Delta < \Delta_{inf}$ then*

$$\hat{v}_{AP}(s^* + \Delta, x') < \hat{v}_{AP}(s^* + \Delta_{inf}, x') .$$




Figure 2: The gain function (top) and the utility of the company in the Common Policy (bottom), as a function of the salary.

PROOF. Assume that g is continuously differentiable, thus v_{CP} is also continuously differentiable in s . Also, $\hat{v}_{CP}(s, x^*) = g(s + f(x^*)) - s - kx^*$ (as a function of s) has a maximum in s^* , i.e., its derivative in $s = s^*$ is 0, and is strictly positive in any $s < s^*$ (from concavity). By differentiating g we have that

$$\frac{\partial \hat{v}_{CP}(s, x^*)}{\partial s} = \frac{\partial g(s + f(x^*)) - s - kx^*}{\partial s} = \frac{\partial g(s + f(x^*))}{\partial s} - 1,$$

thus the slope of g in $u = s + f(x^*)$ is *higher than 1* in any $s < s^*$. Consider the interval $[u_1, u_2]$, where $u_1 < u_2 \leq s^* + f(x^*)$. The slope of the straight line connecting the values of g in both edges of the interval (i.e., $\frac{g(u_2) - g(u_1)}{u_2 - u_1}$) cannot be lower than 1, as 1 is a lower bound of the derivative of g in the interval (see figure 2), thus for any $a > 0$ we get that

$$g(s^* + f(x^*)) - g(s^* + f(x^*) - a) > s^* + f(x^*) - (s^* + f(x^*) - a) = a. \quad (4)$$

We now add that Equation (4) holds even if g is *not* continuously differentiable, since its slope is still lower bounded by 1 (although it requires some technical work to show that).

We now take $a > 0$ to be the difference $\Delta_{inf} - \Delta$.

$$\begin{aligned} & \hat{v}_{AP}(s^* + \Delta, x') - \hat{v}_{AP}(s^* + \Delta_{inf}, x') \\ &= [g(s^* + \Delta_{inf} - a + f(x')) - kx' - (s^* + \Delta_{inf} - a)] \\ & \quad - [g(s^* + \Delta_{inf} + f(x')) - kx' - (s^* + \Delta_{inf})] \quad (\text{from (2)}) \\ &= g(s^* + f(x^*) - f(x') + kx' - a + f(x') - kx') + a \\ & \quad - [g(s^* + f(x^*) - f(x') + kx' + f(x') - kx')] \\ &= a + g(s^* + f(x^*) - a) - g(s^* + f(x^*)) \\ &< a - a = 0 \quad (\text{from (4)}) \end{aligned}$$

□

As an immediate corollary we get that for the best response s' , Δ' must be at least $\Delta_{inf} = f(x^*) - f(x') + kx'$, which means (as seen in the proof of Proposition 2) that

$$u_{AP}(s', x') = u_{AP}(s^* + \Delta', x') \geq u_{CP}(s^*, x^*).$$

That is, the employee is indeed not harmed in the new equilibrium.

As for the company, we have that

$$\hat{v}_{AP}(s', x') \geq \hat{v}_{AP}(s^* + \Delta', x') > \hat{v}_{CP}(s^*, x^*),$$

where the first inequality is due to the fact that s' is the best response to x' , and the second is due to (3). Thus the company is even better off with the new equilibrium.

5. MULTIPLE EMPLOYEES

We now turn to an extension to multiple employees. As noted, adding *identical* employees makes no difference, as the equilibrium described in previous sections will satisfy all of them independently. Unfortunately (at least from an analytic point of view) different people do have different preferences, which are reflected in our model as different functions f_i for each employee. Proposition 2 cannot be extended to this case, as the following example shows:

EXAMPLE 6. Suppose $k = 1$. For the first employee, $x_1^* = 10$; $f_1(x_1^*) = 10$; $x_1' = 5$; $f_1(x_1') = 8$. For the second employee $x_2^* = 20$; $f_2(x_2^*) = 20$; $x_2' = 12$; $f_2(x_2') = 15$. We can easily see that

$$f_2(x_2^*) - f_2(x_2') + kx_2' = 17 > 10 = kx_1^*,$$

so any fixed salary increase Δ will either disappoint employee 2, or will make the leasing deal of employee 1 less desirable for the company (or both).

