Vestlandsforsking-rapport nr. 10/2012

Fuel consumption in heavy duty vehicles

A report from the Transnova project: "Energy- and environmental savings in Lerum Frakt BA"

Vestlandsforsking, Pb 163, 6851 Sogndal • Tlf.: 906 33 600 • Faks: 947 63 727

Vestlandsforsking rapport

Tittel Fuel consumption in heavy duty vehicles. A report from the Transnova-project: "Energy- and environmental savings in Lerum Frakt BA"	**Rapportnummer** 10/2012 **Dato** 01.01.2012 **Gradering** Open
Prosjekttittel Energi- og miljøbesparende tiltak i Lerum Frank BA	**Tal sider** 95 **Prosjektnr** 6198
Forskar(ar) Morten Simonsen	**Prosjektansvarleg** Morten Simonsen
Oppdragsgivar Statens Vegvesen v/Transnova-programmet	**Emneord** Transport godstransport
Samandrag	
Andre publikasjonar frå prosjektet	
ISBN: 978-82-428-0324-5	**Pris:** Gratis tilgjengeleg på http://www.vestforsk.no

Content

Introduction

Western Norway Research Institute is a partner in the project "Energy- and environmental savings in Lerum Frakt BA" together with Lerum Frakt, a local transport provider responsible for distributing products from Lerum, a maker of jam and lemonade based in Sogndal, Norway. Lerum is located in a region outside the most densely populated areas in Norway where the biggest markets for their products are. Consequently, they need to transport their goods to these markets. Lerum Frakt is responsible for providing this service for Lerum.

The goods are shipped by truck transport. There is no railway available in the Sogndal region and transport by ship to destinations in southern Norway takes too long [1]. Consequently, truck transport is the only transport alternative for Lerum Frakt. The road from Sogndal to the biggest market in the Oslo-region crosses a mountain pass and the trucks have to climb to an altitude of 1100 meter above sea level. There is an alternative route, but it also crosses a mountain pass in an altitude of about 1000 meter above sea level. Both Sogndal and Oslo are situated at sea level so the altitude difference is equal to the highest point on the roads over the mountain passes.

The object of the project "Energy- and environmental savings in Lerum Frakt BA" is to analyze how transport providers like Lerum Frakt can reduce barriers for implementing more environmentally sound transport solutions. Such barriers can be of different kinds. They can be institutional in the sense that public policies make it difficult for transport providers to choose better transport solutions. An example can be taxes on alternative fuels. These taxes inhibit spread of solutions that may be beneficial for a reduction in energy consumption or for mitigation of emission of greenhouse gases or other environmentally harmful gases. Barriers can also be of a practical kind, i.e. distribution of alternative fuels is hindered by lack of relevant facilities. Or barriers can be of a cognitive kind, lack of knowledge about harmful behavioural effects, i.e. from driving patterns, can make it difficult to modify that behaviour.

This document will discuss how driving behaviour can influence energy consumption pr vehicle km. By behaviour we mean driving patterns. Other behavioural factors will also have an influence such as environmental attitudes or economic motivation. They are not discussed in this document, we will only focus on drivers' use of driving patterns that may have an impact on energy consumption per vehicle km. Such driving patterns may be use of cruise control, the amount of idle running, use of automatic gear shift and use of highest gear shifts.

Emissions of greenhouse gases and other environmentally harmful gases are a direct function of fuel consumption since emissions are found by multiplying fuel consumption by some emission factor pr litre fuel. For diesel, the emission factor for CO_2 is 3,17 gram CO_2 per kg which, with a density of 0,84 kg pr litre, gives about 2,66 gram per litre [2]. So if we can understand the driving forces for fuel consumption, we will also understand the related emissions of greenhouse gases and other

[1] For transport to northern Norway, Lerum Frakt uses the ship route called Hurtigruta which visits all major ports from Bergen to Kirkenes in the far north. Destinations for Lerum Frakt north of Trondheim are served by Hurtigruta.

[2] Toutain, J.E.W, Taarneby, G., Selvig, E.,*Energiforbruk og utslipp til luft fra innenlandsk transport*, Statistisk Sentralbyrå, Rapport 2008/49, Table 2.39 and Table 2.1. http://www.ssb.no/emner/01/03/10/rapp_200849/rapp_200849.pdf, see also United States Environment Protection Agency: *Emission Facts: Average Carbon Dioxide Emissions Resulting From Gasoline and Diesel Fuel*, http://www.epa.gov/oms/climate/420f05001.htm

environmentally harmful gases. The energy consumption is the key to understand the complete environmental impacts of transporting goods by truck.

The main research questions addressed in this document are: What are the actual effects of observed driving patterns on fuel consumption? What are the expected effects of any behavioural change in driving patterns? In order to answer these questions, we also need to control for the impact of road infrastructure, road curvature and terrain so that the independent impact of driving behaviour can be assessed. Controlling for these factors means that we compare indicators for driving behaviour for the same infrastructure, road curvature and terrain.

The analysis model

The dependant variable in the analysis, the effect to be explained, is fuel consumption per vehicle km. Another key indicator is energy consumption pr tonne km where an average weight load pr km is assumed. Trucks can have a high energy consumption pr vehicle km and still perform favourable pr tonne km if they on average can carry more load than smaller trucks with lower energy consumption per vehicle km. In general, a higher load factor will yield lower energy consumption pr tonn-km.

Energy consumption is influenced by several factors. Vehicle type is an obvious one. Newer vehicles with improved technology will have less energy consumption pr vehicle km than older vehicles, all other things equal. Maximum torque will also have an impact, we expect trucks with more horse power and torque to have higher energy consumption pr vehicle-km.

Infrastructure is another factor. Better roads with less curvature and less climbing of steep hills will give lower energy consumption per vehicle km. Wider roads and better road surface will also have a positive effect. In Norway, use of spike tyres is very common in the winter season. These tyres wear down the road surface faster than normal tyres. For society as a whole, better infrastructure often implies more transport volume which will cause the positive environmental impact of better infrastructure to be counter-balanced by more energy consumption from a growing number of cars and trucks.

The type of goods being transported is also an important factor. Some of the trucks in the fleet we analyse will often transport bottles back to the lemonade factory in Sogndal. Bottles have a high volume but less weight pr volume. This should reduce the energy consumption pr vehicle km but not pr tonn-km since loads with a higher density (more weight pr volume) will have more weight to distribute the energy consumption on. We have no indicators that can control for this, so we assume a constant weight pr volume ratio for all weight variables used in the analysis.

Load factor is another obvious factor that influences the energy consumption. A higher load factor will mean more weight pr km. This should increase energy consumption pr vehicle km but not pr tonne-km. Generally, a higher load factor is desirable since more goods can be transported with the same amount of trucks, thereby leading to less energy consumption pr tonne of transported good in total.

This reasoning leads to an important observation: A reduction in energy consumption pr vehicle km is in itself not desirable if this reduction comes from less amount of goods being transported. A truck is used for transporting goods. In society as a whole, there is a certain amount of goods that needs to be transported during a specific time span, say a year. We regard this amount as given. The options available to public authorities are therefore to reduce the environmental impacts from transport of goods by selecting the correct transport mean for the correct transport and by using the applied transport means as optimal as possible. Optimal use of trucks for transport purposes will also be beneficial for freight companies, truck owners and drivers who will have more secure and profitable jobs.

The population in our analysis is the theoretical set of all trips that can be made with the group of trucks included in the analysis. This set is theoretical in the sense that it is impossible to observe it, it will nevertheless have a specific meaning as a demarcation of the set of all observations we can possibly make.

In this document, we will analyze energy consumption relative to these driving pattern indicators:

Table 1 List of explanatory variables

- the average driving speed (idle running excepted),
- the relative amount of driving time per day the vehicle is running idle,
- the relative amount of driving time per day the vehicle uses cruise control,
- the relative amount of driving time per day the vehicle is driven with an engine load above 90% of maximum torque,
- the relative amount of driving time per day the vehicle is driven with the highest gear shift,
- the relative amount of driving time per day the vehicle is driven by automatic gear shift,
- the relative amount of driving time per day the vehicle rolls without using engine power,
- the amount of brake applications per 100 km per day
- the relative amount of driving time per day spent driving with a high weight load.

The basic time unit used in this document is one day. This means that we have collected data day by day for each vehicle and for each driver. All data values in the above list are measured as mean per day. All relative time amounts are calculated in percentages.

We will analyze fuel distribution both pr vehicle and pr driver. Some vehicles are shared between drivers. This means that we cannot interpret differences between trucks as differences between drivers. Driving patterns are related to drivers and not trucks. Still, any driver's behaviour is modified through vehicle attributes such as engine size, model type and emission class. Therefore, the interaction of driver and truck will have an impact on drivers' behaviour and also on fuel consumption. For example, to what extent cruise control can be applied is not only determined by driver's intention but also by the vehicle, its load capacity as well as infrastructure and driving conditions. The interaction between driver and vehicle is studied by analyzing fuel consumption both pr truck and pr driver.

For fuel consumption as well as for each of the variables above, we will analyse distributional attributes such as mean, standard error and CV-values.

Equation 1 Mean value (M) for variable F

$$M = \frac{\sum_{i=1}^{n} V_i}{n}$$

Equation 2 Standard error (SE) for variable F

$$SE = \frac{\sqrt{\frac{\sum_{i=1}^{n}(V_i - M)^2}{n-1}}}{\sqrt{n}}$$

Equation 3 CV-Value (Coefficient-of-Variation)

$$CV = \frac{SE}{M}$$

Equation 1-Equation 3 show formulas for computing mean, standard error and CV-value for fuel consumption or any driving pattern indicator in the list above. The symbol n is number of observations included in the analysis. We use the standard error as the dispersion measurement since we are interested in the location of the mean and not any individual observation in the sample distribution. The standard error is the standard deviation for the distribution of theoretical means that can be calculated as we draw repeated samples from the same population.

Table 2 shows data for each vehicle included in the analysis.

Table 2 Vehicle type, vehicle emission class and vehicle engine type

Vehicle	Type	Emission class	Engine type	Horsepower
A	FH13 62T	EM-EC01	D13A480	480
B	FH13 62T	EM-EU5	D13C500	500
C	FH13 62T	EM-EU5	D13C500	500
D	FH13 62T	EM-EU5	D13C500	500
E	FH13 62T	EM-EU5	D13C500	500
F	FH13 62T	EM-EU5	D13C500	500
G	FH13 62R	EM-EU5	D13C540	540
H	FH16 62T	EM-EU5	D16G700	700
I	FH16 64R	EM-EU5	D16G700	700
J	FH13 62T	EM-EU5	D13C500	500
K	FH16 62T	EM-EC01	D16C550	550
L	FH13 62T	EM-EU5	D13C540	540
M	FH13 62T	EM-EC06	D13A520	520
N	Scania	Euro 4		
O	Scania	Euro 5		

We will analyze how fuel consumption vary between different vehicles, between different *types* of vehicles and between vehicles measured by their emission classes and engine type. We will also analyze how fuel consumption vary between drivers. In order to do this we will use analysis of variance which tests whether there are statistically significant differences between group means. The analysis breaks total variance into two components, the variance between groups and the variance inside or within each group.

Equation 4 Variance within groups

$$SSE = \sum_{i=1}^{n_j} \left(x_{ij} - \overline{x}_j\right)^2$$

Equation 5 Variance between groups

$$SSB = n_j * \left(\overline{x}_j - \overline{\overline{x}}\right)^2$$

Equation 4-Equation 5 show the formulas for calculating between-groups and within-group variances for a set of groups in an analysis of variance. In the equations, $\overline{\overline{x}}$ symbolizes the mean for all observations, x_{ij} is observation i for variable x in group j, $\overline{x}_j$ is the mean for group j and n_j is number of observations in group j. A group can consist of vehicles, vehicle types, emission classes, engine types or drivers.

Equation 6 Calculation of F-statistic for variance analysis

$$F = \frac{\frac{SSB}{(k-1)}}{\frac{SSE}{(n-k)}}$$

Equation 6 shows the formula for the test statistic for the analysis of variance. The test statistic is F-distributed with two types of degrees of freedom. The degree of freedom for between-group variance is number of groups (k) minus 1 while the degrees of freedom for variance within groups is number of observations in total minus number of groups.

In a analysis of variance, the null hypothesis is always that the population means for the different groups are equal. The alternative hypothesis is that the means are different in some way. Since the alternative hypothesis does not state any direction, we perform the test as a two-sided test. As for any test, we reject the null hypothesis if the observed probability of getting a test statistic as big as or bigger as the one observed, is less than some predefined level which we call the significance level. The significance level is the maximum acceptable chance of making a Type I error which is rejecting a null hypothesis that is correct. The Type I error is parallel to giving an innocent a guilty verdict in a court, it is the worst possible outcome. The probability of making a Type I error is called significance probability while the predefined maximum acceptable threshold value for this probability is called the significance level. In all our analysis we will use a significance level of 0,05 which means that we accept at most a 5% chance of making a Type I error.

The analysis of variance can tell us whether there are significant differences in mean fuel consumption between drivers or vehicles or vehicle types. In order to analyze which drivers or which vehicles that are different from each other we will extend the analysis by using four measures on so-called post hoc effect analysis. These analysis test what drivers or vehicles are different from each other by performing a set of pairwise tests between each of them. We cannot use pairwise t-tests between means to obtain the same information since the individual error rate involved in any individual paired t-test is not equal to the cumulative error rate involved in making several paired tests simultaneously with the same data material. The null hypothesis applied in all tests assume that all pairs come from the same population and that there are no significant differences between all possible pairs. When performing multiple pairwise tests with the same null hypothesis we must correct the cumulative error rate for the number of comparisons made.

The error rate is the probability that the observed difference between any two means can arise from pure random variation. If this error rate is less than some preset theoretical significance level we will reject the null hypothesis.

The post-hoc analysis will be tested using four measures [3]:

- LSD-test, least-significant-difference which for any pairwise test calculates the minimum difference that will yield a statistical significant difference between the effects in the pair. The formula assumes equal sample size for all samples being tested. If sample sizes are not equal (as is the case in our analysis) a harmonic mean [4] of all sample sizes can be applied.
- Bonferoni test which corrects the significance level for the number of comparisons made, thereby producing a lower real significance level when all comparisons are taken into consideration. Any significance probability (the empirical error rate as opposed to the theoretical significance level) will then be compared to this lower significance level. With the Bonferoni significance level effects must therefore be greater in order to be statistically significant than effects evaluated by the original non-corrected significance level.
- Scheffe's test which produce a critical value for a difference between any pair that must be exceeded in order for that difference to be statistically significant. The calculation for this critical value is different than the LSD-test since the test statistics follow the F-distribution rather than the t-distribution and since sample sizes for any pair are entered directly in the formula rather than as a mean size for all sample sizes.
- Tukey. This is a correction of the LSD-test which uses a special probability distribution for the test statistic. This probability distribution is called the studentized range statistic. The critical value for obtaining statistical significant differences between any pair is based on this probability distribution. The test assumes equal sample size for all samples. This can be corrected for by using a harmonic mean for all sample sizes.

We will also study whether fuel consumption is influenced by driving patterns. Driving pattern indicators will be used as independent variables while energy consumption pr vehicle km is the dependent variable or response variable. The effects of the independent variables are studied using bivariate and multivariate regression analysis. The bivariate models employ only one of the independent variables as explanatory variable against the dependent while the multivariate uses all independent variables simultaneously in the same model.

This gives us Equation 7 for the bivariate case:

Equation 7 The linear bivariate regression model

$$Y = \alpha + \beta X_i + \varepsilon$$

where α and β are parameters in the population and ε denotes the residuals in the model. The subscript i denotes a model with the i'th independent variable where the complete list of all independent variables are displayed above.

[3] See Stevens, J.J.: *Post Hoc Tests in Anova*, http://pages.uoregon.edu/stevensj/posthoc.pdf and Newsom, J.T.: *Post Hoc Tests*, http://www.upa.pdx.edu/IOA/newsom/da1/ho_posthoc.pdf

[4] A harmonic mean of 4 different sample sizes can be calculated as 4/(1/An+1/Bn+1/Cn+1/Dn) where An is number of observations in sample A and so on.

The residuals are deviations from the estimated regression line. The less the magnitude of the deviations, the better is the model fit. We will use OLS-estimates [5] of the regression coefficients, these estimate will yield a smaller sum of squared deviations from the estimated regression line than any other estimates of the same regression coefficients.

The model in Equation 7 is linear in the sense that we expect all observations to lie along a straight line relating values for the independent variable with values for the dependent. We will also look at models with other expected patterns between the independent and the dependent variable.

Equation 8 The inverse bivariate regression model

$$Y = \alpha + \beta \frac{1}{X_i} + \varepsilon$$

Equation 8 shows an inverse bivariate regression model. This model is more appropriate if we expect the process to have some limiting or threshold value that the actual values do not exceed or fall below. This can be a more reasonable model since fuel consumption must have a lower limit, it is possible to save fuel but only up to a point. The marginal effect of fuel consumption is also smaller the less fuel consumption there is to start with. This is also captured by the inverse regression model.

We will start the analysis by analyzing fuel consumption pr vehicle, pr vehicle type and pr emission class. In addition we will analyze fuel consumption pr driver. Then we will present different scatterplots which show the relationship between the dependent and each independent variable. These are the bivariate regression models. The idea is that the bivariate models can tell us something basic and intuitive about the relationship between fuel consumption and the different independent variables or indicators. The bivariate models will be easier to interpret and can therefore be a sensible starting point for further analysis with a more comprehensive multivariate model. Also, the bivariate plots can tell us what functional form to use for the independent variable, whether the effect of the variable is linear or inverse.

We will present scatterplots based on data from trucks and drivers. In the scatterplots based on truck data we will identify each truck in the scatterplot with a colour code. When we present the same scatterplot based on drivers' behaviour each driver will be identified by a colour code.

The bivariate analysis capture the total effect of one independent variable on fuel consumption. This means that when one indicator, say use of cruise control is changing, so does the values of other indicators. When a driver uses more cruise control, he or she is also likely to use more automatic gear shift, spend more of the driving time in highest gear, increase average speed and reduce the amount of time spent with a high engine load and so on. The bivariate plots are appropriate for capturing the total effect of say more use of cruise control on fuel consumption since we allow for other driving behaviour indicators to change their values as well. In this manner, all the effects of increased use cruise control are shown in the bivariate plots. These effects include the direct effect, the indirect effects of cruise control on other indicators who also have an impact on fuel consumption as well as spurious effects of other behavioural indicators which have an effect on both use of cruise control and fuel consumption. Of all these effect, only the direct effect is assessed in the multivariate regression model. This is the separate, independent effect of a change in use of cruise

[5] OLS stands for Ordinary Least Squares

control alone while the values of other behavioural indicators are assumed to be constant. This direct effect can be larger or smaller than the total effect depending on the sign and magnitude of the indirect and spurious effects that make up the total effect together with the direct effect.

Consequently, the bivariate plots are giving us important information about the relationship between one indicator and fuel consumption, but is limited information in the sense that we have to clear the bivariate effect of the indirect and spurious effects in order to find the direct effect of increased use of that driving behavioural indicator.

The multivariate regression model allows for the fact that there is a relationship between the independent variables. If a vehicle is using both cruise control and automatic gear shift, it can be difficult to estimate precisely what is the effect of each independent variable if we only study one at the time. Let's say we find a relationship between automatic gear shifts and fuel consumption. If a vehicle is using more cruise control while also using more automatic gear shift, how can we be sure that the effect we estimate is really for automatic gear shift and not for cruise control? The answer is that we cannot be sure so long as we only estimate fuel consumption by one of the explanatory variables at the time and leave other explanatory variables out of the equation. If we want to know what is the precise effect of one indicator or independent variable we will have to control for the others in the same model.

Similar reasoning can be done for average speed and use of torque above 90% of maximum torque. Obviously higher speed requires more torque. But more torque outtake is also required when trucks are climbing steep hills in low speed. How do we know whether torque or speed is the determining factor? By controlling for the effect of the other variable (say speed) when assessing the impact of one specific variable (say torque). This control is done in a multiple regression model.

Equation 9 shows a multivariate regression model where we apply statistical control for all independent variables in the same model. This means we use all the independent variables at the same time. This will give us a more precise estimate of the separate effect of each of the independent variables. Turning to the example above, we can compare the effect of automatic gear shifts for all vehicles that have the same use of cruise control. If we compare two trucks where one uses automatic gear shifts and the other does not while the both uses cruise control, we know that the difference in fuel consumption between them will be attributable to the difference in gear shift modus and not to use of cruise control since the last indicator does not vary. This is the idea behind using a multivariate model, we can control for the influence of other independent variables while assessing the effect of a single one of them.

Equation 9 The multivariate regression model

$$Y = \alpha + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 \dots \dots + \beta_n X_n + \varepsilon$$

In the multiple regression models we do not control for cargo type. This means that we assume that cargo type is distributed randomly between different vehicles. We also assume that vehicles are driving routes with the same type of infrastructure, road curvature and driving conditions. Any violation of these assumption, i.e. if one vehicle is systematically used more on roads that are steeper and have more curvature than other vehicles, will undermine the conclusions we draw on the effect of driving patterns. We assume that all vehicle have the same chance of having the same cargo type and using the same road.

Data collection

Data are collected using the fleet management system Dynafleet. According to Volvo which has designed and developed the system, it is "(a) system for transport planning combined with vehicle planning, message handling and automatic reporting of vehicle status and driver times." Also according to Volvo, the system "..makes it possible to have a better transport administration and follow-up of the running costs of the vehicle, the work contribution of the driver and how economically the driver drives." [6] It is the last point, to which degree the system allows for driving patterns that are more economical for the driver, the transport operator and not the least the vehicle's owner, that is addressed in this document.

Table 3 Registrations per vehicle

Vehicle	Number of days
A	191
B	300
C	278
D	289
E	288
F	307
G	294
H	315
I	292
J	267
K	142
L	183
M	25
N	10
O	73

The Dynafleet system is implemented in fifteen trucks used by the transport operator Lerum Frakt. We use an anonymous identifier for each vehicle and each driver since we do not want them to be publicly identified. Table 3 shows each vehicle and number of days we have registrations for each of them. Table 4 shows the same data for drivers. We do not show the start date for registrations since this can identify a vehicle or a driver. The last date for data collection was January 14th, 2012. Data items are sent from each vehicle to the transport operator's office using GSM, a mobile phone network. The system is, according to Volvo, developed for Volvo FH and FM vehicles.

Table 4 Registrations per driver

Driver	Count
A	90
AB	1
AC	3
AD	6
AE	1

[6] http://www.dynafleetonline.com/fm/dynafleet-static/online_documents/online_documents.html

AF	14
AG	2
AH	26
AI	24
AJ	12
AK	101
AL	91
AM	62
AN	37
AP	29
AQ	4
AR	1
AS	1
AT	42
AU	22
AV	2
AW	9
AX	15
AY	82
AZ	79
B	143
BA	75
BB	1
BC	5
BD	2
BE	77
BF	65
BG	21
BH	70
BI	61
BJ	36
BK	50
BM	9
BN	1
BO	55
BP	50
BQ	11
BR	12
BS	2
BT	5
C	15
D	136
E	141
F	95
G	114

H	34
I	34
J	143
K	14
L	51
M	186
N	212
O	49
P	37
R	22
S	74
T	41
U	5
V	4
W	3
X	53
Y	16
Z	2
AA	1

Results

The analysis are performed on a per day basis. This indicates that i.e. mean fuel consumption is calculated as mean per day per vehicle or per driver. The same applies to other attributes, the amount of time spent using cruise control is for example calculated relative to total driving time per vehicle or per driver per day. Only vehicles or drivers travelling more than 100 km per day are included in the analysis.

Fuel consumption

Table 5 [7]shows statistics for fuel consumption by vehicle. We note that there are some differences between different vehicles. Vehicle I has the largest mean and it is 1,5 litre per 10 km times bigger than the mean for the vehicle with the lowest mean, vehicle L. Vehicle L has fewer observations than most other vehicles, so the result for this vehicle could be more uncertain. Still, the CV-value for the vehicle does not suggest that it has more variations than other vehicles, rather the opposite seems to be the case.

Table 5 Fuel consumption in litre per 10 km per vehicle

Vehicle	Number of observations	Mean	Std	Cv	Min	Max	Horse-power
A	181	4,3	0,062	1,4	2,2	6	480
B	300	4,6	0,04	0,9	3	6,3	500
C	278	4,6	0,047	1	2,7	6,9	500

[7] We have deleted 10 observations for vehicle A because of unreasonable data items. We refer to appendix A where the observations are documented.

D	289	4,7	0,046	1	2,4	6,8	500
E	288	4,6	0,043	0,9	2,9	6,7	500
F	307	4,7	0,044	0,9	2,2	7,8	500
G	294	4,9	0,048	1	2,6	7,2	540
H	315	4,9	0,045	0,9	2,7	7,4	700
I	292	5,1	0,044	0,9	3,2	7,7	700
J	267	4,7	0,047	1	3	8,3	500
K	142	4,2	0,066	1,6	2,3	5,6	550
L	183	3,6	0,019	0,5	2,2	4	540
M	25	4,3	0,135	3,1	3,4	6,2	520
N	10	5,8	0,36	6,3	3,7	7,9	
O	73	4,3	0,1	2,4	2,6	5,9	
P	64	5,1	0,108	2,1	3,3	7,1	
Total	3308	4,6	0,015	0,3	2,2	8,3	

It should be noted that vehicle I with the larges mean fuel consumption also has the largest engine. There is obviously a trend towards higher consumption the more horsepower a vehicle has. Figure 1 shows this relationship with mean consumption per vehicle on the Y-axis and the vehicle's horsepower on the X-axis. The figure suggests that engine size has an important impact on fuel consumption. Still, the differences between vehicles with same number of horsepower is striking. Vehicles G and L have both 540 horsepower but vehicle G uses 1,3 litre per 10 km more on average. Vehicle G's mean consumption is closer to vehicles with a lot more horsepower such as vehicles I and H. From the database with fuel consumption registered pr driver we can observe that vehicle G is driven by several drivers which could be an important explanation for the vehicle's fuel consumption.

Figure 1 Horsepower vs fuel consumption

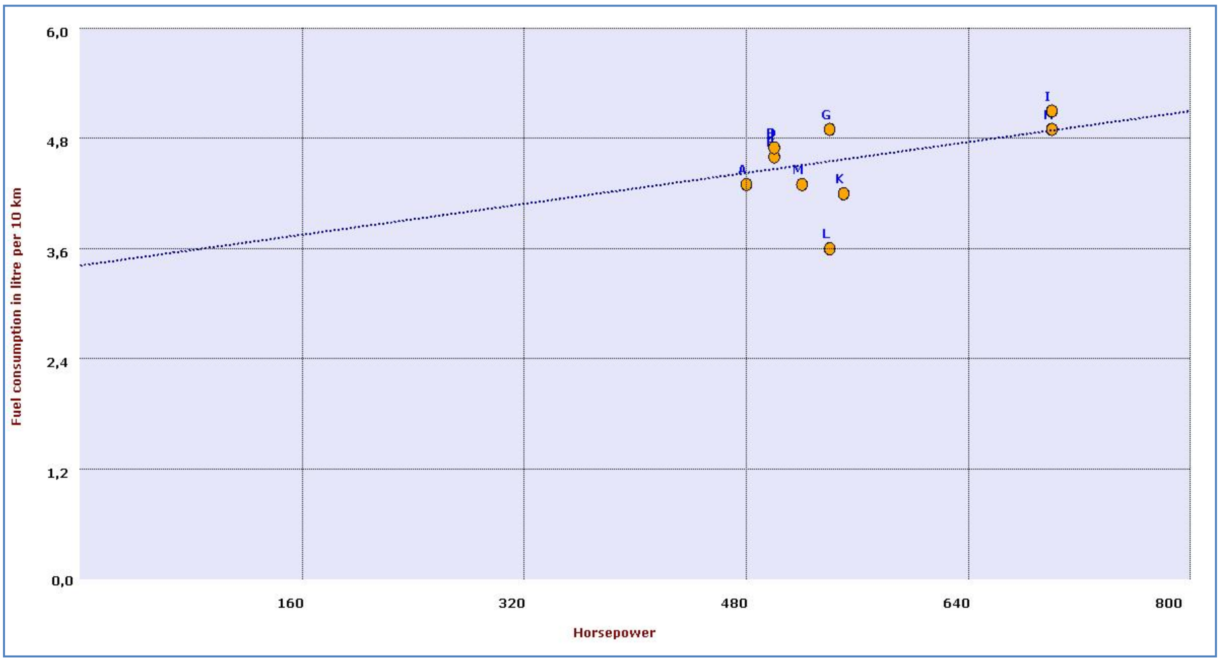

Table 6 Descriptive statistics for fuel consumption per 10 km

N	3308
Mean	4,6
Standard deviation	0,0150
CV	0,3
Min (P0)	2,2
P10	3,6
P25	4,0
P50 (median)	4,6
P75	5,2
P90	5,7
Max	8,3

Table 6 shows descriptive statistics and percentiles for the distribution of fuel consumption per 10 km per vehicle. The P10 value is such that that 10% of all observations in the distribution have lower or equal value to the P10 value. The P25 value is such that 25% of the distribution have lower or equal value and so on. The P90 value is also such that only 10% of the distribution have *higher* values than this value, consequently the P75 value is such that 25% of the distributions have higher value and so on.

As the table shows, 90% of all observations lie above 3,6 litre per 10 km. With a large data set, we can interpret this such that there is only 10% chance of getting a fuel consumption *lower* or equal to 3,6 litre. Accordingly, there is a 50% chance of getting a fuel consumption lower or equal to 4,6 litre per 10 km and there is a 10% chance of getting a fuel consumption *higher* than 5,7 litre per 10 km.

Figure 2 shows a histogram for the distribution of fuel consumption per 10 km. A histogram shows how many percent of a distribution that falls within different value intervals with equal size. Number of intervals is determined as square number of observations included in the analysis, this is the same approach used by i.e. Excel. The histogram shows endpoints for each of the intervals. When there are many intervals, the endpoint of every second interval will be shown. The intervals are found by dividing the value range into a specified number of intervals with equal distance between the endpoint in each interval. The histogram in Figure 2 shows a fairly symmetrical distribution where most values for fuel consumption falls in the intervals from 3 to 6,6 litre per 10 km. It also shows that a very small percentage of the observations (0,7%) have a *higher* value than 6,6 litre per 10 km.

Figure 2 Histogram fuel consumption per 10 km

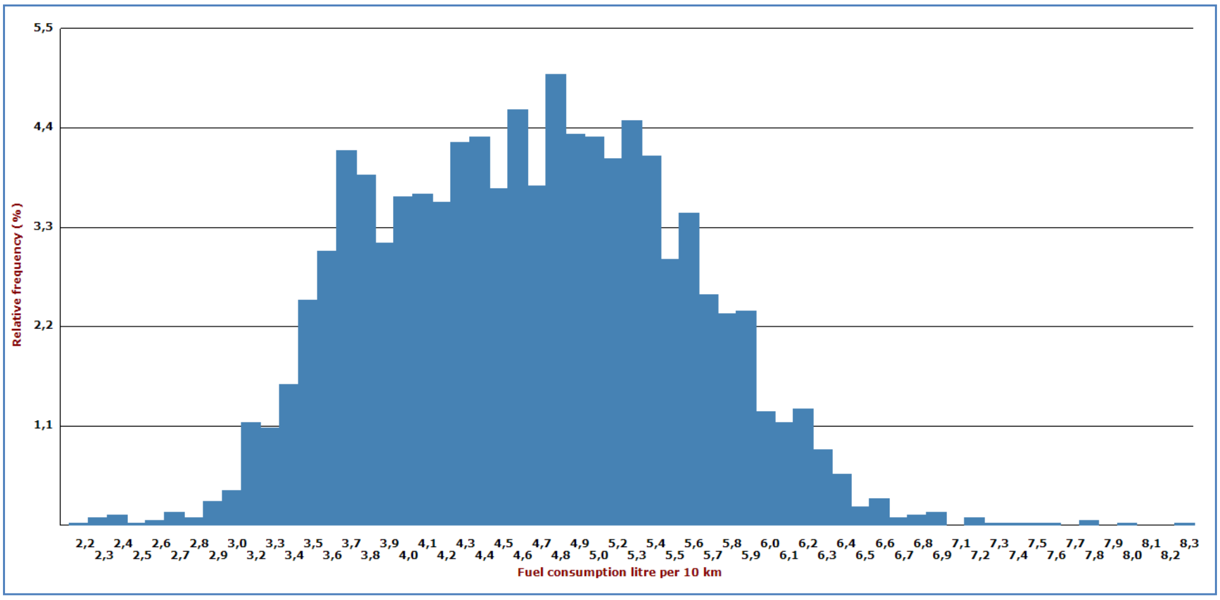

Analysis of variance

Vehicles

Table 7 shows the test statistics for an analysis of variance of fuel consumption by vehicle. The test statistics tell us that there are highly significant differences between the vehicles, some vehicles have a systematic higher fuel consumption than others. In order to test which vehicles this concerns, we will perform post hoc analysis with a pairwise comparison of each effect corrected for the number of comparisons made as described above.

Most vehicles are driven by more than one driver. The results presented for vehicles should therefore not be interpreted as differences between drivers.

Table 7 Analysis of variance (ANOVA) of fuel consumption by vehicle

Variance source	Sum-of-Squares (A)	Degrees-of-freedom (B)	Mean sum of squares (C=A/B)	Test statistic (f)	Significance probability (p)
Between groups	410,614138	14	29,3295813	50,9446619	6,408E-129
Within groups	1890,07075	3283	0,57571451		

Table 8 Differences between vehicles' mean fuel consumption

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,562	4,169	4,263	4,280	4,318	4,576	4,586	4,589	4,678	4,693	4,695	4,909	4,944	5,106
P	5,137	1,575	0,968	0,874	0,857	0,819	0,561	0,551	0,548	0,459	0,443	0,442	0,228	0,193	0,031
I	5,106	1,544	0,937	0,843	0,826	0,788	0,530	0,520	0,517	0,428	0,413	0,411	0,197	0,162	
G	4,944	1,382	0,775	0,681	0,664	0,626	0,368	0,358	0,355	0,266	0,251	0,249	0,035		
H	4,909	1,347	0,740	0,645	0,629	0,591	0,332	0,322	0,320	0,231	0,215	0,214			
F	4,695	1,133	0,526	0,432	0,415	0,377	0,119	0,109	0,106	0,017	0,001				

D	4,693	1,132	0,524	0,430	0,414	0,376	0,117	0,107	0,105	0,015					
J	4,678	1,117	0,509	0,415	0,399	0,360	0,102	0,092	0,090						
C	4,589	1,027	0,420	0,325	0,309	0,271	0,012	0,002							
E	4,586	1,025	0,417	0,323	0,307	0,268	0,010								
B	4,576	1,015	0,407	0,313	0,297	0,258									
A	4,318	0,756	0,149	0,055	0,038										
M	4,280	0,718	0,111	0,016											
O	4,263	0,702	0,094												
K	4,169	0,607													

We have 15 vehicles [8] that can be compared to each other. This gives (15*14)/2 = 105 possible pairwise comparisons. Table 8 shows these pairwise comparisons. The table is ordered by magnitude of mean so that the rows shows the mean in decreasing order with the smallest mean left out. The columns are the same means ordered in increasing order, starting with the smallest and with the highest left out. So the smallest mean is left out in the rows and the largest mean is left out in the columns [9]. The row and column headings include the mens themselves to make the comparisons more transparent.

There are an unequal number of observations between each group or vehicle. To correct for this we have calculated the harmonic mean for number of observations for all vehicles.

Equation 10 Harmonic mean

$$\frac{n}{\sum_{i=1}^{n} \frac{1}{n_i}}$$

Equation 1 shows how the harmonic mean can be calculated for all groups in the analysis. This calculation gives 126,9 number of observations on average per group.

We start by using the LSD-test which gives a critical value for the differences in Table 8. Every difference larger than this critical value is assessed as significant.

$$LSD = t_{\propto} * \sqrt{2 * MSE * n'}$$

where MSE is the mean square error, the mean sum of squares for the within group variance from Table 7 and n' is the harmonic mean for all group sizes. The value $t_{\propto}$ is the critical value from the t-distribution with within-group degrees of freedom from Table 7. The test statistic for the LSD-test is 0,1868 which means that all absolute differences between means greater than this critical value is evaluated as significant.

Table 9 shows which differences are significant according to the LSD-test. Starting with the rows, the table shows that vehicle I has higher fuel consumption than all other vehicles except H and G.

[8] Dynafleet is installed in 16 vehicles, but we have only 10 registrations for vehicle N and they are all from 2010. We have not included vehicle N in the analysis since we consider the number of registrations to be too small over a too limited time span.
[9] See http://pages.uoregon.edu/stevensj/posthoc.pdf

Starting with the columns vehicle L has lower fuel consumption than all other vehicles. Since the table is sorted in increasing order along the columns and deceasing order along the rows, the differences will be smaller as we read the table from left to right. The highest amount of significant differences will therefore be on the upper left side of the table.

Table 9 Significant differences (marked with *) according to the LSD-test

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	*	*	*	*	*	*	*	*	*	*	*	*	*	
I	5,1	*	*	*	*	*	*	*	*	*	*	*	*		
G	4,9	*	*	*	*	*	*	*	*	*	*	*			
H	4,9	*	*	*	*	*	*	*	*	*	*	*			
F	4,7	*	*	*	*	*									
D	4,7	*	*	*	*	*									
J	4,7	*	*	*	*	*									
C	4,6	*	*	*	*	*									
E	4,6	*	*	*	*	*									
B	4,6	*	*	*	*	*									
A	4,3	*													
M	4,3	*													
O	4,3	*													
K	4,2	*													

The Bonferoni test finds the familywise significance level which is the cumulative significance level for a set of groups, a family, which is tested against itself [10]. The individual pre-set significance level is the level for any pairwise comparison, the familywise significance level is the level that all individual pairwise tests must satisfy in order for the individual significance level to be the same for all tests when they are performed together.

Equation 11 Familywise significance level

$$\propto_f = 1 - (1 - \propto)^c$$

where $\propto_f$ is the familywise significance level and c is number of comparisons. This number c is calculated as (n(n-1))/2 where n is number of groups tested [11]. In order to test which pairwise comparisons satisfy the familywise significance level, we calculate the critical t-value for the new familywise significance level. The critical t-value is compared to the empirical t-values from the pairwise t-tests. A t-value for a pairwise test between two groups is found by applying Equation 12.