In the general formulation of the multi-employee problem, there are n employees. Together with the company, the game now has $n + 1$ players. The strategy space of each employee is her mileage, as in Section 4. As for the company, it is possible in theory to give a different salary increase Δ_i to every employee. This would break down the game to n independent games that can be solved as in Section 4. However, this would be unfair and impractical, since this increase would be based on personal habits—some employees will get more just because they like to drive more. Moreover, employees might behave strategically by increasing their mileage before the new policy takes effect, thus manipulating the value of x_i^* and Δ_i .

The salary increase may also be based on the distance from the employee's home to her workplace. However, this idea only supplies us with a partial solution, since the pre-paid fuel is also used for private needs, and also because some alternative commuting solutions may be unavailable or inconvenient for some of the employees.

We therefore assume that the strategy of the company has a single parameter, Δ , which is the salary increase given to all employees in the Alternative Policy. Thus the utility function of the employees remains the same:

$$u_{i AP}(\Delta, x_i) = f_i(x_i) + s_i^* + \Delta - kx_i,$$

where s_i^* is the salary of the employee in the Common Policy (assumed to be fixed). The utility of the company from each interaction is $\hat{v}_{i AP}(\Delta, x_i) = g(u_{i AP}(\Delta, x_i)) - s_i^* - \Delta$, and

its total utility in the game is

$$\mathbf{v}_{AP}(\Delta, x_1, \dots, x_n) = \sum_{i=1}^n \hat{v}_i{}_{AP}(\Delta, x_i) .$$

The dominant strategy of each employee does not depend on Δ , nor on the mileage of the other employees. Thus, we can continue to assume that employee i drives x'_i kilometers in the Alternative Policy.

We take the simple approach of computing the *average* value $\bar{\Delta} = \frac{1}{n} \sum \Delta_i$, where Δ_i is determined according to the two-player game between the company and employee i , as in Section 4. We compute the social welfare in the new policy (the company's utility is computed separately and is *not considered* part of the social welfare). We find that

$$\begin{aligned} \sum_{i=1}^n u_i{}_{AP}(s_i^* + \bar{\Delta}, x'_i) &= \sum_i (s_i^* + \bar{\Delta} + f_i(x'_i) - kx'_i) \\ &= n\bar{\Delta} + \sum_i (s_i^* + f_i(x'_i) - kx'_i) \\ &= \sum_i \Delta_i + \sum_i (s_i^* + f_i(x'_i) - kx'_i) \\ &= \sum_i (s_i^* + \Delta_i + f_i(x'_i) - kx'_i) \\ &= \sum_i u_i{}_{AP}(s_i^* + \Delta_i, x'_i) \geq \sum_{i=1}^n u_i{}_{CP}(s_i^*, x_i^*) , \end{aligned} \quad (5)$$

i.e., the social welfare still improves in the Alternative Policy (although some employees may still be unhappy).

We now compute the utility \mathbf{v} of the company, when summing over all interactions. Since $\Delta_i < kx'_i$, we get that

$$n\bar{\Delta} < \sum_{i=1}^n kx'_i , \quad (6)$$

thus the company still saves money w.r.t. the Common Policy. This still does not mean that the utility of the company improves, since we did not consider the gain (g) yet.

Unfortunately, even though the expenses of the company are lower *and* the social welfare increased (suggesting employees are happier), it is not guaranteed that the overall utility of the company increases. This is due to the non-linearity of the gain function g . It suggests that there may potentially be an embittered employee whose productivity now deteriorates significantly, dragging down the average gain. A closer look at this scenario reveals that not every employee can have such a negative effect. The happier employees are (before the change), the smaller their effect on the change in the average gain (due to the concavity of g). If indeed the Alternative Policy is more profitable to those who are initially worse off, then the increase in social welfare will induce an increase in the average gain—and hence in the utility of the company. Moreover, it is quite reasonable to assume that in reality, the employees who benefit the most from the common leasing policy are indeed those who exploit it the most by using a very high mileage. Thus these employees will indeed benefit less than others from the new policy, as in the case in Example 6. We now formalize and prove this intuition.