Equation 12 T-test for a pairwise comparison

[10] See http://www.upa.pdx.edu/IOA/newsom/da1/ho_posthoc.pdf
[11] A vehicle or a driver is also a group since there are many registrations for each vehicle and driver.

$$t = \frac{\bar{X}_1 - \bar{X}_2}{\sqrt{\frac{S_1^2}{n_1} + \frac{S_2^2}{n_2}}}$$

Equation 13 Bonferoni new pairwise significance level

$$\propto_p = \frac{\propto_f}{c}$$

Equation 13 shows the new pairwise significance level corrected for the number of comparisons made. The new pairwise critical t-value is found by using the new significance level from Equation 13 and errorwise degrees of freedom, 3283, from Table 7.

Table 10 Pairwise t-tests

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	14,43	7,68	5,94	4,98	6,59	4,89	4,75	4,67	3,91	3,79	3,80	1,95	1,64	0,26
I	5,1	32,21	11,84	7,70	5,83	10,33	8,92	8,39	8,02	6,63	6,44	6,58	3,12	2,48	
G	4,9	26,86	9,53	6,13	4,65	7,97	5,91	5,53	5,29	3,96	3,76	3,82	0,54		
H	4,9	27,41	9,26	5,87	4,43	7,66	5,51	5,13	4,89	3,52	3,31	3,37			
F	4,7	23,62	6,64	3,94	2,93	4,94	2,00	1,75	1,65	0,26	0,02				
D	4,7	22,65	6,52	3,90	2,91	4,84	1,92	1,69	1,59	0,23					
J	4,7	22,07	6,30	3,75	2,80	4,62	1,66	1,44	1,35						
C	4,6	20,29	5,19	2,94	2,17	3,47	0,20	0,04							
E	4,6	21,64	5,29	2,96	2,17	3,53	0,17								
B	4,6	23,11	5,30	2,90	2,11	3,50									
A	4,3	11,63	1,64	0,46	0,26										
M	4,3	5,28	0,74	0,10											
O	4,3	6,88	0,79												
K	4,2	8,89													

Table 10 shows the pairwise t-test for the vehicle tests. The new significance level is 0,0095 which gives a new critical t-value of 2,6 with 3283 degrees of freedom. Table 11 shows which pairwise comparisons are evaluated as significant with the new critical t-value from the Bonferoni test. We see much the same picture as in Table 9. All in all there are 76 significant differences with the Bonferoni test while there were 80 significant tests with the LSD-test. L has one more significant difference against K while vehicle M has three less significant tests (against vehicles C,E,B) and vehicles H and G has one less significant test each (both against vehicle P).

Table 11 Significant differences (marked with *) according to the Bonferoni-test

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
	Diff	3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	*	*	*	*	*	*	*	*	*	*	*			
I	5,1	*	*	*	*	*	*	*	*	*	*	*	*		

G	4,9	*	*	*	*	*	*	*	*	*	*	*			
H	4,9	*	*	*	*	*	*	*	*	*	*	*			
F	4,7	*	*	*	*	*									
D	4,7	*	*	*	*	*									
J	4,7	*	*	*	*	*									
C	4,6	*	*	*		*									
E	4,6	*	*	*		*									
B	4,6	*	*	*		*									
A	4,3	*													
M	4,3	*													
O	4,3	*													
K	4,2	*													

The Scheffe test calculates a new critical value for the pairwise differences between the vehicles. If the actual difference is greater than this critical value, the difference is statistically significant with the individual significance level satisfied for all comparisons simultaneously. Scheffe's test statistic calculates one critical difference for each pair and not one critical value for all comparisons. Also, Scheffe's test statistic is the only test statistic so far for post hoc comparisons that take the actual pairwise sample sizes into the calculation. Scheffes test statistic does not use an approximation of an equal sample size for all comparisons as the other test statistics do.

Equation 14 The critical Scheffe value for difference between mean i and mean j

$$\sqrt{k * f_{crit}} * \sqrt{MSE\left(\frac{1}{n_i}+\frac{1}{n_j}\right)}$$

Equation 14 shows the formula for the calculation of the critical value for the difference between mean i and mean j where k is number of groups minus one (which is 14) and f_{crit} is the critical value from the f-distribution with degrees of freedom equal to k and v where v is within group variance degrees of freedom from Table 7 (3283). The critical f-value is 1, 1,6948 in this case. MSE is *within group variance mean sum of squares* from Table 7 (0, 0,576).

Table 12 Scheffe critical values for pairwise differences

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	0,537	0,556	0,633	0,872	0,537	0,509	0,511	0,512	0,514	0,511	0,508	0,507	0,510	0,510
I	5,1	0,348	0,378	0,484	0,770	0,350	0,304	0,307	0,310	0,313	0,307	0,302	0,300	0,305	
G	4,9	0,348	0,378	0,483	0,770	0,349	0,303	0,306	0,309	0,312	0,306	0,302	0,300		
H	4,9	0,344	0,374	0,480	0,768	0,345	0,298	0,301	0,304	0,307	0,301	0,296			
F	4,7	0,345	0,375	0,481	0,769	0,346	0,300	0,303	0,306	0,309	0,303				
D	4,7	0,349	0,379	0,484	0,770	0,350	0,305	0,308	0,310	0,314					
J	4,7	0,355	0,384	0,488	0,773	0,356	0,311	0,314	0,317						
C	4,6	0,352	0,381	0,486	0,772	0,353	0,308	0,311							

E	4,6	0,349	0,379	0,484	0,771	0,351	0,305								
B	4,6	0,347	0,376	0,482	0,769	0,348									
A	4,3	0,387	0,414	0,512	0,789										
M	4,3	0,788	0,802	0,856											
O	4,3	0,512	0,532												
K	4,2	0,413													

Table 12 shows the critical values for the pairwise differences. If we contrast these critical values with the actual differences in Table 8 we get Table 13 which shows what differences are statistically significant. These are marked with an asterisk (*) in the table. We observe that this is a more conservative test since there are fewer statistically significant differences than what we found with the LSD test and the Bonferoni test.

Table 13 Significant differences (marked with *) according to the Scheffe-test

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	*	*	*		*	*	*	*						
I	5,1	*	*	*	*	*	*	*	*	*	*	*			
G	4,9	*	*	*		*	*	*	*						
H	4,9	*	*	*		*	*	*	*						
F	4,7	*	*			*									
D	4,7	*	*			*									
J	4,7	*	*			*									
C	4,6	*	*												
E	4,6	*	*												
B	4,6	*	*												
A	4,3	*													
M	4,3														
O	4,3	*													
K	4,2	*													

At last we perform the Tukey test on the same pairwise comparisons. The Tukey test uses a critical value from the studentized range distribution [12]. The degrees of freedom for this critical value are number of groups (in this case 15) and *within group variance degrees of freedom* from Table 7 (which is 3283). By using a look-up table for the studentized range distribution [13] we find the critical value q to be 5,45.

Equation 15 The Tukey HSD test statistics

[12] http://en.wikipedia.org/wiki/Tukey's_range_test

[13] See http://www.watpon.com/table/studen_range.pdf . For the same degrees of freedom, this web site http://cse.niaes.affrc.go.jp/miwa/probcalc/s-range/srng_tbl.html reports smaller critical value for i.e. k=15 and v=Infinity. We have assumed that the smaller critical value is due to one-sided as opposed to two-sided test. With this assumption, we have used the two-sided critical value in this case. The value for v is set to infinity since the table stops at specific values for v = 120.

$$HSD = q * \sqrt{\frac{MSE}{n'}}$$

Equation 15 shows the formula for the Tukey HSD test statistic. HSD stands for honest significant difference. The test statistic is a critical value so that any pairwise differences exceeding this test statistic are evaluated as significant. The Tukey test statistic requires equal sample size, but for unequal sample sizes a common sample size can be approached by using the harmonic mean of all sample sizes as discussed above [14].

Table 14 Significant differences (marked with *) according to the Tukey hsd-test

		L	K	O	M	A	B	E	C	J	D	F	H	G	I
		3,6	4,2	4,3	4,3	4,3	4,6	4,6	4,6	4,7	4,7	4,7	4,9	4,9	5,1
P	5,1	*	*	*	*	*	*	*	*	*	*	*			
I	5,1	*	*	*	*	*	*	*	*	*	*	*			
G	4,9	*	*	*	*	*	*								
H	4,9	*	*	*	*	*									
F	4,7	*	*	*	*	*									
D	4,7	*	*	*	*	*									
J	4,7	*	*	*	*										
C	4,6	*	*												
E	4,6	*	*												
B	4,6	*	*												
A	4,3	*													
M	4,3	*													
O	4,3	*													
K	4,2	*													

Table 14 shows the result of the Tukey test. We find that this result is roughly the same as for the Scheffe test. Interestingly, the Tukey test statistics gives more statically significant results for vehicles M and O which have the smallest number of observations . Thus, a reasonable proposition is that Scheffe's test is more accurate since it takes sample sizes into account while the other test statistics use an approximation of an equal sample size. The harmonic mean used for this approximation gives a sample size that is far from the actual sample size of most vehicles since the three vehicles M, P and O have very few observations but are still given the same weight as the other vehicles for all test statistics except Scheffe's.

Table 15 Number of significant comparisons with the different post hoc tests

	Number of significant comparisons	Relative
LSD	81	77,1 %

[14] See Stevens, J., J.: *Post Hoc Tests In Anova*, http://pages.uoregon.edu/stevensj/posthoc.pdf

Bonferoni	76	72,4 %
Scheffe	50	47,6 %
Tukey	57	54,3 %

Table 15 gives a summary of the post hoc tests presented here. The table shows the number of comparisons evaluated as statistically significant with the different tests. The number of possible comparisons is 105. The LSD test evaluates 77,1% of all comparisons as significant while the same number for the Scheffe test is 47,6% . This is quite a difference and it is striking that the only test which takes the actual sample sizes into consideration is the most conservative one. We have very unequal sample sizes, and based on this judgment we select Scheffe's test as the most appropriate one. In the following, we will therefore only use Scheffe's test for post hoc comparisons.

Vehicle engine type

Table 16 shows descriptive statistics for fuel consumption per 10 km distributed on vehicle engine type. There are seven different types. Each engine type has a distinct number of horsepower. As the table shows, the engine type with the highest number of horsepower also have the highest mean fuel consumption, except for Iveco where number of horsepower is unknown. Still, the engine type with the smallest consumption is not the engine type with lowest number of horsepower. Also, the fuel consumption for the vehicle type with 550 horsepower is smaller on average than the fuel consumption for the vehicle type with 480 horsepower.

Table 16 Fuel consumption by vehicle engine type

Engine type	Horse - power	Count	Mean	Std	Cv	Min	Max
D13A480	480	181	4,3	0,838	19,4	2,2	6,0
D13A520	520	25	4,3	0,673	15,7	3,4	6,2
D13C500	500	1729	4,6	0,758	16,3	2,2	8,3
D13C540	540	477	4,4	0,945	21,4	2,2	7,2
D16C550	550	142	4,2	0,783	18,8	2,3	5,6
D16G700	700	607	5,0	0,788	15,7	2,7	7,7
Scania		83	4,4	1,013	22,8	2,6	7,9
Iveco		64	5,1	0,860	16,7	3,3	7,1

Table 17 shows the test statistics for the analysis of variance of fuel consumption by vehicle engine type. The table shows that we find a significant difference in fuel consumption between different vehicle engine types since the significance probability is lower than the usual threshold value (significance level) of 0,05. This indicates that the differences between the different vehicle types are systematic and not random. Some vehicle engine types have lower fuel consumption than others.

Table 17 Analysis of variance (ANOVA) of fuel consumption by vehicle engine type

Variance source	Sum-of-Squares (A)	Degrees-of-freedom (B)	Mean sum of squares (C=A/B)	Test statistic (f)	Significance probability (p)	Variance source

Between groups	177,186017	7	25,3122881	38,8905592	8,5795E-53	2,01235393
Within groups	2147,8362	3300	0,65085945			
Total	2325,02222	3307				

Table 18 shows pairwise comparisons for the vehicle engine types. In all, there are 8 engine types which gives 8*(8-1)/2=28 possible comparisons. As discussed above, we restrict ourselves to a Scheffe post hoc analysis of the possible comparisons. To find the Scheffe critical value for each comparison, the value that when exceeded will give a statistically significant result, we use two degrees of freedom. These are number of groups minus 1, which is 7, and degrees of freedom for the within group variance from Table 17, which is 3300. This gives a critical f-value of 2,012. By applying these values on the individual sample sizes as described in Equation 14, we can produce Table 19 which shows the critical value for each pairwise comparison.

Table 18 Pairwise comparisons for vehicle engine types

		D16C550	D13A520	D13A480	D13C540	Scania	D13C500	D16G700
		4,2	4,3	4,3	4,4	4,4	4,6	5,0
Iveco	5,1	0,968	0,857	0,819	0,723	0,693	0,501	0,133
D16G700	5,0	0,835	0,724	0,686	0,590	0,560	0,367	
D13C500	4,6	0,467	0,357	0,318	0,223	0,193		
Scania	4,4	0,274	0,164	0,126	0,030			
D13C540	4,4	0,245	0,134	0,096				
D13A480	4,3	0,149	0,038					
D13A520	4,3	0,111						

Table 19 Scheffe critical values for pairwise differences

		D16C550	D13A520	D13A480	D13C540	Scania	D13C500	D16G700
		4,2	4,3	4,3	4,4	4,4	4,6	5,0
Iveco	5,1	0,456	0,714	0,440	0,403	0,504	0,385	0,398
D16G700	5,0	0,282	0,618	0,256	0,185	0,354	0,143	
D13C500	4,6	0,264	0,610	0,237	0,157	0,340		
Scania	4,4	0,418	0,691	0,401	0,360			
D13C540	4,4	0,289	0,621	0,264				
D13A480	4,3	0,339	0,646					
D13A520	4,3	0,657						

Table 20 Significant differences (marked with *) according to the Scheffe-test for vehicle engine types

		D16C550	D13A520	D13A480	D13C540	Scania	D13C500	D16G700
		4,2	4,3	4,3	4,4	4,4	4,6	5,0
Iveco	5,1	*	*	*	*	*	*	
D16G700	5,0	*	*	*	*	*	*	

D13C500	4,6	*		*	*			
Scania	4,4							
D13C540	4,4							
D13A480	4,3							
D13A520	4,3							

Table 20 shows the result table from the Scheffe test where each statistically significant pairwise comparison is marked with an asterisk (*). Of the total 21 possible comparisons, 15 are statistically significant. The engine type with the highest number of horsepower and the Iveco vehicle have a significantly higher fuel consumption than all other engine types except for the comparison between themselves. The engine type with 500 horsepower has a significantly higher fuel consumption than three of the four engine types with the lowest fuel consumption.

Figure 3 shows the relationship between number of horsepower and mean fuel consumption for the different engine types. As already mentioned, the engine type D16G700 with 700 horsepower has a significantly higher consumption than all other engine types. A difference of 220 in number of horsepower (the difference between the highest and lowest group) yields a difference in fuel consumption of almost 0,7 litre per 10 km. On the other hand, the engine type with next highest number of horsepower has the lowest mean fuel consumption.

Figure 3 Mean fuel consumption and number of horsepower for different engine types

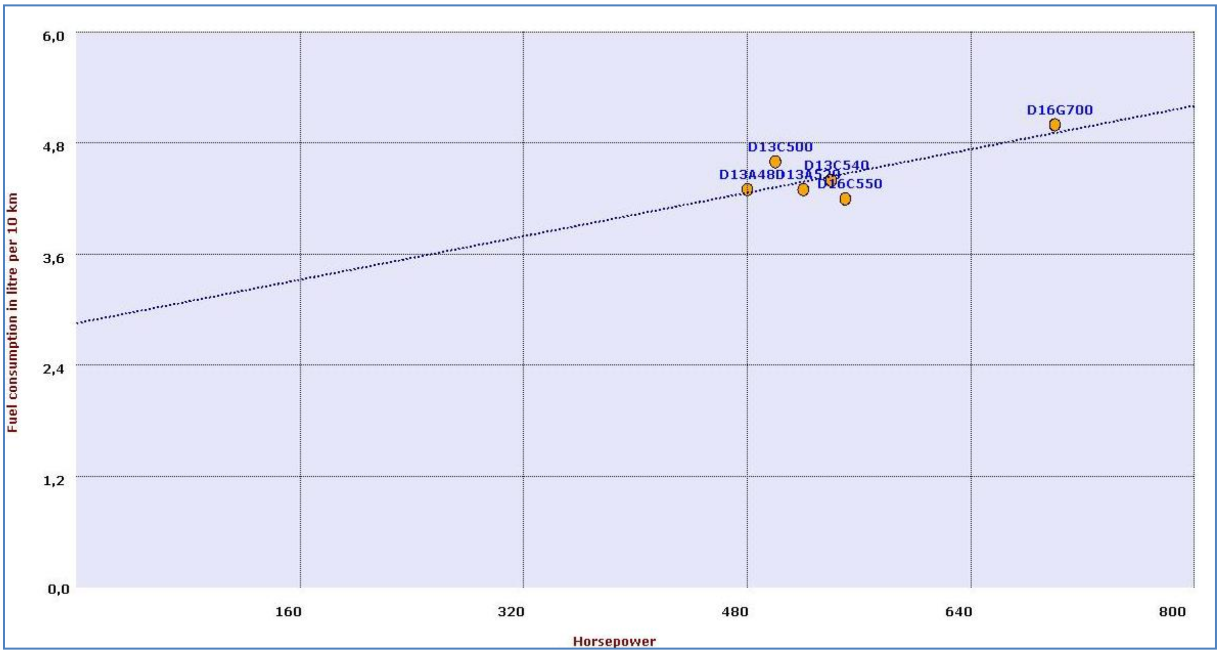

Drivers

We have already analyzed differences in fuel consumption between vehicles. In this section, we will look at differences in fuel consumption between drivers. Is there a significant difference between drivers? First we will look at drivers independent of the vehicles they drive. Then we will analyze whether there is a difference in fuel consumption when we distribute drivers on vehicles. Since drivers drive more than one vehicle, it is reasonable to control for differences in vehicles when we analyze differences between drivers. If not, we risk assessing differences in vehicle properties as differences in driving behaviour.

We have registered fuel consumption, driving distance and driving behaviour such as use of cruise control, use of highest gear etc both for vehicles and drivers. Consequently, we have one database for drivers and one for vehicles [15]. There are many drivers working for Lerum Frakt AB during one year. Some of them drive only for a short period. In the analysis of drivers we have only included drivers that have more than 100 registrations in Dynafleet in the interval January 2011 to January 2012. Of course, the number of registrations will depend on when the Dynafleet system was installed in the vehicles. Also, as for vehicles, only trips longer than 100 km per day are included.

Table 21 Fuel consumption litre per 10 km pr driver

Driver	Count	Mean	Std err	Cv	Min	Max
D	134	4,1	0,066	1,6	3,1	6,6
E	130	4,7	0,059	1,3	3,1	6,5
G	108	4,7	0,086	1,8	2,2	7,7
J	136	4,5	0,064	1,4	2,9	6,5
M	173	4,6	0,064	1,4	2,4	6,9
N	210	4,6	0,047	1	3	6,4

Table 21 shows fuel consumption in litre per 10 km for drivers independent of the vehicles they drive. Driver D has the lowest fuel consumption on average while driver E and G have the highest consumption. The difference between largest and smallest mean fuel consumption is about 0,6 litre per 10 km.