ASSUMPTION 7. *The happier an employee is in the Common Policy CP, the smaller her benefit from the Alternative*

Policy AP (it may be negative). Formally, if

$$u_i{}_{CP}(s_i^*, x_i^*) > u_j{}_{CP}(s_j^*, x_j^*) ,$$

then

$$\begin{aligned} u_i{}_{AP}(s_i^* + \bar{\Delta}, x_i^*) - u_i{}_{CP}(s_i^*, x_i^*) \\ < u_j{}_{AP}(s_j^* + \bar{\Delta}, x_j^*) - u_j{}_{CP}(s_j^*, x_j^*) . \end{aligned}$$

ASSUMPTION 8. *The gain of the company g is uniform and does not depend on the identity of the employee.*

The justification of Assumption 8 is as follows. Unlike f_i , which reflects the private enjoyment of each employee from using her car, the gain function g depends more on the specific job requirements. Although it is unlikely that one gain function will fit *all* employees, it is still reasonable to make this assumption for a group of employees in a similar position. Thus the leasing agreement can be retuned for each such group separately.

PROPOSITION 9. *Both social welfare and the utility of the company increase in the Alternative Policy. Formally,*

$$\sum_i u_i{}_{AP}(\bar{\Delta}, x'_i) \geq \sum_i u_i{}_{CP}(s_i^*, x_i^*) , \quad (7)$$

$$\mathbf{v}_{AP}(\bar{\Delta}, x'_i) > \sum_i \hat{v}_i{}_{CP}(s_i^*, x_i^*) . \quad (8)$$

PROOF. We get (7) directly from Equation (5), so we only need to prove the company's side.

LEMMA 10.

$$\sum_i g(u_i{}_{AP}(s_i^* + \bar{\Delta}, x'_i)) > \sum_i g(u_i{}_{CP}(s_i^*, x_i^*)) .$$

PROOF. We denote $a_i = u_i{}_{CP}(s_i^*, x_i^*)$ and $b_i = u_i{}_{AP}(s_i^* + \bar{\Delta}, x'_i)$. Assume w.l.o.g. that employees are sorted according to a_i (increasing). From Assumption 7, this also means that the difference $b_i - a_i$ is *decreasing*. That is, the employee with the lowest index has the largest benefit from AP, then the benefit gets smaller and smaller (and possibly negative) for larger i .

We now define a new set of points, b'_i , in the following way. We take every j s.t. $b_j < a_j$, and “push” it up toward a_j , until $b'_j = a_j$. We then compensate by pushing b_1 (down) toward a_1 . If $b'_1 = a_1$ already, we continue to push the next point b_2 and so on, until $\sum b'_i = \sum b_i$. This step is demonstrated in Figure 3. We repeat the process as long as there are points such that $b'_j < a_j$. Note that the process must end, since $\sum b'_i = \sum b_i > \sum a_i$ (from Equation 5).

From the concavity of g , when two points $j < i$ are pushed in opposite directions by ϵ , we get that

$$g(b_j - \epsilon) + g(b_i + \epsilon) \leq g(b_j) + g(b_i) ,$$

since g is steeper around j . Thus, after all steps are performed $\sum_i g(b'_i) \leq \sum_i g(b_i)$. Also, after the final step there are no points such that $b'_i < a_i$, thus

$$\sum_i g(a_i) \leq \sum_i g(b'_i) < \sum_i g(b_i) ,$$

as required. \square




Figure 3: The relation between the different sets of points before and after translation. The points b_1 and b_2 were moved left to compensate for moving b_j right. It is easy to see that $|b'_j - b_j| = |b'_1 - b_1| + |b'_2 - b_2|$, but $g(b'_j) - g(b_j) < (g(b_1) - g(b'_1)) + (g(b_2) - g(b'_2))$ which means $g(b_1) + g(b_2) + g(b_j) > g(b'_1) + g(b'_2) + g(b'_j)$.

We continue with the proof:

$$\begin{aligned}
\mathbf{v}_{AP}(\bar{\Delta}, x_1, \dots, x_n) &= \sum_i \hat{v}_i \text{AP}(s_i^* + \bar{\Delta}, x'_i) \\
&= \sum_i (g(u_i \text{AP}(s_i^* + \bar{\Delta}, x'_i)) - (s_i^* + \bar{\Delta})) \\
&> \sum_i g(u_i \text{CP}(s_i^*, x'_i)) - \sum_i (s_i^* + \bar{\Delta}) \quad (\text{from lemma 10}) \\
&= \sum_i g(u_i \text{CP}(s_i^*, x'_i)) - \sum_i s_i^* - n\bar{\Delta} \\
&> \sum_i g(u_i \text{CP}(s_i^*, x'_i)) - \sum_i s_i^* - \sum_i kx_i^* \quad (\text{from (6)}) \\
&= \sum_i g(u_i \text{CP}(s_i^*, x'_i)) - s_i^* - kx_i^* = \sum_i \hat{v}_i \text{CP}(s_i^*, x_i^*) .
\end{aligned}$$

□

6. DROPPING THE CONTRACT

So far we assumed that the strategy of the employee is limited to the mileage she drives in her car. However, an employee who does not believe that her leasing deal is profitable, will simply drop the contract. By doing so, our employee will typically turn to use her private car.⁶

To incorporate this type of behavior in our model, we will formulate the utility of the agent when not engaged in a leasing deal at all:

$$u_0(x) = s_0 + f(x) - k_0 \cdot x .$$

The base salary is of course higher, since no money is deducted from it, thus $s_0 > s^*$. However, the employee now needs to pay even more per kilometer, as there are other expenses on top of fuel, thus $k_0 > k$.

It is easy to see that the maximum of u_0 is reached for some $x'' < x'$, due to the increased cost per kilometer. Now,

⁶It is quite unlikely that a leasing deal was profitable in the first place for a person who can manage without a car at all.

if we keep our analysis restricted to an interaction with a single employee, there is no problem. We know that if the employee used a leased car in the first place (i.e., in the Common Policy), then $u_{CP}(s^*, x^*) > u_0(s_0, x'')$. As we already showed in Section 4, in the Alternative Policy the employee only ends up happier, and there is no reason for her to drop the contract after the policy has been switched.

Of course, in a typical company with many employees, some of them might become disappointed with the new deal (as we saw in Section 5) and renounce it altogether. At least from the environmental point-of-view, this is not at all bad, as the mileage will decrease even more to x'' . Moreover, if an employee decides to drop the leasing deal, it is because this decision is better for him, which means social welfare increases even more. Also, leasing deals are typically subsidized by companies; the company is almost never damaged when an employee returns her car. To formalize these statements, we will add some notations. c is the fixed cost of each leasing deal to the company (not including fuel). $v_0(x)$ represents the basic utility of the company when the employee does not use a leased car:

$$v_0(x) = g(u_0(x)) - s_0 + c .$$

The $+c$ represents the fixed leasing cost which is saved when the employee does not lease a car. As leasing deals are subsidized (on average), we assume that

$$c \geq s_0 - (s^* + \bar{\Delta}) . \quad (9)$$

We now incorporate the new strategy (of the employee) $z \in \{TAKE, DROP\}$ into the utility functions:

$$\bar{u}_{AP}(s, x, z) = \begin{cases} u_0(x) & , \text{if } z = DROP \\ u_{AP}(s, x) & , \text{if } z = TAKE \end{cases} , \quad (10)$$

and

$$\bar{v}_{AP}(s, x, z) = \begin{cases} v_0(x) & , \text{if } z = DROP \\ v_{AP}(s, x) & , \text{if } z = TAKE \end{cases} . \quad (11)$$

Clearly, there are only two possible outcomes: either the employee takes the deal, in which case she plays x' and the company plays $s^* + \bar{\Delta}$ (as described in previous sections). Otherwise, the employee drops the deal and plays x'' , while the company pays the base salary s_0 . We assume that the employee chooses the better possibility for her from these two options. We denote by $(BEST_i)$ the strategy vector preferred by the employee i (i.e., $(BEST_i)$ is either $(s_i^* + \bar{\Delta}, x'_i, TAKE)$ or $(s_{i0}, x''_i, DROP)$, whichever maximizes $\bar{u}_{i AP}$).