Table 22 Analysis of variance (ANOVA) of fuel consumption by driver

Variance source	Sum-of-Squares (A)	Degrees-of-freedom (B)	Mean sum of squares (C=A/B)	Test statistic (f)	Significance probability (p)
Between groups	34,2080005	5	6,84160011	11,7023103	5,6831E-11
Within groups	517,403482	885	0,5846367		
Total	551,611483	890			

Table 22 shows the analysis of variance for fuel consumption by driver. The ANOVA summary statistics show clearly that there are significant differences between drivers.

Table 23 Differences between drivers in fuel consumption (litre per 10 km)

		D	J	N	M	G
		4,1	4,5	4,6	4,6	4,7
E	4,7	0,626	0,243	0,154	0,134	0,026

[15] Technically speaking, we have two different tables in the same database.

G	4,7	0,600	0,217	0,128	0,108	
M	4,6	0,492	0,109	0,020		
N	4,6	0,472	0,090			
J	4,5	0,382				

Table 23 shows the pairwise differences between all drivers. As already mentioned, the highest difference is about 0,626 litre pr 10 km while the smallest is negligible, 0,026 litre per 10 km. The table is ordered so that the biggest differences are in the upper left part of the table. The table also shows that the biggest differences are between driver D on one hand and drivers G and E on the other.

Table 24 Scheffe critical values for pairwise differences

		D	J	N	M	G
		4,1	4,5	4,6	4,6	4,7
E	4,7	0,314	0,313	0,285	0,296	0,332
G	4,7	0,330	0,329	0,302	0,313	
M	4,6	0,293	0,292	0,262		
N	4,6	0,282	0,281			
J	4,5	0,310				

Table 24 shows Scheffes critical difference for the different pairwise comparisons. If we compare the critical differences in Table 24 with the actual ones in Table 23, we see that there are statistically significant differences between driver D and all the other drivers. No other pairwise comparisons have significant effects which mean the differences may just as well be pure random variations.

What if we include the vehicles driven by the drivers in the analysis? Table 25 shows descriptive statistics for the combination of driver and vehicle. Drivers and vehicles are coupled by selecting on same date, same distance, same fuel consumption and same amount of time driven in the two database tables, one for vehicles and one for drivers. Not every trip the drivers make can be identified by a vehicle, therefore the amount of days driven will not be identical per driver in Table 25 and Table 21,

Table 25 Fuel consumption litre per 10 km pr driver and vehicle

Driver	Vehicle	Count	Mean	Std err	Cv	Min	Max	Horse-power
D	E	2	5,1	0,159	3,1	4,9	5,2	500
D	J	38	4,6	0,113	2,5	3,1	5,9	500
D	L	6	3,5	0,075	2,1	3,2	3,7	540
E	B	8	5	0,173	3,5	4,1	5,6	500
E	C	4	4,8	0,323	6,8	4,1	5,5	500
E	F	103	4,7	0,068	1,5	3,1	6,5	500
G	F	98	4,7	0,086	1,8	2,2	6,4	500
J	E	124	4,4	0,065	1,5	2,9	6,5	500
M	C	150	4,6	0,063	1,4	3,2	6,9	500

N	B	500	191	4,5	0,046	1	3	500

We can now perform an analysis of variance on the combination of driver and vehicle. This amounts to specifying interaction terms in the analysis since the effect of driving behaviour per driver is modified by the vehicles they are driving. Driver and vehicle interact to produce the outcome of the analysis. All in all we have 10 combinations and we can indentify each combination by the driver's letter and the vehicle's so that DE is driver D in vehicle E while NB is driver N in vehicle B and so on. First letter is the driver's and the second is the vehicle's. An interesting point is whether vehicles driven by different drivers score very different on mean fuel consumption. Is there any difference between fuel consumption for vehicle B when it is driven by driver N and by driver E? Or is there a difference in vehicle F's performance when it is driven by two different drivers E and G?

Table 26 Analysis of variance for fuel consumption by combination driver/vehicle

Variance source	Sum-of-Squares (A)	Degrees-of-freedom (B)	Mean sum of squares (C=A/B)	Test statistic (f)	Significance probability (p)
Between groups	15,2608796	9	1,69565329	3,26066285	0,00067389
Within groups	371,303782	714	0,52003331		

Table 27 Pairwise comparisons of driver/vehicle interaction

		DL	JE	NB	DJ	MC	GF	EF	EC	EB
		3,5	4,4	4,5	4,6	4,6	4,7	4,7	4,8	5,0
DE	5,1	1,561	0,658	0,553	0,471	0,456	0,400	0,377	0,319	0,081
EB	5,0	1,480	0,577	0,473	0,391	0,376	0,320	0,297	0,238	
EC	4,8	1,242	0,339	0,235	0,153	0,138	0,082	0,059		
EF	4,7	1,184	0,280	0,176	0,094	0,079	0,023			
GF	4,7	1,161	0,257	0,153	0,071	0,056				
MC	4,6	1,105	0,201	0,097	0,015					
DJ	4,6	1,090	0,186	0,082						
NB	4,5	1,008	0,104							
JE	4,4	0,903								

Table 27 shows the pairwise comparisons for all combinations of driver and vehicle. It must be stressed that many of the combinations have very few observations, hence the table should be interpreted with caution.

Table 28 shows the critical values for pairwise comparisons for combinations of driver and vehicles. The critical values are large since there are few observations behind many combinations. As the formula for Scheffe critical values in Equation 14 shows, if the number of observations (n_i or n_j) are small, the resulting fraction 1 over number of observations will be larger and the resulting critical value will also be larger. Therefore, with a smaller number of observations in each group it takes a larger difference for the effect to be statistically significant.

An inspection of the actual pairwise comparisons from Table 27 with the critical ones in Table 28 show that no pairwise comparisons are significant. The differences between driver/vehicle pairs may just as well arise from pure chance. The differences we found between drivers disappear when vehicles are taken into the analysis. It could be that number of observations in each combination of driver and vehicle is so small that it is harder to find statistically significant differences between them. Therefore, it may be that some differences turn out to be significant at a later stage when more data are available.

Table 28 Scheffe critical values for pairwise comparisons driver/vehicle

		DL	JE	NB	DJ	MC	GF	EF	EC	EB
		3,5	4,4	4,5	4,6	4,6	4,7	4,7	4,8	5,0
DE	5,1	2,430	2,122	2,116	2,159	2,119	2,126	2,125	2,578	2,353
EB	5,0	1,608	1,086	1,074	1,158	1,080	1,094	1,092	1,823	
EC	4,8	1,921	1,512	1,504	1,565	1,508	1,518	1,517		
EF	4,7	1,250	0,397	0,364	0,565	0,381	0,420			
GF	4,7	1,252	0,402	0,370	0,569	0,387				
MC	4,6	1,239	0,361	0,325	0,541					
DJ	4,6	1,308	0,552	0,529						
NB	4,5	1,234	0,343							
JE	4,4	1,244								

Table 29 Fuel consumption pr driver with more than 50 working days

Driver	Count	Mean	Std	Cv	Min	Max
A	86	4,9	0,08	1,6	3,5	6,4
AK	90	4,9	0,073	1,5	3,1	7,7
AL	89	5	0,076	1,5	3,5	6,3
AM	56	3,7	0,062	1,6	3,1	5,2
AY	79	5	0,088	1,8	3,4	6,6
AZ	73	5	0,088	1,8	3,2	6,9
B	89	3,7	0,029	0,8	3,2	4,7
BA	74	4,7	0,089	1,9	2,6	6,6
BE	72	5,1	0,107	2,1	3,3	7,1
BF	58	4,3	0,117	2,7	2,7	6
BH	58	4,5	0,087	1,9	3,3	6,2
BI	59	5,2	0,119	2,3	2,9	7,1
BO	51	4,7	0,103	2,2	3,5	6,3
D	134	4,1	0,066	1,6	3,1	6,6
E	130	4,7	0,059	1,3	3,1	6,5
F	77	4,6	0,107	2,3	2,3	6,5
G	108	4,7	0,086	1,8	2,2	7,7
J	136	4,5	0,064	1,4	2,9	6,5
M	173	4,6	0,064	1,4	2,4	6,9
N	210	4,6	0,047	1	3	6,4

S	74	4,6	0,084	1,8	3	6,4
X	50	3,6	0,06	1,7	2,4	5,3

Table 29 shows fuel consumption for drivers with 50 working days in the registration period. We include this table in order to show how representative the drivers with 100 working days are relative to other drivers. We note that the variation is bigger when we include drivers with fewer working days. The difference between maximum and minimum mean value is about 0,6 litre pr 10 km when we analyze drivers with 100 working days while the same difference is 1,6 litre per 10 km when we include drivers with 50 working days.

Winter season

We define winter season as the months December, January, February and March. Since Sogndal is located as far north as Fairbanks, Alaska, the winter season is long with many snow storms and cold weather. Since the vehicles are travelling over mountain passes with narrow, winding steep roads the winter season can cause very challenging driving conditions. In cold weather the challenge is to keep the vehicle and driver's cabin warm during delivery stops.

Table 30 Analysis of variance (ANOVA) of fuel consumption by time of year

Variance source	Sum-of-Squares (A)	Degrees-of-freedom (B)	Mean sum of squares (C=A/B)	Test statistic (f)	Significance probability (p)
Between groups	78,7450186	1	78,7450186	120,723817	1,6645E-27
Within groups	1761,14006	2700	0,6522741		
Total	1839,88508	2701			

Table 30 shows the result of a analysis of variance with fuel consumption as dependent variable and time of year as factor or independent variable. The table shows a clear seasonal variations in fuel consumption, the effect of winter months is highly significant. A regression analysis with fuel consumption as dependent and a dummy for winter months as independent shows that in winter months, the fuel consumption is 0,40 litre higher per 10 km on average. The t-value for the regression coefficient is the square root of the f-statistic from the ANOVA-analysis in Table 30 with degrees of freedom equal to 2700 and the p-value (significance probability) is exactly the same.

In this analysis we have not controlled for other driving indicators like average speed, use of cruise control, automatic gear shift, weight load and so on. We have only taken time of year into consideration. This is the bivariate, total effect of winter months. The driving indicators will also be influenced by time of year, it is highly probable that i.e. use of cruise control will be easier in summer months than in winter months. Therefore there are several indirect effects of winter season, effects that are mediated through other independent variables. The direct effect of winter months is the effect that comes only from winter months and is not mediated through other independent variables. We will estimate this direct effect in the section on multivariate regression analysis.

Bivariate scatter plots

We will now turn to bivariate regression analysis in order to study how fuel consumption is related to indicators for driving patterns. We have so far confirmed that there are differences between vehicles and drivers, but we have not given any indications as to how these differences may be explained. Only daily distances bigger than 100 km are included in the analysis, both for drivers and vehicles.

A scatterplot is a plot where the dependent variable (the effect variable) is on the Y-axis and the independent variable (explanatory variable) is on the X-axis. The Y-axis is fuel consumption in litre per 10 km in all plots. The plot will show how and to what extent the dependent variable varies with the independent. By using such plots, we can detect trends in the co-variation between the two variables. Such trends can be important for specifying the functional form for the effect of X on Y. The functional form tells us whether the effect is linear or curvi-linear in some way. A linear effect is such that an increase of one unit in X always generates the same average response in the Y-variable. A curvi-linear effect is such that the average effect on Y of one unit change in X is dependent on where on the X-axis the effect is evaluated.

Scatterplots only show the relationship between two variables. There is no control for third variables in the statistical sense. In order to perform such a control we have to use multivariate regression where the effect of several explanatory variables are evaluated simultaneously. In spite of this, scatterplots are meaningful in order to visualize the impact on any X on Y. A scatterplot can show us what form the relationship between the X and Y variables have and the degree of variation in that relationship.

We will present two scatterplots for each independent variable. One is for vehicles and the other for drivers. Any discrepancy between the two plots can give us an idea of the significance of either drivers or vehicles for the relationship we study. If the effect of increasing use of cruise control on fuel consumption is different for vehicles and for drivers, this can tell us whether driver behaviour or vehicle attributes are most important for the impact of cruise control use on fuel consumption. The interaction of driver and vehicle can be better assessed by presenting comparable plots for vehicles and drivers.

We will present plots where each vehicle or driver is represented in the plot with a colour code. When we evaluate the relationship between, say, average speed on fuel consumption per 10 km, each truck or driver will be marked with a distinct colour in the plot. This allow us to assess whether some vehicles or drivers constantly lie above or below the estimated regression line.

Cruise control

Table 31 shows descriptive statistics for the relative amount of driving time cruise control was in use for vehicles and drivers. On average, cruise control was used in 5,6% of total driving time for vehicles and 6,4% of driving time for drivers. The maximum relative amount of driving time was about 86-87% for both drivers and vehicles while the minimum was zero percent for both. The symbol P10 stands for the 10th percentile which is the value which is such that 10% of the vehicles have a lower or equal use of cruise control. Correspondingly ,75% of the vehicles have a use of cruise control that is lower or equal to the P75 value and only 10% of the trucks have a use of cruise control that is greater than the P90 value. These values tell us something about how skewed the distribution is. If we look at drivers' data, Table 31 shows that 75% of them use cruise control less than 7,2% of driving time while 10% of them use cruise control more than 17,1%. Consequently, at some days the cruise control has been applied extensively while most of the days use of cruise control is limited.

Table 31 Descriptive statistics on relative amount of time using cruise control

	Pr vehicle	Pr driver
N	3308	2702
Mean	5,6	6,4%
Standard error	0,1913	0,2290
CV	3,4	3,6
Min (P0)	0,0	0,0%
P10	0,0	0,0%
P25	0,0	0,0%
Median (P50)	1,7	1,9%
P75	6,0	7,2%
P90	14,3	17,1%
Max	86,2	86,9%

The table shows no great differences between the distribution for vehicles and drivers. This suggests that vehicle attributes are no limiting factor for increased use of cruise control. Also, there seems to be a potential for more use of cruise control since 50% of vehicles and drivers use it less than 2% of total driving time. This can be seen from the median values in the table.

Figure 4 shows a histogram of use of cruise control as percentage of driving time per day per vehicle. The figure shows that the distribution is very skewed. At most days the use of cruise control is well below well 20%. For 90% of all vehicles, cruise control is used less than 15% of total driving time. All in all we might say that cruise control is not used extensively in vehicles.

Figure 4 Histogram use of cruise control as % of driving time per day pr vehicle

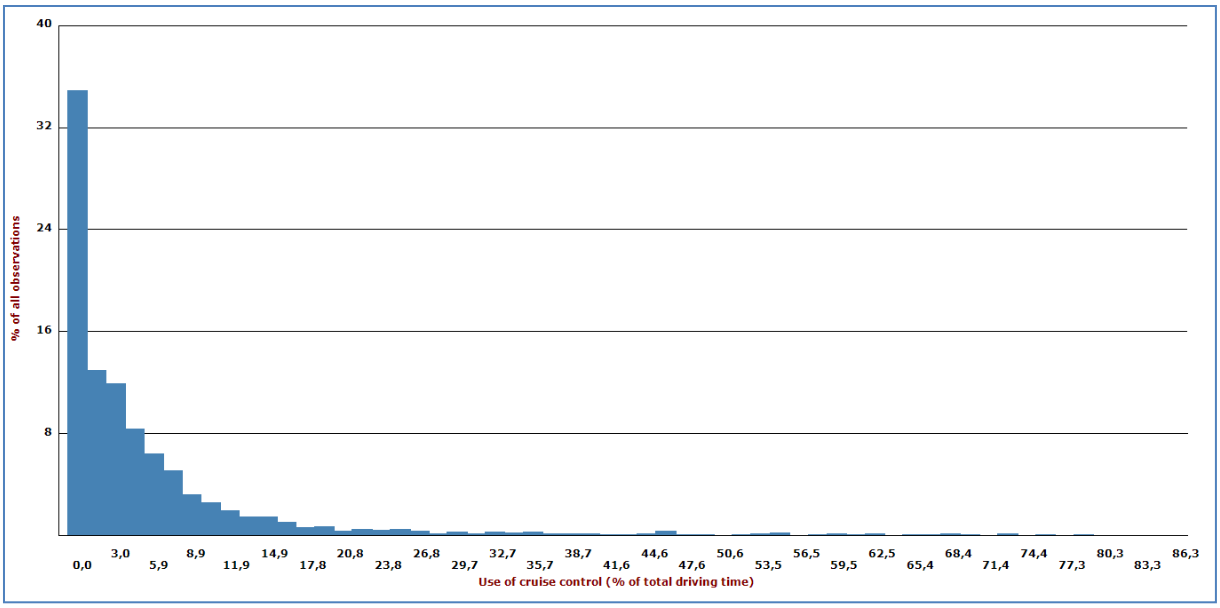

Figure 5 shows the same histogram for use of cruise control for drivers. As already mentioned, there are no obvious differences in use of cruise control between vehicles and drivers.

Figure 5 Histogram use of cruise control as % of driving time per day pr driver

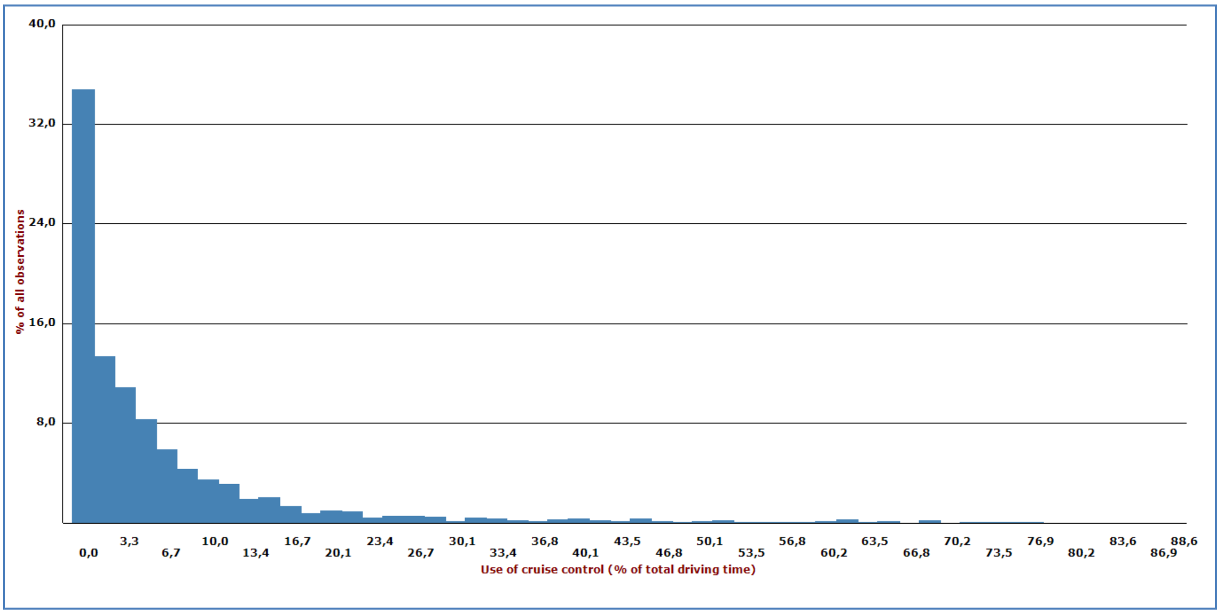

Figure 6 shows fuel consumption per 10 km by use of cruise control per vehicle. There is a great variation in use of cruise control. The figure shows a linear regression line fitted to the observations. Each truck has a letter code and this code is identified with the number of observation for that vehicle. This means that G_160 means the 160th observation for truck G. Each vehicle also has its own colour code. The colour code for each vehicle is shown in Figure 7. The colour identifier for each vehicle is the same in all plots.