PROPOSITION 11. *Even if employees can choose to drop the leasing deal, the Alternative Policy still increases both social welfare and the company's utility (under all the assumptions described so far). Formally:*

$$\sum_i \bar{u}_{i AP}(BEST_i) \geq \sum_i u_{i CP}(s_i^*, x_i^*) ,$$

and

$$\sum_i \bar{v}_{i AP}(BEST_i) > \sum_i \hat{v}_{i CP}(s_i^*, x_i^*) .$$

PROOF. We begin with the social welfare. Since for every employee

$$\bar{u}_{i AP}(BEST_i) \geq \bar{u}_{i AP}(s_i^* + \bar{\Delta}, x'_i, TAKE) = u_{i AP}(s_i^* + \bar{\Delta}, x'_i) ,$$

then from Equation (7) of Proposition 9,

$$\sum_i \bar{u}_{i AP}(BEST_i) \geq \sum_i u_{i CP}(s_i^*, x_i^*) .$$

We also get from monotonicity of g that

$$g(\bar{u}_{i AP}(BEST_i)) \geq g(u_{i AP}(s_i^* + \bar{\Delta}, x'_i)) . \quad (12)$$

This means that the company is never damaged even if the employee returns the car, since in this case $(BEST_i) = (s_{i0}, x''_i, DROP)$ and

$$\begin{aligned} \bar{v}_{i AP}(BEST_i) &= v_{i0}(x''_i) = g(u_{i0}(x''_i)) - s_{i0} + c \\ &= g(\bar{u}_{i AP}(BEST_i)) - s_{i0} + c \\ &\geq g(u_{i AP}(s_i^* + \bar{\Delta}, x'_i)) - s_{i0} + c \quad (\text{from (12)}) \\ &\geq g(u_{i AP}(s_i^* + \bar{\Delta}, x'_i)) - (s^* + \bar{\Delta}) \quad (\text{from (9)}) \\ &= \hat{v}_{i AP}(s_i^* + \bar{\Delta}, x'_i) . \end{aligned}$$

If the employee keeps the car, then $\bar{v}_{i AP}(BEST_i) = \hat{v}_{i AP}(s_i^* + \bar{\Delta}, x'_i)$ by Definition (11).

Finally, when we sum up the gains, then from Equation (8) of Proposition 9,

$$\sum_i \bar{v}_{i AP}(BEST_i) \geq \sum_i \hat{v}_{i AP}(s_i^* + \bar{\Delta}, x'_i) > \sum_i v_{i CP}(s_i^*, x_i^*) ,$$

as required. \square

7. DISCUSSION

We showed that transferring the fuel cost from the company to the employee has more benefits than “just” helping the environment and reducing congestion. It will actually leave both employees and their employer richer and happier. This result holds under quite weak and realistic assumptions on the preferences of the involved parties. Moreover, results still hold when we consider stability issues, multiple interactions, and the option to return the car. The underlying idea that is responsible for this situation is *marginal benefit vs. marginal cost*. When an employee does not pay for

fuel, her marginal cost of driving more is 0, which gives her an incentive to use her car even when the marginal benefit from it is negligible. On the other hand, using the car does not really come for free—it does have a cost, which is externalized and incurred on the company (and on the environment). The company, in turn, rolls some of this cost back on the employee, “hidden” inside the salary deduction of the leasing deal.

7.1 Equilibrium and Commitments

In Section 4 we showed that the Alternative Policy not only enables a situation that is better for everyone, but that this situation is also an equilibrium (in iterated dominant strategies). In Sections 5 and 6 we also suggested a strategy for the company that makes the Alternative Policy better for all the involved parties. However, this strategy (i.e., $\bar{\Delta} = \frac{1}{n} \sum_{i=1}^n \Delta_i$) is not necessarily the optimal strategy of the company, and so the new state does not have to be an equilibrium. This means that employees may have a justified objection to the transition, if they do not trust the company to carry out the suggested strategy.