Figure 6 Fuel consumption by use of cruise control per vehicle

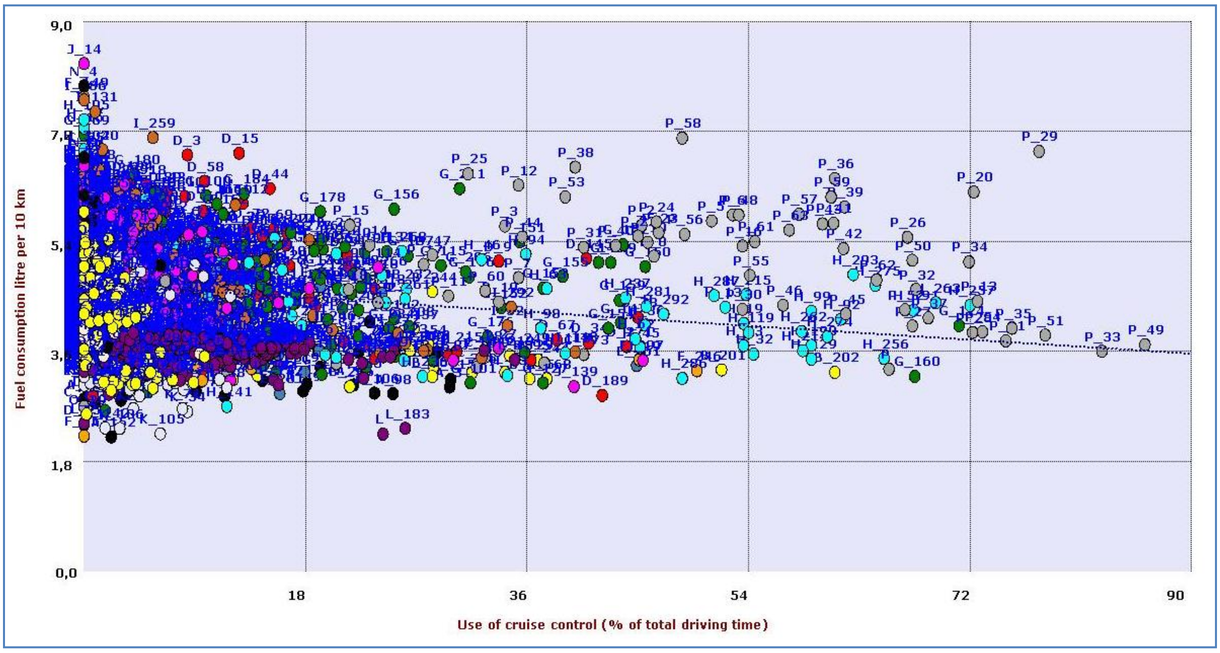

The colours in the plot suggests that vehicles H and P has more use of cruise control than other vehicles.

Figure 7 Colour legend for vehicles

Figure 8 Fuel consumption by use of cruise control per driver

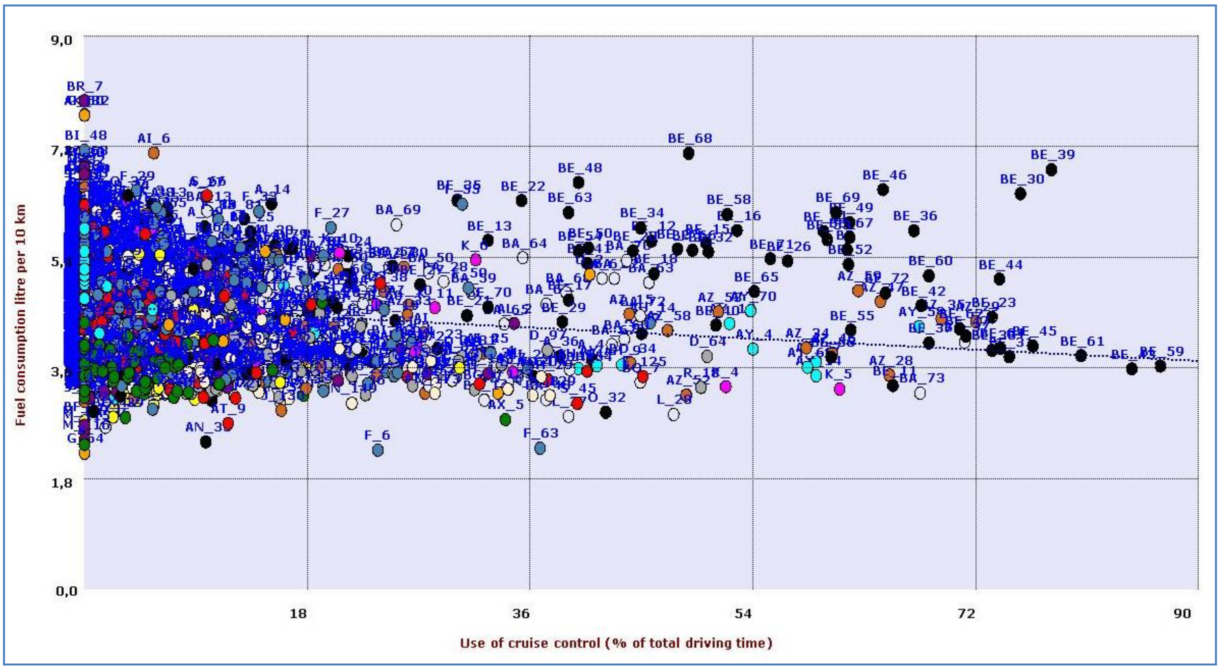

Figure 9 Colour legend for drivers

We will also present bivariate plots with data distributed on drivers and not on vehicles. Figure 8 shows the relationship between use of cruise control and fuel consumption for drivers. Again, there is no obvious difference between the relationship between use of cruise control and fuel consumption for vehicles and for drivers. In the bivariate plot for drivers, each driver is identified with a colour code, just as vehicles are. Figure 9 shows the colour code for drivers used in the plots.

When we assess the effect of cruise control on fuel consumption, we use data from the driver's table. We will do that for all assessments of regression coefficients from scatterplots if not otherwise stated in the text.

The linear effect of cruise control use is slightly negative, more use of cruise control seems to lower fuel consumption. The effect from the bivariate regression analysis is small but significant and the regression coefficient suggests that if the use of cruise control is increased by 10 percentage points (of total driving time) the expected average reduction in fuel consumption is 0,10 litre pr 10 km.   If a truck drives 100 000 km a year, the expected fuel savings by increasing the amount of driving time using cruise control by 10% is 1 000 litre or about 2,7 tonnes of CO_2 since there is 2,6628 kg CO_2 [16] for each litre diesel. All in all, use of cruise control explains about 2 percent of total variations in fuel consumption.

The effect of driving behaviour course focusing on fuel consumption

Some of the drivers participated in a driving behaviour course 17th and 18th of June 2011. The purpose of the course was to encourage fuel savings by focusing on driving behaviour. Use of cruise control was one of the indicators for driving behaviour that the course focused on. The question is then whether there is any change in use of cruise control before and after the course for drivers who took part in it.

Table 32 shows the result. We have performed a t-test for unpaired samples with assumed unequal variance in use of cruise control between the drivers [17]. The first sample is use of cruise control for the drivers before participating in the course while the second sample is the same use after the course took place. For one driver the use of cruise control was lower after participating in the course, but this effect is not statistically significant. For three drivers we find a statistically significant result, these drivers are marked with an asterisk in the table. For one driver the use of cruise control increased with nearly 15 percentage points. Before the course, this driver's fuel consumption was on average 4,9 litre per 10 km per day. After the course his fuel consumption is 4,5 litre per 10 km. This effect cannot alone be attributed to use of cruise control, but is is a sign that modifying driving behaviour, among them use of cruise control, may have a positive impact on fuel consumption. If the driver drives 100 000 km a year his yearly fuel consumption will be reduced by 4000 litre and the reduction in CO_2 emission would be 10,6 tonnes.

Table 32 Use of cruise control for drivers participating in driving behaviour course

Driver	Before	After	Difference	t-value	Degrees-of freedom	p-value	Signi-ficance
AK	2,9	4,1	1,2	0,86	71	0,3912	

[16] Toutain, J.E.W, Taarneby, G., Selvig, E.,*Energiforbruk og utslipp til luft fra innenlandsk transport*, Statistisk Sentralbyrå, Rapport 2998/49, Table 2.39 and Table 2.1. http://www.ssb.no/emner/01/03/10/rapp_200849/rapp_200849.pdf

[17] Driver K is not included in the analysis even though the driver participated in the course. There are only two registrations for this driver before the course and nine registrations after the course. We consider this to be too few registrations to give a meaningful comparison.

AY	9,7	5,9	-3,7	-0,89	46	0,3787	
BA	11,0	25,6	14,6	4,28	69	0,0001	*
E	2,0	3,5	1,5	1,92	105	0,0572	
L	0,3	5,9	5,6	2,69	33	0,0112	*
M	0,6	2,6	2,0	3,30	90	0,0014	*

Table 33 Effect on fuel consumption in litre per 10 km of driving behaviour course

	Before	After	Difference	t-value	Degrees-of freedom	p-value	Signi-ficance
AK	5,0	4,8	-0,1	-0,91	87	0,3635	
AY	4,7	5,1	0,4	2,09	52	0,0415	*
BA	4,9	4,5	-0,4	-2,31	70	0,0239	*
E	4,9	4,6	-0,3	-2,05	99	0,0432	*
L	4,3	4,2	-0,1	-0,54	18	0,5980	
M	4,8	4,3	-0,4	-3,49	162	0,0006	*

Table 33 shows the effect on fuel consumption for the drivers who participated in the course. As the table shows, four of six drivers have significantly different fuel consumption after they participated in the course. For one of the drivers the effect is positive which means that the fuel consumption after the course is *higher* than what it was before. Therefore, for three of six drivers we can conclude that fuel consumption is significantly *lower* after they participated in the driving behaviour course. For these drivers the reduction in fuel consumption is about 0,3-0,4 litre per 10 km which means a total reduction over one year of 3000-4000 litre assuming a total driving length of 100 000 km. This also implies a reduction in emissions of CO_2 in the order of 8-10,7 tonnes a year.

Automatic gear shift

Table 34 shows descriptive statistics for the relative amount of driving time the vehicles spent using automatic gear shifts. On average, about 90% of the driving time was spent using automatic gear shifts. The minimum relative amount of driving time was 5,2% while the maximum was 100% for vehicles. The percentiles (P-values) suggest that the distribution is skewed towards the right. Some days use of automatic gear shifts is low while most of the days the use is extensive. For both vehicles and drivers, there is a 50% chance that automatic gear shift is used more than roughly 98% of total driving time (the median).

Table 34 Descriptive statistics on relative amount of time using automatic gear shifts

	Pr vehicle	Pr driver
N	2838	2566
Mean	93,5%	94,4%
Standard error	0,1939	0,1700
CV	0,2	0,2
Min (P0)	5,2%	35,5%
P10	81,1%	83,1%

P25	91,1%	92,1%
Median (P50)	97,9%	98,4%
P75	99,8%	99,9%
P90	100,0%	100,0%
Max	100,0%	100,0%

Figure 10 Histogram use of automatic gear shift pr vehicle

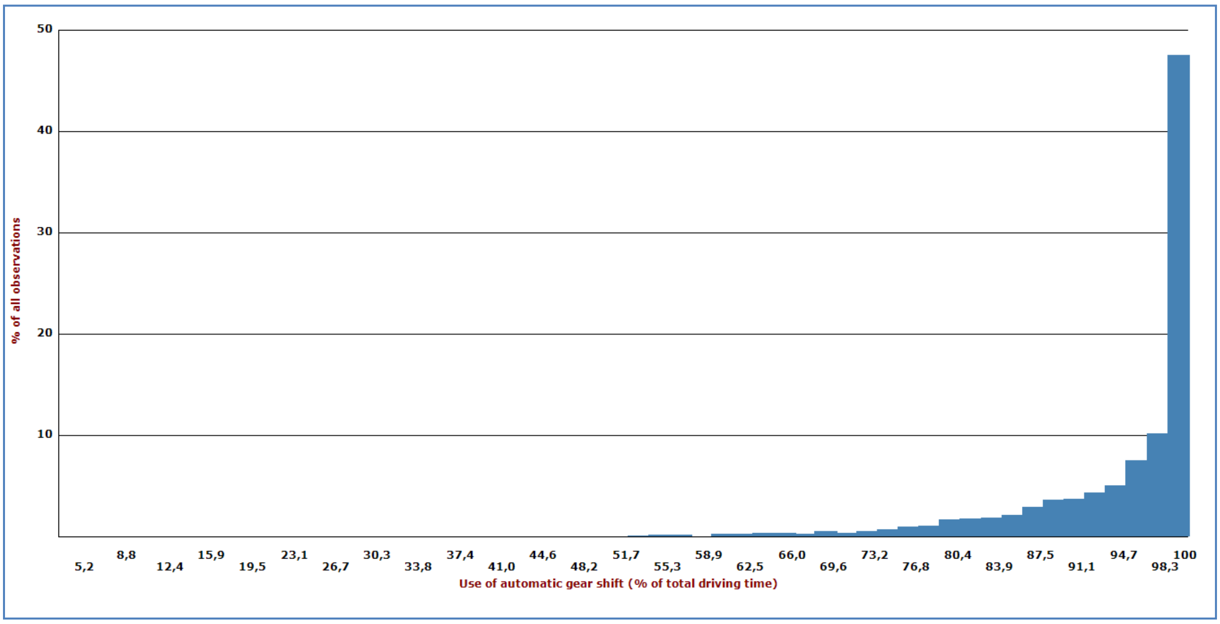

Figure 11 Histogram use of automatic gear shift pr driver

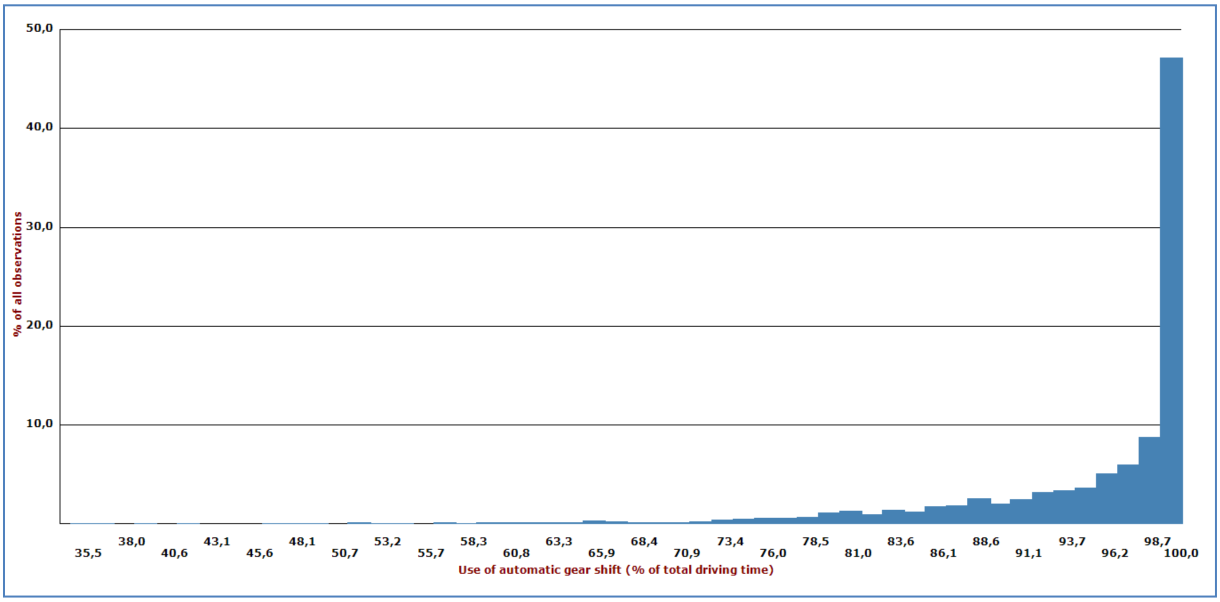

Figure 10 shows a histogram for use of automatic gear shift per vehicle. The figure shows that 50% of all vehicles used automatic gear shift more than 95% of driving time. Figure 11 shows the same histogram for use of automatic gear shift distributed for drivers. There are no obvious discrepancy between the two histograms. Table 34 shows that the minimum use of automatic gear shift is quite higher for drivers than for vehicles, probably reflecting the fact that not all vehicles are equally suitable for use of automatic gear shift under equal driving conditions.

Figure 12 Fuel consumption by use of automatic gear shift by vehicle

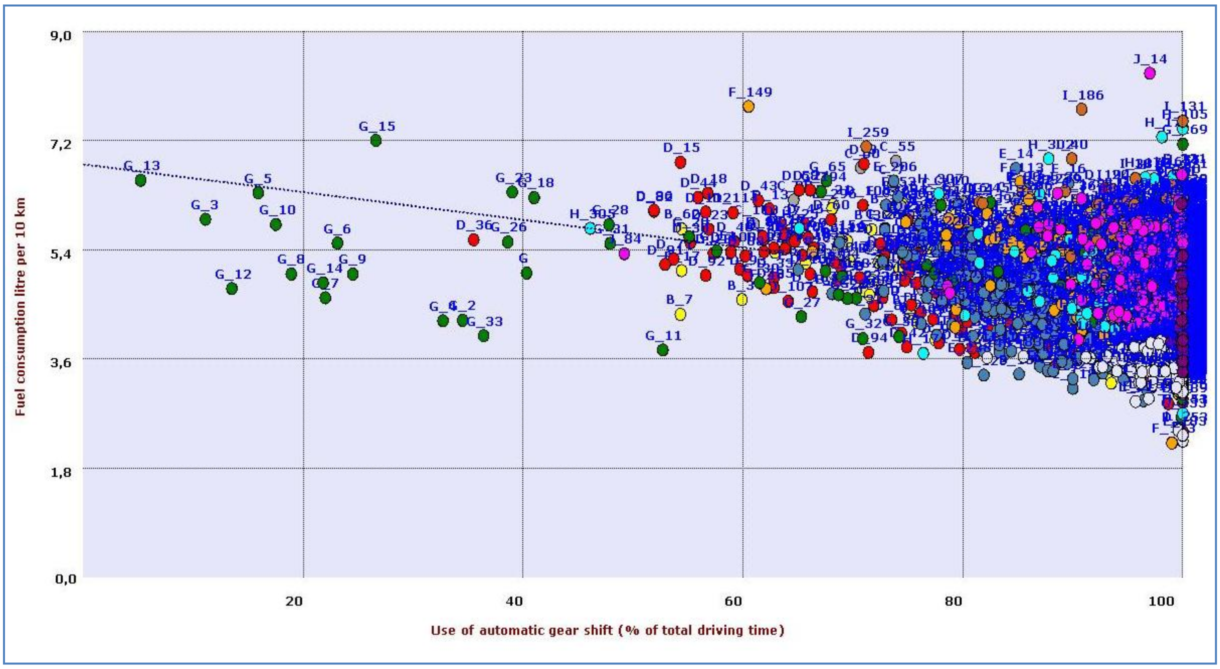

Figure 10 shows a histogram for use of automatic gear shift per vehicle. The figure shows that 50% of all vehicles used automatic gear shift more than 95% of driving time. Figure 11 shows the same histogram for use of automatic gear shift distributed for drivers. There are no obvious discrepancy between the two histograms. Table 34 shows that the minimum use of automatic gear shift is quite higher for drivers than for vehicles, probably reflecting the fact that not all vehicles are equally suitable for use of automatic gear shift under equal driving conditions.

Figure 12 shows the relationship between fuel consumption and use of automatic gear shift. The estimated linear regression line in the figure suggests that fuel consumption decreases slightly with increased use of automatic gear shift. This picture is not unambiguous, the highest registrations of fuel consumptions are found among the highest registrations of automatic gear shift. We have a clustering of data observations to the right of the figure since most of the time automatic gear shifts is in use over 95% of driving time. For vehicles that use automatic gear shift more than 80% of driving time there is a considerable spread in fuel consumption. Also, the spread around the regression line increases with increasing values of X which indicates heteroskedasticity, a condition that makes the regression model inappropriate [18]. Figure 12 suggests that differences in fuel consumption must be explained by more than one independent variable.

Figure 13 shows the same relationship distributed on drivers and not vehicles. We see the same pattern as for vehicles. There are more observations and bigger variance in the the left part of the figure for vehicles than for drivers. This confirms what we stated above, all vehicles may not be equally suitable for use of automatic gear shift.

[18] If heteroskedasticity is present, the coefficients are unbiased but their variance is larger than optimal so the coefficients are not the best among all unbiased coefficients.

Figure 13 Fuel consumption by use of automatic gear shift by driver

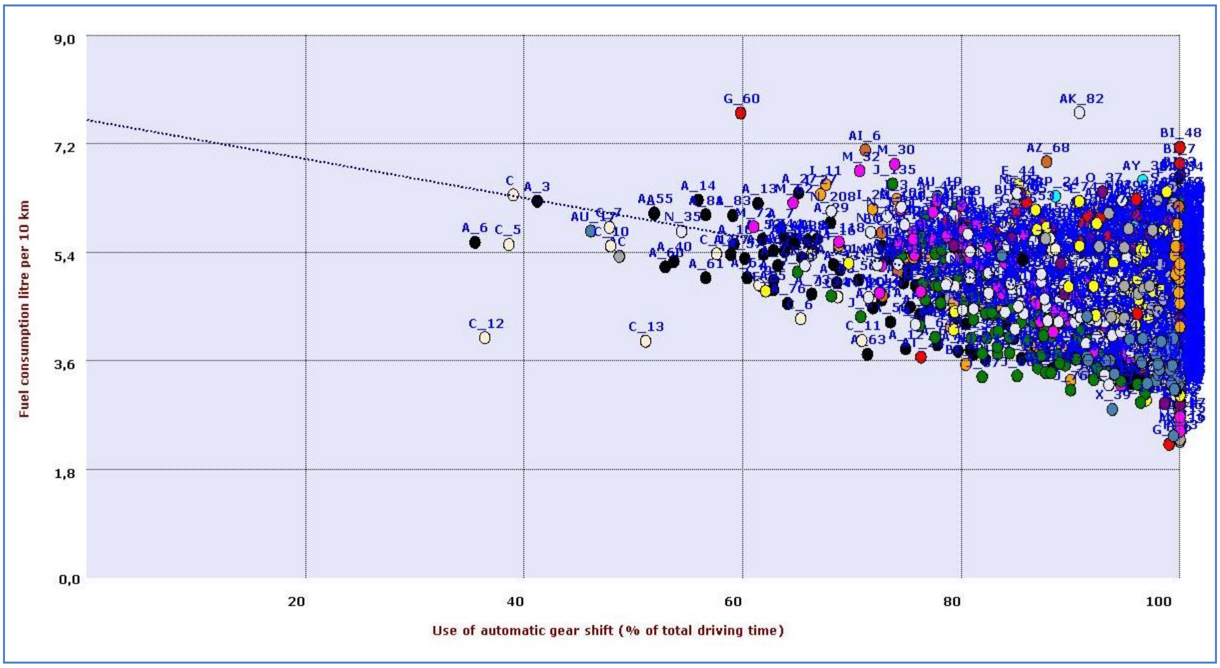

On average there seems to be an effect of increased use of automatic gear shift. According to the regression line, a 10 percentage points increase in time spent using automatic gear shifts can reduce the fuel consumption on average by 0,32 litre per 10 km. An increase of about 31 percentage points in use of automatic gear shift will reduce the fuel consumption by 1 litre per 10 km. To illustrate, assuming a driver drives 100 000 km a year, the effect on fuel consumption from a 10 percent increase in time spent using automatic gear shift would be 3200 litre per year. With 2,6628 kg CO2 pr litre diesel this would imply a reduction in emissions of CO_2 of almost 8,5 tonnes. The variation in use of automatic gear shifts explains about 12% of the variations in fuel consumption.

It should be noted that the regression effect described above is calculated for drivers. If we calculate the same effect for vehicles the effect is smaller, a 10 percentage points increase in use of automatic gear shifts will reduce the fuel consumption with 0,23 litre per 10 km. The model for vehicles explains about 8% of variations in fuel consumption. This is because the differences in fuel consumption between vehicles are greater than the corresponding differences between drivers.

We assume that more use of cruise control and automatic gear shifts will give a more even, steady speed and less variations in engine load. This is the assumed mechanism for the effects observed in the figures above.

Table 35 Use of automatic gear shifts for participants in driving behaviour course

	Before	After	Difference	t-value	Degrees-of-freedom	p-value	Signi-ficance	Fuel saving litre per 10 km
AK	98,8	98,4	-0,4	-0,95	88	0,345		-0,1
AY	98,5	99,8	1,3	1,84	27	0,076		0,4
BA	99,6	99,9	0,3	2,20	45	0,033	*	-0,4
E	90,4	96,5	6,1	4,84	72	0,000	*	-0,3
L	99,9	100,0	0,1	1,16	12	0,268		-0,1

M	88,5	92,9	4,4	3,53	170	0,001	*	-0,4

Table 35 shows use of automatic gear shift for participants in the driving behaviour course before and after the course was taken. For driver L there are only a few observations so the registrations for this driver are very uncertain. Three of six drivers have significantly increased their use of automatic gear shifts after the course. We have also included the numbers for fuel savings taken from Table 33 above. These numbers show that for the three drivers who have significantly increased their use of automatic gear shift after the course, the fuel savings are about 0,3-0,4 litre per 10 km. This is *not* due to increased use of automatic gear shifts alone, but the table suggests it may have some impact on fuel savings.

Average driving speed per day

Table 36 shows descriptive statistics for average driving speed pr day. The average driving speed was 62,7 km pr hour pr vehicle pr day and 62,9 km per driver in the registration period. The maximum average driving speed was 82,8 for both vehicles and drivers The minimum speed was 36,5 km per hour per day for vehicles and 42,2 km per hour per day for drivers. This indicates that driving speed varies somewhat less between drivers than between vehicles. The percentiles (P-values) suggest that 10% of vehicles had an average speed per day less than or equal to 55,8 km/hour while 10% had an average speed over 69,8 km/hour. The numbers for drivers are roughly the same. Half the vehicles and drivers had an average speed over 62,7 km/hour per day.

Appendix A documents some unreasonable speed registrations from the drivers' data table that were excluded from the analysis.

Table 36 Descriptive statistics on average driving speed in km per hour

	Pr vehicle	Pr driver
N	3308	2701
Mean	62,7	62,9
Standard error	0,0984	0,1272
CV	0,2	0,2
Min (P0)	36,5	42,2
P10	55,8	54,7
P25	59,3	59,0
Median (P50)	62,7	62,7
P75	66,3	66,8
P90	69,8	72,0
Max	82,8	82,8

Figure 14 Histogram average driving speed km per hour per vehicle per day

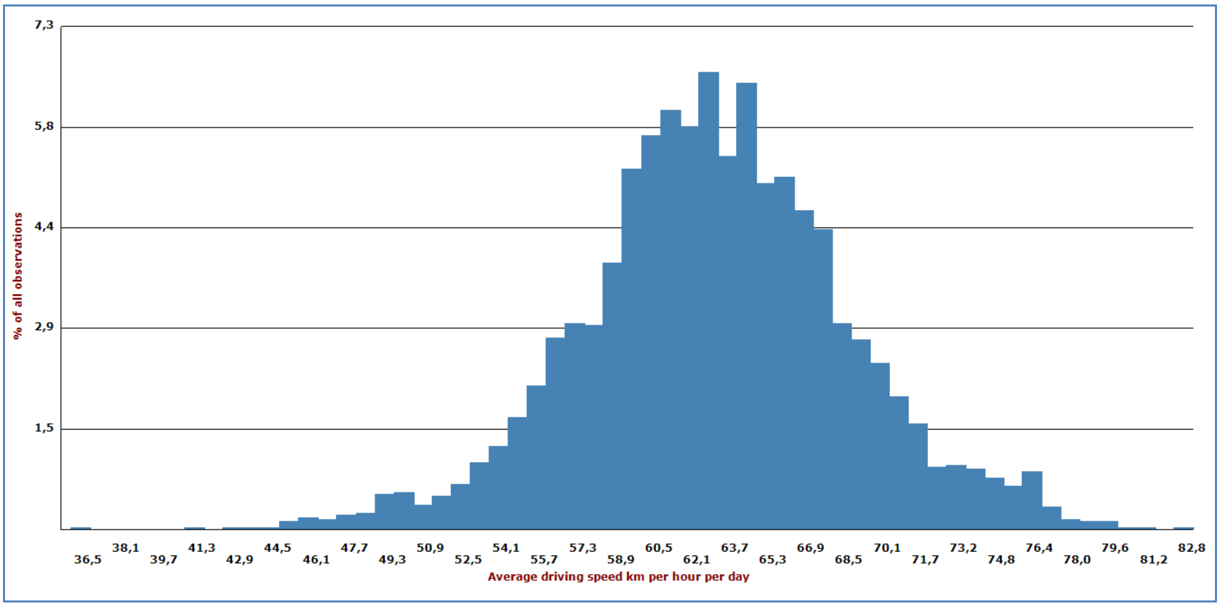

Figure 14 shows a histogram for average driving speed per vehicle per day. The figure shows that most vehicles at most days have a driving speed between roughly 50 and 75 km per hour. About 90% of all observations are in this interval. Figure 15 shows the same histogram for drivers. Both distributions are quite symmetric, accordingly there is an equal chance of driving slower or faster than the mean. Thus, drivers' intention may have a lesser impact than other indicators such as terrain, road quality and driving conditions.

Figure 15 Histogram average driving speed km per hour per driver per day

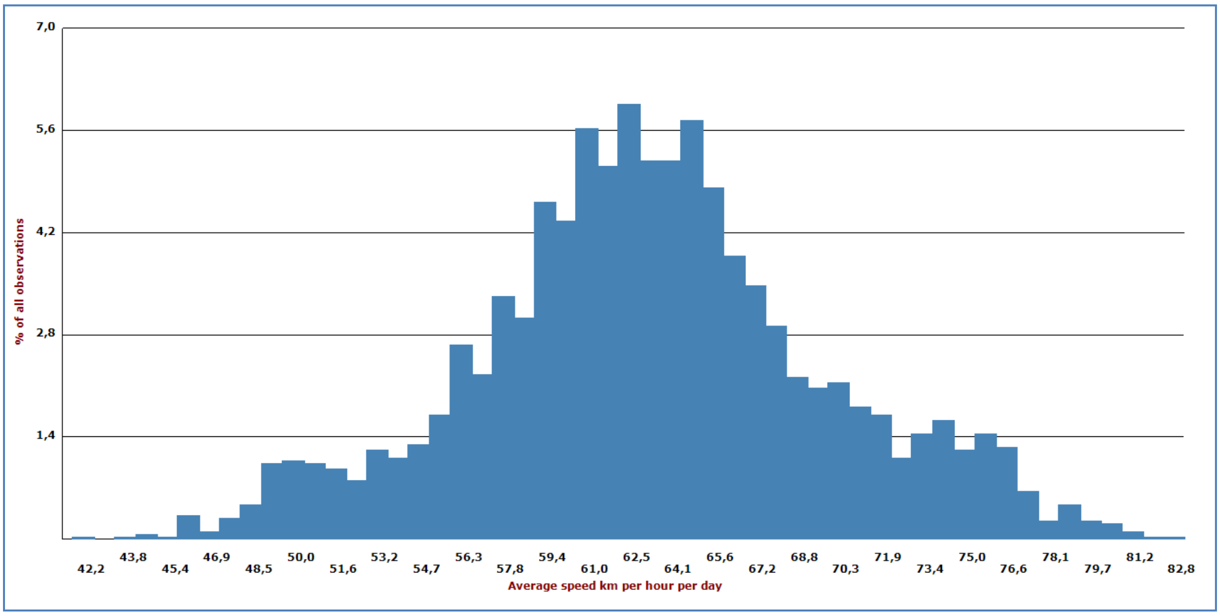

The average speed is not very high due to frequent climbing of steep hills on narrow, winding roads. Also, driving conditions in the winter months prohibit high average speed.

Figure 16 shows the relationship between average speed and fuel consumption. As vehicles climb steep hills, they need to go slower and use more fuel pr km. Driving on better roads with less curvature and less steep hill climbing lowers the fuel consumption and allows for higher average speed. This is presumably the effect captured in the figure. The regression line is an inverse regression line, it is non-linear and moves towards a threshold line. The more the X-value increases, the more we move towards this threshold line. The shape of the regression line suggests that we can reduce the fuel consumption by increasing the speed, but only up to a certain point. Above that point, there is no reduction in fuel consumption. Also, the effect of increasing average speed is larger when the average speed is low to start with.

Figure 16 Fuel consumption by average speed for vehicles

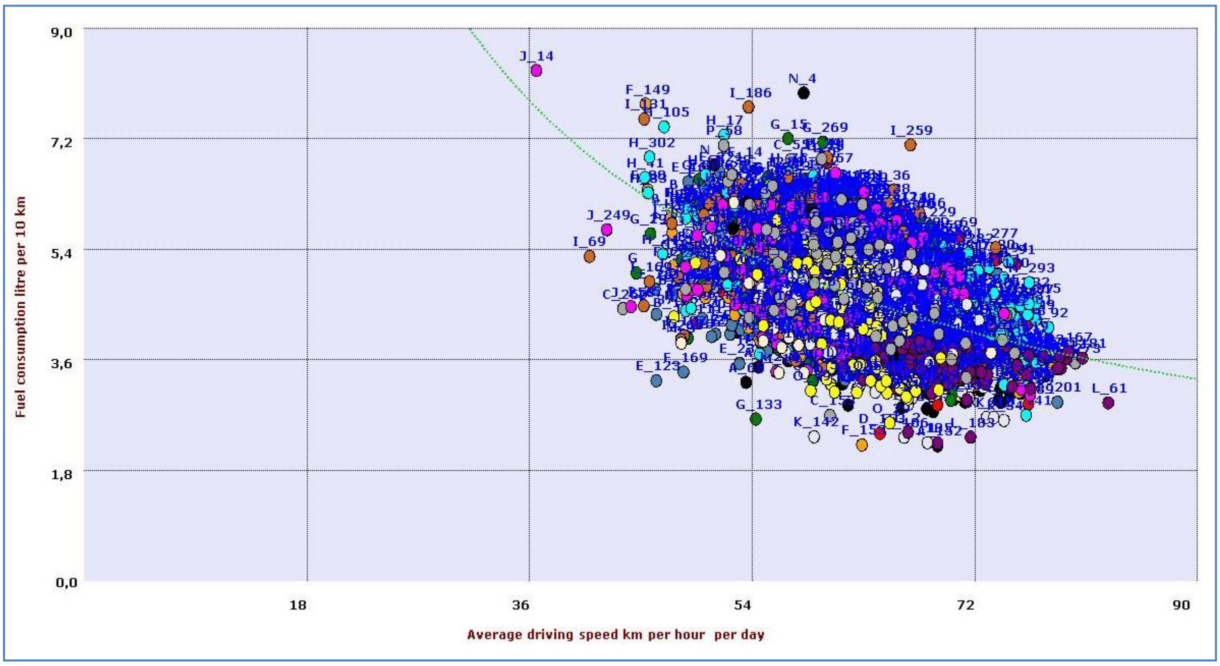

Figure 17 Fuel consumption by average speed for drivers

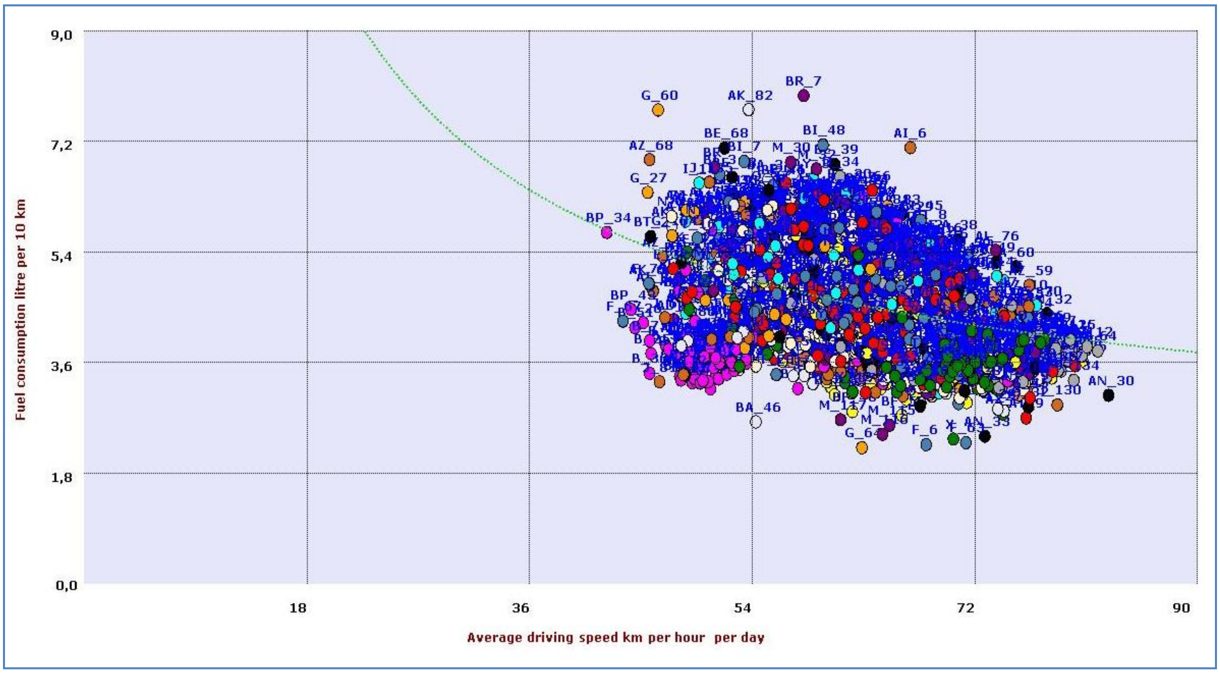

Figure 15 shows the same figure for drivers. We see the same trend, but for an average driving speed of 45-55 km per hour there seems to be slightly more variations in fuel consumption for drivers than for vehicles.