In order to solve this issue, we will use the notion of *commitments*. A player that commits to a strategy (i.e. *limits* his own freedom) can in certain situations create an equilibrium where it does not exist, or shift an existing equilibrium in a game towards one that is more preferred to him [11]. Our intent is slightly different, as the commitment is supposed to convince the other players to play a different game (i.e. AP instead of CP). A company that is interested in performing such transition (to the Alternative Policy), could alleviate the suspicions of its employees by *committing* to the aforementioned strategy. That is, the company will publicly announce the intended raise Δ due to the policy shift, and a binding contract can be used to enforce such a commitment. This makes the outcome we analyzed in the last two sections an equilibrium in dominant strategies, as only the employees are free to change their strategy (and they have no incentive to do so).

7.2 The Real World

The immediate question that arises from our results is about their validity in the real world. If the Alternative Policy is indeed so desirable, we would expect companies in the market to have adopted it by now. We supply two possible explanations for the current situation, although these should be taken only as preliminary suggestions, as this question is not the focus of this paper (we do intend to study this question more deeply in future research).

The first reason is that local or national taxation policy makes the Common Policy (with pre-paid fuel) more profitable for companies, thus effectively subsidizing fuel that is paid for by the company. The taxation system in the UK until recently was shown to act in this manner[23].⁷

A second reason, that is perhaps harder to verify, is that the benefit of free fuel *is perceived* (in the eyes of the employee) as better than it really is. While the implicit assumption that players in a game behave rationally is usually valid for *companies* (which seek to maximize their profit), employees as individuals may be much more affected by irrational thought patterns. Default-bias and loss-aversion could pos-

⁷Taxation policy was also pointed out by Windsor and Omer [24] as one of the contributors to the ubiquity of the leasing arrangement in Israel.

sibly account for the reluctance of employees to adopt the Alternative Policy, whereas companies refrain from a policy change that is perceived as hurting its employees.

7.3 Sustainable Transportation

The validity of our argument crucially depends on the existence and availability of alternative transportation methods, and alternative commuting solutions in particular. In the absence of these, employees will not be able to reduce their mileage,⁸ and will not be able to benefit from the Alternative Policy. Direct and indirect subsidies for private cars were pointed out as having a negative effect on the development and embedding of alternative means of transportation [16]. Similarly, it is likely that the Common Policy creates a “vicious circle” in a way, since the large number of employees that have pre-paid fuel lowers the demand (and in turn the availability) of alternative solutions such as public transportation and car-pooling. The low availability of alternatives is used to justify the benefit of pre-paid fuel, and so on. Steps that are taken at the national level to support the alternatives, such as subsidies for trains or taxes on fuel, are less effective since they are not relevant for a significant part of the population.

Making *all* drivers face the true costs of their behavior should also assist in breaking this cycle and promote the availability of alternatives that would benefit the rest of the population.

Our paper demonstrates how economic theory supports an environment-friendly policy by eliminating externalities that affect both the players and the environment. In the words of P. B. Goodwin, this is an example of how a “gold-green” coalition can emerge [15].

7.4 Conclusion and Future Work

This paper deals with the optimal behavior of *rational* players under “pure” conditions, that is, with no external intervention. We proved that under these conditions, pre-paid fuel should not be considered a benefit. We hope that arguments such as the one presented in this paper can assist in removing obstacles such as the ones described in Section 7.2, by highlighting the negative role of a given taxation policy and convincing policymakers of the benefits of the transition.

We emphasize the fact that we do not suggest adding a new mechanism that will help reduce congestion, but rather to *remove* any external intervention in the form of a fuel subsidy, and let the market do its work.

We believe that some of the assumptions used in our model are perhaps too strong, and we intend to obtain stronger theoretical results by relaxing these assumptions where possible. It may also be possible to take into account irrational (yet predictable) aspects in the behavior of the involved parties in order to better understand the implications of each policy (similar to the approach taken by Bazzan et al. w.r.t. recommendation systems [7]). Finally, our work should be complemented by an experimental study on the effects of the suggested policy transition on real companies.

⁸In our model, this means that the function f is almost linear. Conversely, when there are more alternatives, f becomes “more concave” and the potential savings for all sides increases.

Acknowledgments

We are indebted to the Transportation Today and Tomorrow organization and to its director Tamar Keinan for sharing with us their databases and ideas, and for their willingness to turn theory into practice. We also thank Ilan Nehama, Adi Meir, and Giora Alexandron for commenting on a draft of this paper.

8. REFERENCES

- [1] Richard J. Arnott, André de Palma, and Charles Robin Lindsey. Departure time and route choice for the morning commute. *Transportation Research, Part B*, 24(3):209–228, 1990.
- [2] Richard J. Arnott and Kenneth Small. The economics of traffic congestion. *American Scientist*, 82:446–455, 1994.
- [3] F. Balbo and S. Pinson. Dynamic modeling of a disturbance in a multi-agent system for traffic regulation. *Decision Support Systems*, 41(1):131–146, 2005.
- [4] Ana Bazzan, Felipe Boffo, and Franziska Klügl. Reducing the effects of the Braess paradox with information manipulation. *Application of Agent Technology in Traffic and Transportation*, pages 85–98, 2005.
- [5] Ana Bazzan and Robert Junges. Congestion tolls as utility alignment between agent and system optimum. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 126–128, 2006.
- [6] Ana Bazzan and Franziska Klügl. Case studies on the Braess paradox: Simulating route recommendation and learning in abstract and microscopic models. *Transportation Research, Part C*, 13:299–319, 2005.
- [7] Ana Bazzan, Joachim Wahle, and Franziska Klügl. Agents in traffic modelling - from reactive to social behavior. *Advances in Artificial Intelligence, no. 1701 in LNAI*, pages 303–306, 1999.
- [8] William R. Black. *Sustainable Transportation: Problems and Solutions*. The Guilford Press, 2010.
- [9] Dietrich Braess. On a paradox of traffic planning. *Transport. Sci.*, 39(4):446–450, 2005.
- [10] John Mc Breen, Pablo Jensen, and Fabrice Marchal. An agent based simulation model of traffic congestion. In *ATT '06: Proceedings of the 4th workshop on Agents in Traffic and Transportation*, pages 43–49, 2006.
- [11] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 82–90. ACM, 2006.
- [12] Claudio Cubillos, Franco Guidi-Polanco, and Claudio Demartini. MADARP: Multi-agent architecture for passenger transportation systems. In *Proceedings of the 8th International IEEE Conference on Intelligent Transportation Systems*, 2005.
- [13] Paul Davidsson, Lawrence Henesey, Linda Ramstedt, Johanna Törnquist, and Fredrik Wernstedt. Agent-based approaches to transport logistics. *Application of Agent Technology in Traffic and Transportation*, pages 1–16, 2005.

- [14] Klaus Fischer, Jörg P. Müller, Markus Pischel, and Darius Schier. A model for cooperative transportation scheduling. In *ICMAS '95: Proceedings of the First International Conference on Multi-Agent Systems*, pages 109–116. The MIT Press, 1995.
- [15] Phil B. Goodwin. Solving congestion. Inaugural Lecture for the Professorship in Transport Policy, University College London, 1997.
- [16] Kerry Hamilton. If public transportation is the answer, what is the question? In Amanda Root, editor, *Delivering Sustainable Transport*, pages 49–59. Elsevier, 2003.
- [17] Garrett Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.
- [18] Israel's central bureau of statistics, motor vehicles 2008. available online from http://www.cbs.gov.il/publications09/1381/pdf/e_print.pdf.
- [19] Anna Nagurney. *Sustainable Transportation Networks*. Edward Elgar Publishing, 2000.
- [20] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [21] Kagan Tumer, Zachary T Welch, and Adrian Agogino. Aligning social welfare and agent preferences to alleviate traffic congestion. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 655–662, 2008.
- [22] UK department of transport, transport statistics 2008. available online from <http://www.dft.gov.uk/pgr/statistics/datatablespublications>.
- [23] UK department of transport, A New deal for Transport: Better for everyone, 2000.
- [24] Ahuva Windsor and Moshe Omer. Changing commuting patterns (in hebrew). available online from http://www.s-t.org.il/index_en.asp, 2007.
- [25] Tomohisa Yamashita, Kiyoshi Izumi, and Koichi Kurumatani. Analysis of the effect of route information sharing on reduction of traffic congestion. *Application of Agent Technology in Traffic and Transportation*, pages 99–111, 2005.