The effect of an increase in average speed can best be illustrated as in Table 37 which is calculated from drivers' data distribution. The table shows estimated fuel consumption with five different input values for average speed. We calculate the effect of fuel savings for four of the input values relative to the previous one. The table shows that an increase from 30 to 40 km gives a fuel savings of 1,3 litre pr 10 km while the savings from 40 to 50 km is 0,8 litre pr 10 km. The increase in average speed from 60 to 70 km an hour on the other hand will yield savings of 0,4 litre pr 10 km which is considerably less than going from 30 to 40 km an hour. Also, as the table shows, the higher the average speed to start with, the less the effect on fuel savings from an additional increase in speed. If average driving speed was increased from 70 to 80 km per hour the decrease in fuel consumption approaches zero.

Table 37 Estimated effect of speed increase on fuel consumption from regression model

Average speed in km pr hour	Estimated fuel consumption in litre pr 10 km	Savings relative to previous speed	Savings relative to 30 km an hour
30	7,3	0	0
40	6,0	-1,3	-1,3
50	5,2	-0,8	-2,1
60	4,7	-0,5	-2,6
70	4,3	-0,4	-3,0
80	4,0	-0,3	-3,3

The effects in Table 37 are from the bivariate analysis. As discussed above, it may well be that the direct, independent effect of average speed is lower when controlled for other independent variables that are also influenced by average speed in a multivariate regression model.

A driver that drives 100 000 km a year could save 13 000 litre of fuel if average speed could be increased from 30 to 40 km an hour. With an average emission of 2,6628 kg CO2 pr litre diesel [19] this alone would imply a reduction in emissions of CO_2 of more than 34,6 tonnes only for one vehicle. Variations in average speed alone explains 11% of all variations in fuel consumption per 10 km per day. An increase in driving speed is of course more a question about infrastructure in a difficult terrain than about drivers' behaviour.

Table 38 Average driving speed for participants in driving behaviour course

	Before	After	Difference	t-value	Degrees-	p	Signi-

[19] Toutain, J.E.W, Taarneby, G., Selvig, E.,*Energiforbruk og utslipp til luft fra innenlandsk transport*, Statistisk Sentralbyrå, Rapport 2998/49, Table 2.39 and Table 2.1. http://www.ssb.no/emner/01/03/10/rapp_200849/rapp_200849.pdf

					of-freedom		ficance
AK	60,0	61,6	1,6	1,46	73	0,1489	
AY	64,5	61,8	-2,7	-1,55	59	0,1253	
BA	63,5	62,8	-0,7	-0,63	71	0,5315	
E	64,1	63,0	-1,1	-1,37	120	0,1729	
L	59,6	61,8	2,1	0,91	22	0,3738	
M	64,0	63,6	-0,4	-0,57	164	0,5689	

Table 38 shows the average driving speed for participants in the driving behaviour course in June 2011. There is no significant change in driving speed for none of the participants. This result is hardly surprising since it is terrain, roads and driving conditions that limits the driving speed more than drivers' intentions.

Running idle

Table 39 shows descriptive statistics for the relative amount of time the vehicles are running idle. The statistics are distributed on vehicles and drivers. Preferably, the amount of time spent running idle should be as small as possible since the vehicle is not doing any work. On the other hand, stopping and starting the engine may cause higher fuel consumption during cold days than simply letting the engine run during stops.

On average, the vehicles spent about 14% of their driving time running idle wile the corresponding mean for drivers is about 13%. The maximum amount of time spent running idle was 66% for vehicles and a little less, 61%, for drivers. The lowest 10% of vehicles and drivers spent about 4% of total driving time running idle. Half of all vehicles spent more than roughly 12% of total driving time running idle, the number for drivers is about one percentage point lower. The 10% of vehicles with the highest values spent more than roughly 27% of total driving time running idle, the same number for drivers was about two percentage points lower. All in all, this suggests that vehicle attributes have a greater impact on the propensity to let the vehicle run idle than drivers' intention.

Table 39 Descriptive statistics on running idle

	Pr vehicle	Pr driver
N	3308	2702
Mean	14,0%	12,7%
Standard error	0,1617	0,1703
CV	1,2	1,3
Min (P0)	1,0%	0,8%
P10	4,4%	4,0%
P25	7,0%	6,3%
Median (P50)	11,7%	10,5%
P75	18,5%	16,9%
P90	26,7%	24,5%
Max	66,0%	60,6%

Figure 18 shows a histogram for the distribution of percentage of time spent running idle for vehicles. The distribution is skewed to the right, some few vehicles at some days spent a very high percentage of their driving time running idle while most other vehicles at most other days spent roughly 5 to 40 percent of their driving time. We assume that cold weather in winter is the most influencing factor for the relative amount of driving time spent running idle.

Figure 18 shows the same histogram for drivers. The shape of the distribution is identical, but the histogram for vehicles shows some larger values on the X-axis which confirms the proposition above: The relative amount of time spent running idle vary more between vehicles than between drivers.

Figure 18 Histogram running idle pr vehicle

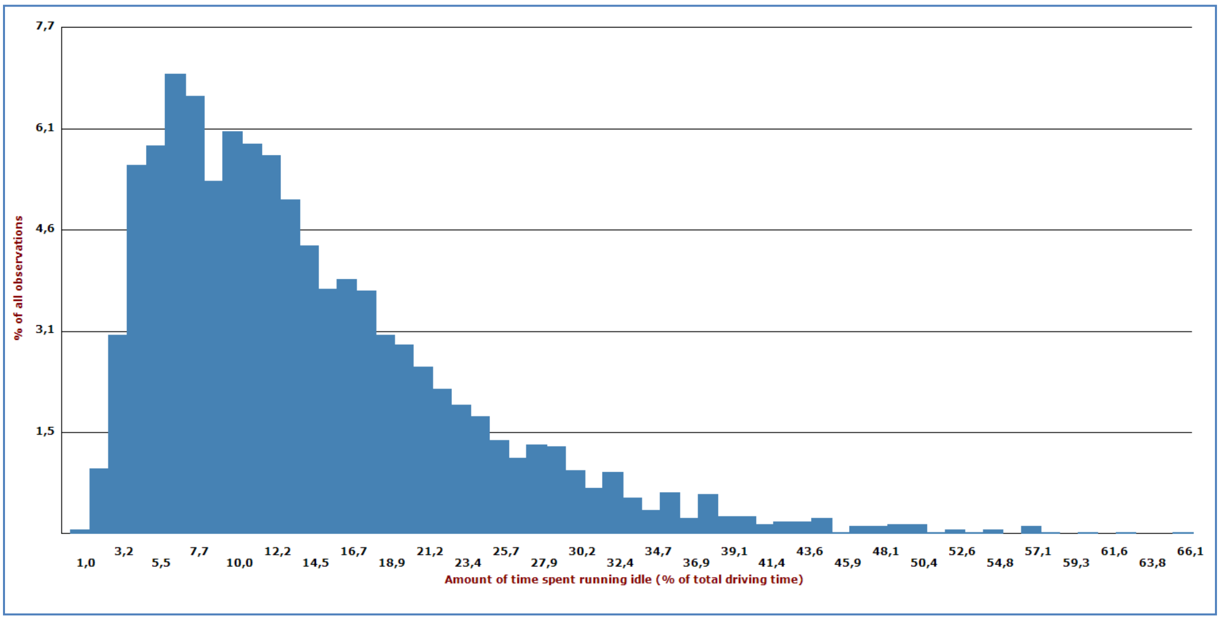

Figure 19 Histogram running idle pr driver

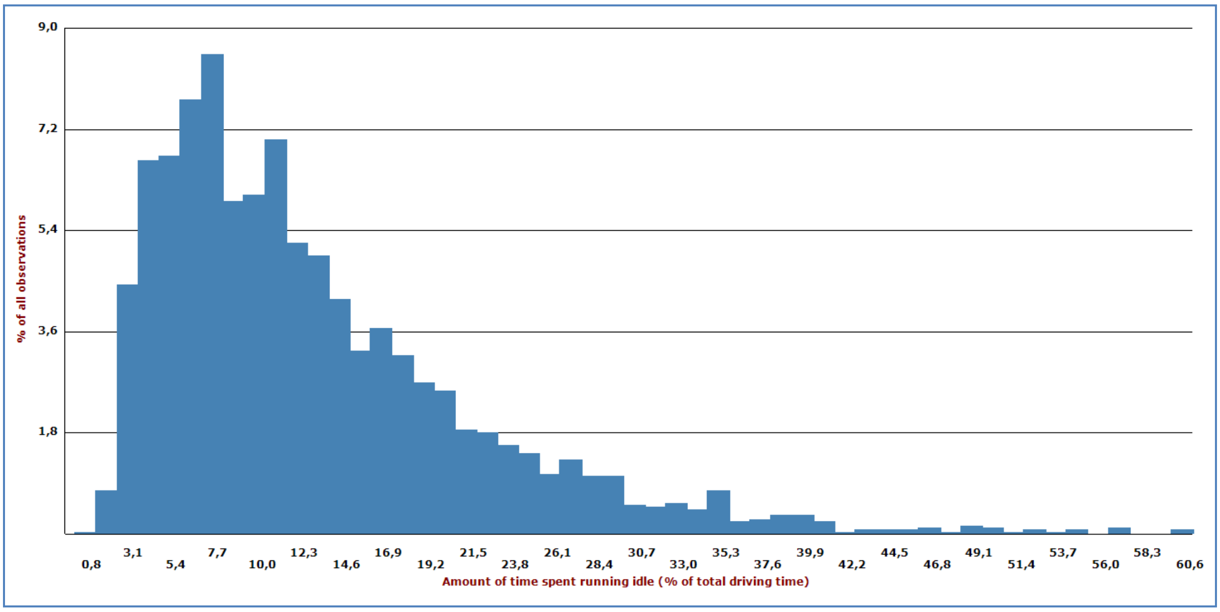

Figure 20 shows fuel consumption vs percentage of total driving time spent running idle pr vehicle. Figure 21 shows the same relationship for drivers. There are no major differences between the two plots.

The effect of an increased relative amount of time spent running idle is positive, statistically significant but weak in the figure based on drivers' data. An increase of 10 percentage points will increase fuel consumption by 0,09 litre pr 10 km. A truck running 100 000 km a year would save 900 litre fuel or about 2,4 tonnes CO_2 by *decreasing* the amount of driving time running idle with 10 percentage points. The variations in relative amount of time spent running idle explains only 1% of total variations in fuel consumption for drivers.

The linear regression effect for vehicles is twice the size of the effect for drivers. Based on vehicle data, if the amount of driving time spent running idle was reduced with 10 percentage point the fuel consumption would be 0,18 litre lower per 10 km. The amount of time spent running idle explains about 4% of total variations in fuel consumption in the data material from vehicles. This again suggests that some vehicles are harder to start or to heat up and that this determines more of the time spent running idle than drivers' intentions.

Figure 20 Fuel consumption by percentage of time running idle for vehicles

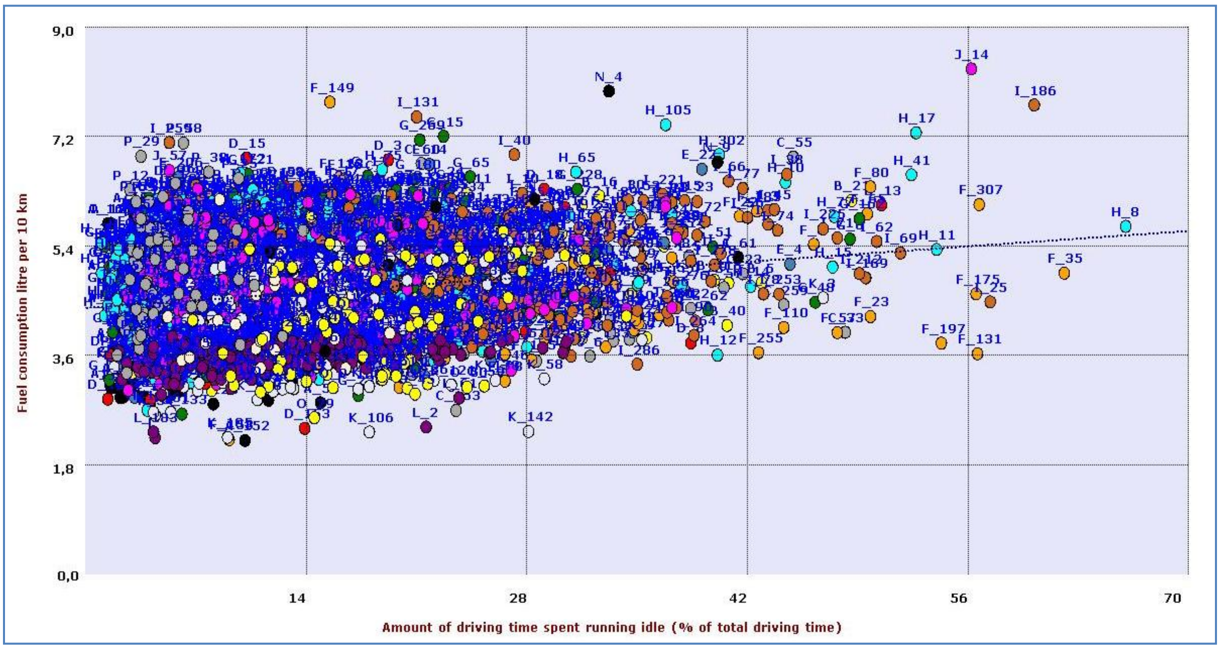

Figure 21 Fuel consumption by percentage of time running idle for drivers

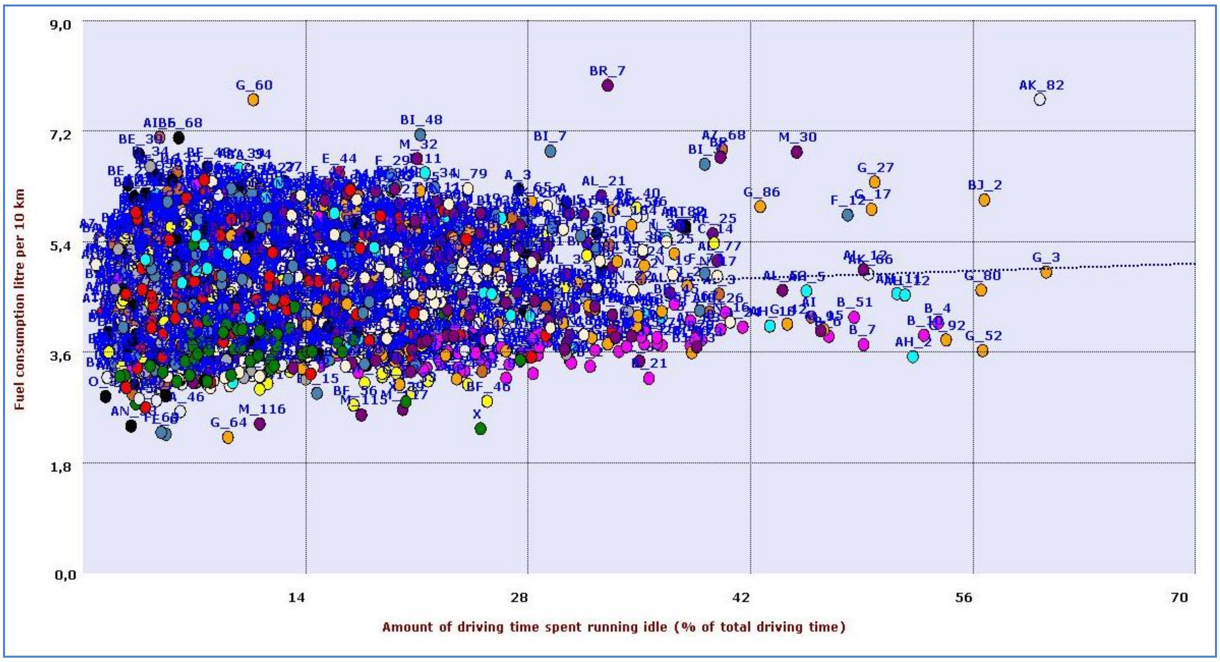

Table 40 shows the amount of time spent running idle for participants in the driving behaviour course. The table shows that five out of six participants have reduced the amount of time spent running idle. For four of them the reduction is statistically significant, which means it is larger than what could be expected from pure chance. The reduction is is in the range of 7,1% to 2,9% for the significant effects. We should mention that we have included winter months both before and after the course was taken.

Table 40 Amount of time spent running idle for participants in driving behaviour course

Driver	Before	After	Difference	t-value	Degrees-of-freedom	p-value	Signi-ficance	Fuel saving
AK	17,5	18,0	0,5	0,27	88	0,7847		-0,1
AY	10,7	7,8	-2,9	-2,49	45	0,0167	*	0,4
BA	9,6	5,5	-4,1	-4,68	71	0,0000	*	-0,4
E	14,3	7,8	-6,5	-7,27	85	0,0000	*	-0,3
L	10,9	9,2	-1,7	-1,82	33	0,0781		-0,1
M	17,9	10,8	-7,1	-6,67	169	0,0000	*	-0,4

Engine load of more than 90% of maximum torque

Table 41 shows descriptive statistics for the relative amount of driving time the engine load is above 90% of maximum torque [20] pr vehicle. On average, 12,5% of driving time was spent driving the vehicles with this engine load. The mean value is about the same for vehicles and drivers. Again, it

[20] Six observations for vehicle K are omitted from the analysis because of extreme values. They all spent 82% percent or more of total driving time with 90% of maximum torque. This is more than twice the maximum value when these observations are discarded. All observations were made within a six day period from 8th to 14th of March, 2011. These observations do not have extreme values on other variables such as driving distance and fuel consumption.

should be noted that the vehicles spend a lot of their time in terrain dominated by steep hill climbing on narrow, winding roads between western and eastern Norway.

Table 41 Descriptive statistics on the relative amount of time the engine load is above 90% of maximum torque

	Pr vehicle	Pr driver
N	3302	2702
Mean	12,5%	12,4%
Standard error	0,1278	0,1472
CV	1,0	1,2
Min (P0)	0,0%	0,0%
P10	3,3%	2,7%
P25	7,0%	6,6%
Median (P50)	11,8%	11,6%
P75	17,8%	18,0%
P90	22,7%	22,9%
Max	40,6%	41,4%

The distribution pr vehicle and pr driver are almost identical. Thus, drivers' intention may be just as decisive as other factors influencing fuel consumption such as terrain, road curvature and quality and driving conditions. To test for that , we need to control for other factors that may be influenced by use of maximum torque and which also has an influence on fuel consumption.

Figure 22 Histogram amount of driving time in % with engine load more than 90% of maximum torque pr vehicle

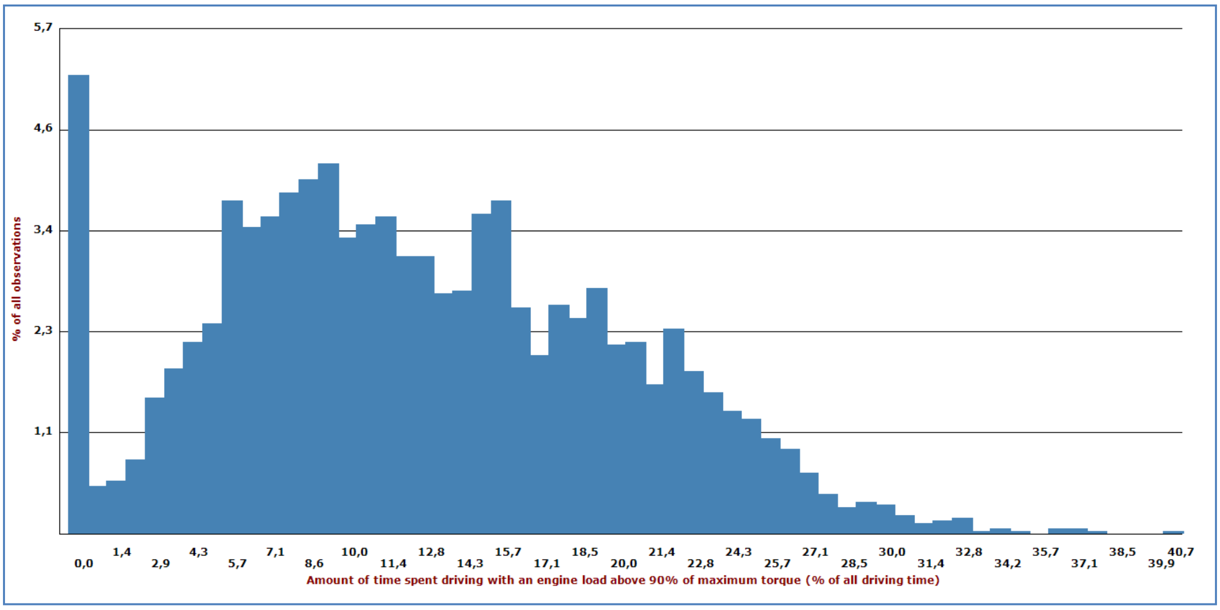

Figure 23 Histogram amount of driving time in % with engine load more than 90% of maximum torque pr driver

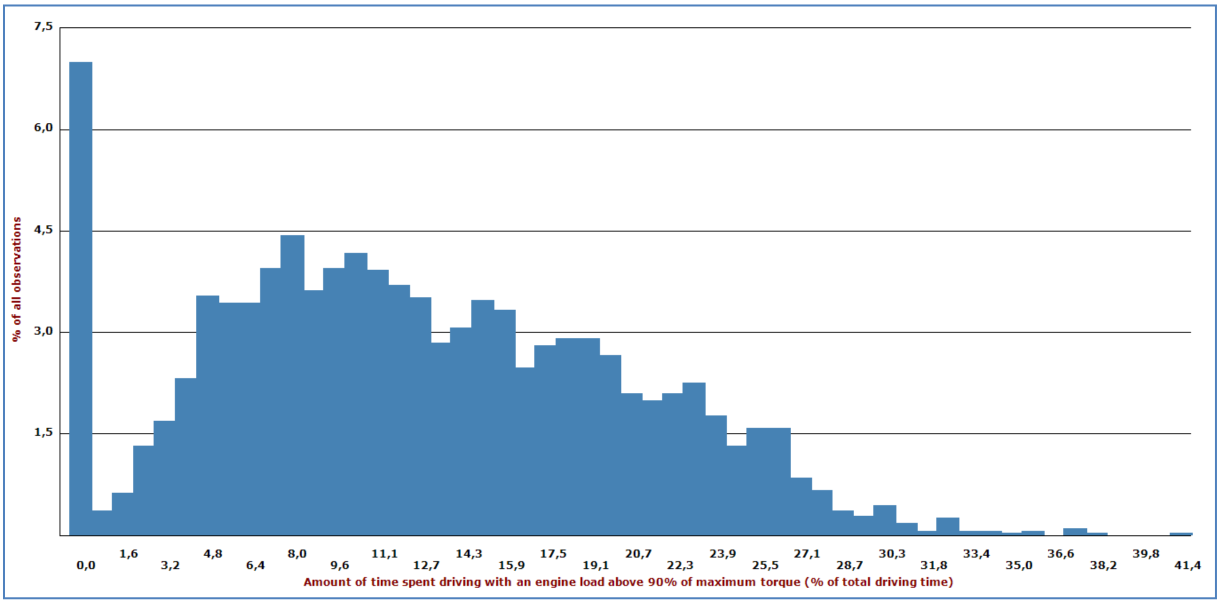

Figure 22Figure 23 shows a histogram for percent of driving time in percent spent with an engine load above 90% of maximum torque per vehicle. The distribution is right skewed, some few vehicles spend a larger percentage of the driving time with a high engine load. Most vehicles (80%) are in the interval from 4 to 25 percent of the driving time. Figure 23 shows the same histogram for drivers. The differences between the two histograms are negligible. It is noteworthy that some vehicles and drivers do not use more than 90% of maximum torque at all on some days. This lead us the the proposition that the routes the vehicles are travelling may be an important factor.

Figure 24 Fuel consumption by percentage of time spent with an engine load over 90% of maximum torque pr vehicle

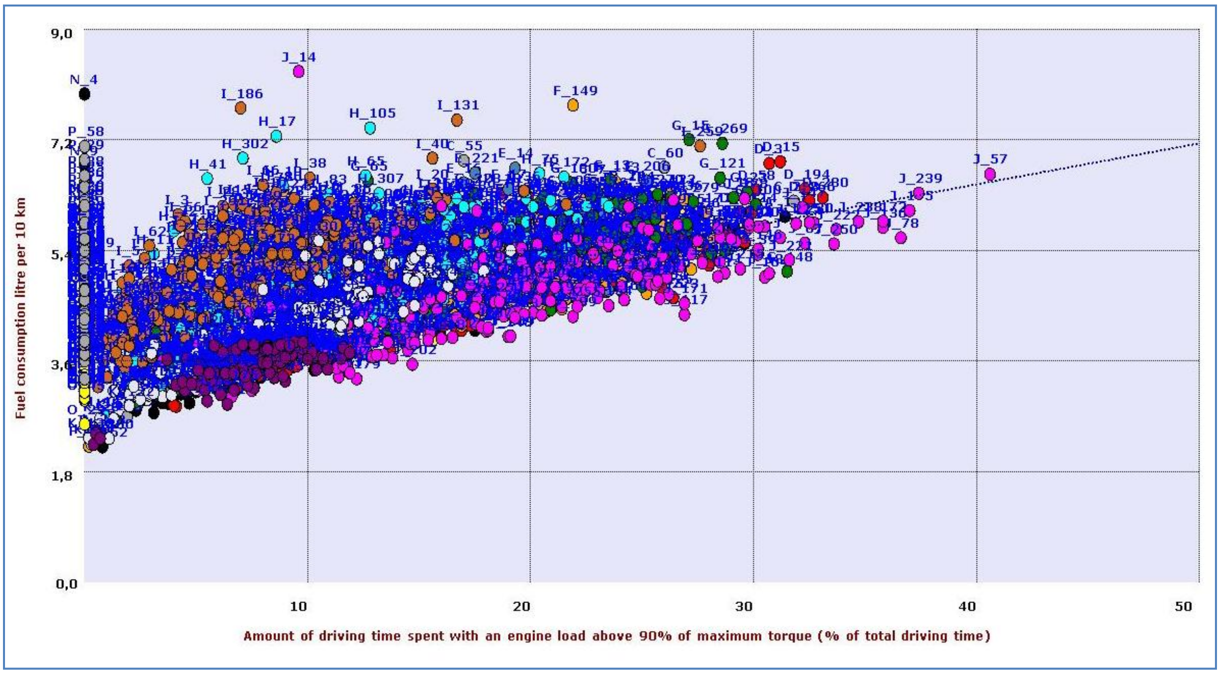

Figure 25 Fuel consumption by percentage of time spent with an engine load over 90% of maximum torque pr driver

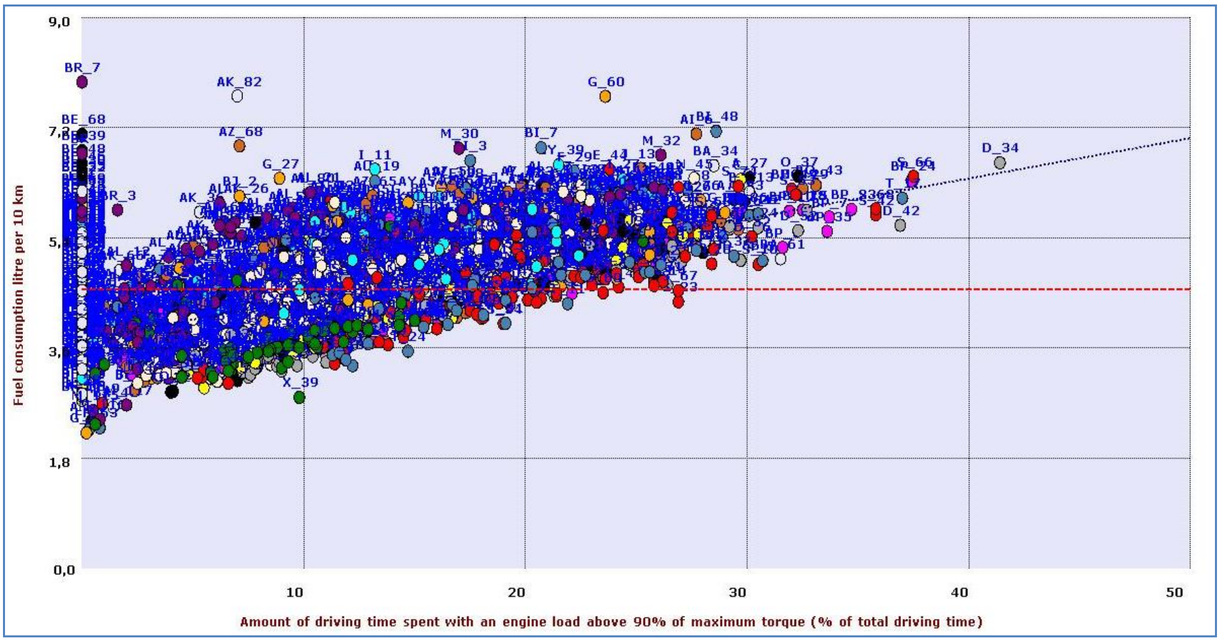

Figure 24 shows the relationship between fuel consumption and the relative amount of time the engine load of the vehicle is above 90% of maximum torque pr vehicle. Figure 25 shows the same relationship for drivers. Again, the two figures are practically identical. There is, not surprising, a significant relationship between engine load and fuel consumption. Using drivers' data, an increase of 10 percentage point in driving time with this engine load will on average increase the fuel consumption by 0,66 litre per 10 km. Given that the mean time spent with this engine load is less than 13%, an increase of 10 percentage points is a lot. Still, if a vehicle could decrease its time spent driving with this engine load from 20% to 10% on average, a yearly driving distance of 100 000 km would imply fuel savings of about 6600 litre and a reduction in the CO_2-emissions of about 17,6 tonnes. Such a decrease in engine load is only possible with improved infrastructure. All in all, variations in use of engine load above 90% of maximum torque alone explains 37% of variations in fuel consumption using drivers' data. Therefore, engine load is among the factors with biggest impact on fuel consumption.

Table 42 Amount of time spent driving with an engine load above 90% of maximum torque for participants in driving behaviour course

Driver	Before	After	Diff	t-value	Degrees-of-freedom	p-value	Signifi-cance	Fuel savings
AK	7,5	7,4	0,0	-0,03	73	0,9735		-0,1
AY	11,3	14,2	2,9	2,67	58	0,0099	*	0,4
BA	18,1	16,6	-1,4	-0,85	68	0,3984		-0,4
E	14,4	16,9	2,5	2,15	112	0,0338	*	-0,3
M	13,4	13,2	-0,3	-0,24	164	0,8127		-0,4

Table 42 shows the relative amount of driving time spent driving with an engine load above 90% of maximum torque for participants in the driving behaviour course summer 2011. Driver L spent no time using more than 90% of maximum torque neither before nor after the course. We assume this is because the driver only drove a specific route where torque demand is lower. This driver is not included in the analysis.

Two drivers have significantly different values before and after the course, but in both cases they spent *more* time with an engine load above 90% of maximum torque after the course than before. This shows that other factors than drivers' intentions are decisive for the amount of driving time spent with a high engine load.

Highest gear

Table 43 shows descriptive statistics on relative amount of driving time per day spent driving in highest gear pr vehicle. On average, almost 50% of the driving time was spent driving in the highest gear for vehicles while the average for drivers was slightly larger. There is a considerable variation among drivers since the minimum is 5,5% while the maximum is above 90% of total driving time. The vehicles with the lowest 10% of the values spend at most 32% of driving time in highest gear and the vehicles with the highest 10% of values spend at least 67% in highest gear [21]. The mean and median is almost identical for vehicles which suggests a symmetric distribution which confirmed by the histogram in Figure 26 for vehicles.

Table 43 Descriptive statistics on relative amount of driving time spent driving in highest gear

	Pr vehicle	Pr driver
N	3308	2702
Mean	49,5%	51,4%
Standard error	0,2343	0,2726
CV	0,5	0,5
Min (P0)	5,5%	5,5%
P10	32,3%	33,7%
P25	40,4%	41,5%
Median (P50)	49,6%	50,7%
P75	58,7%	61,1%
P90	66,6%	71,5%
Max	86,9%	90,1%

Table 43 also shows data pr driver and not pr vehicle. It seems that the distribution pr driver is a bit different from vehicles in the sense that the drivers' distribution has a higher use of the highest gear for the same percentile. For example, the lowest 10% of drivers use the highest gear in at most 34% of total driving time while the same value for vehicles is roughly 32%. Also, the highest 10% of values from the driver table use the highest gear at least 71,5% of the driving time while the corresponding number from the vehicle table is 66,6%. This may be an indicator that drivers' intention more than vehicle attributes determine the choice of highest gear. Variation among drivers in use of the highest gear is larger than what could be expected from vehicles alone.

[21] These are roughly P10 and P90 values

Figure 26 Histogram of amount of time spent using highest gear pr vehicle

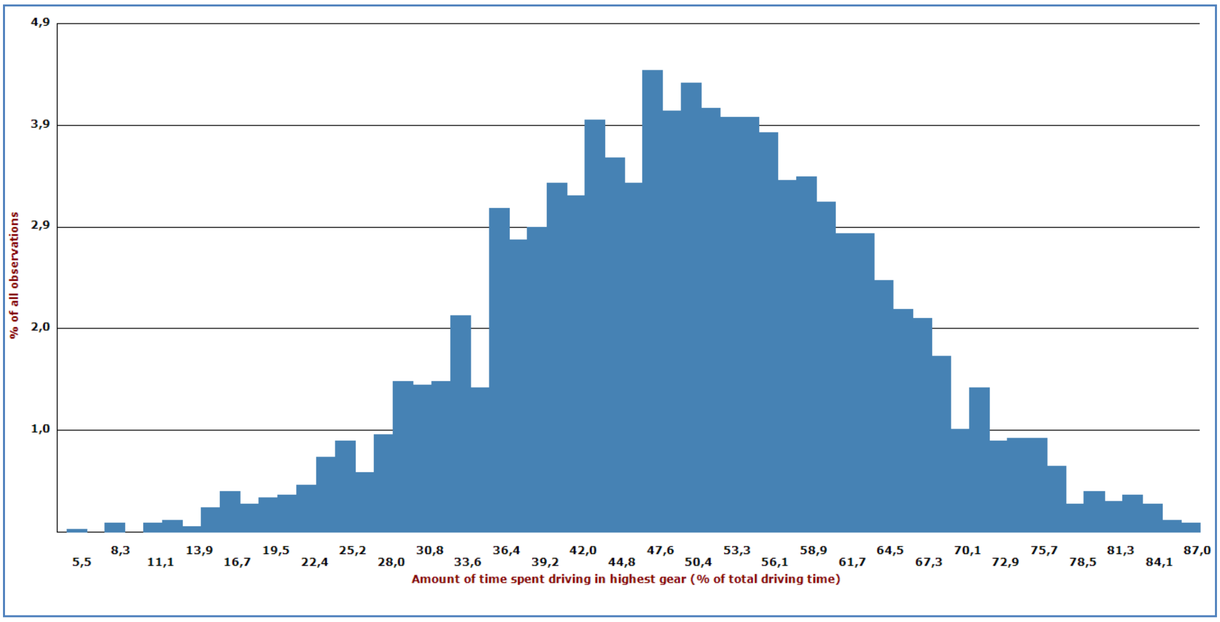

Figure 27 Histogram of amount of time spent using highest gear pr driver

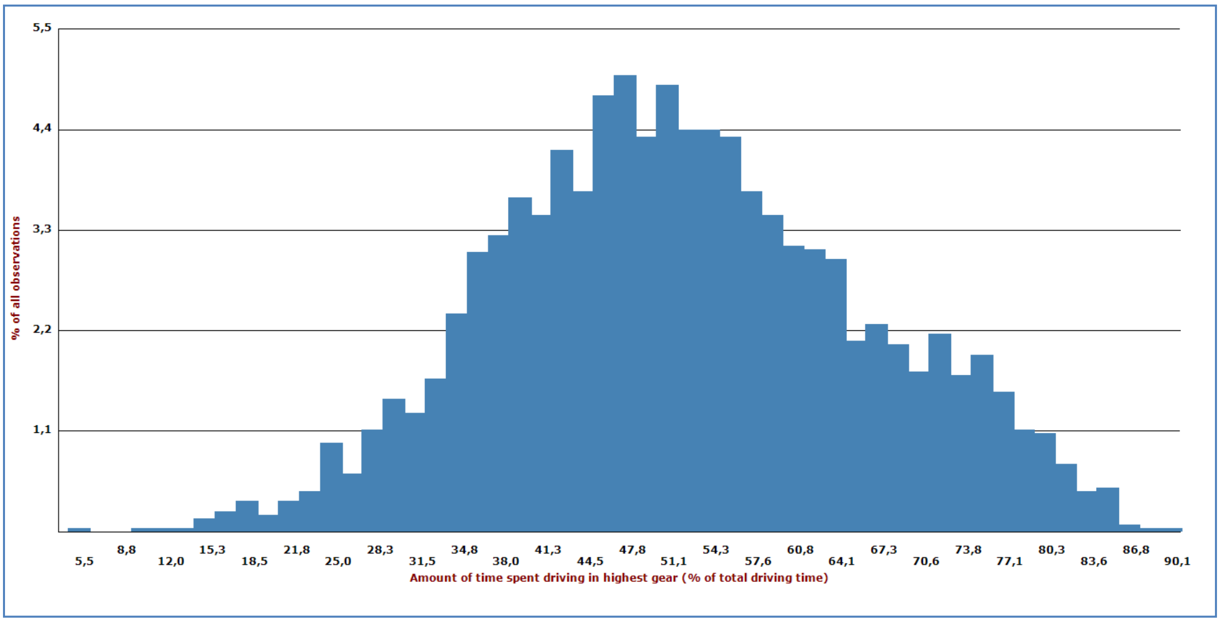

Figure 26Figure 27 shows the histogram for amount of time spent in highest gear for vehicles. Figure 27 shows the same histogram for drivers. There is no discernible differences between the two histograms.

Figure 28 shows the relationship between fuel consumption and the relative amount of time spent driving in highest gear per vehicle. Figure 29 shows the same relationship per driver. The figures are practically identical.

Figure 28 Fuel consumption by percentage of time spent in highest gear per vehicle

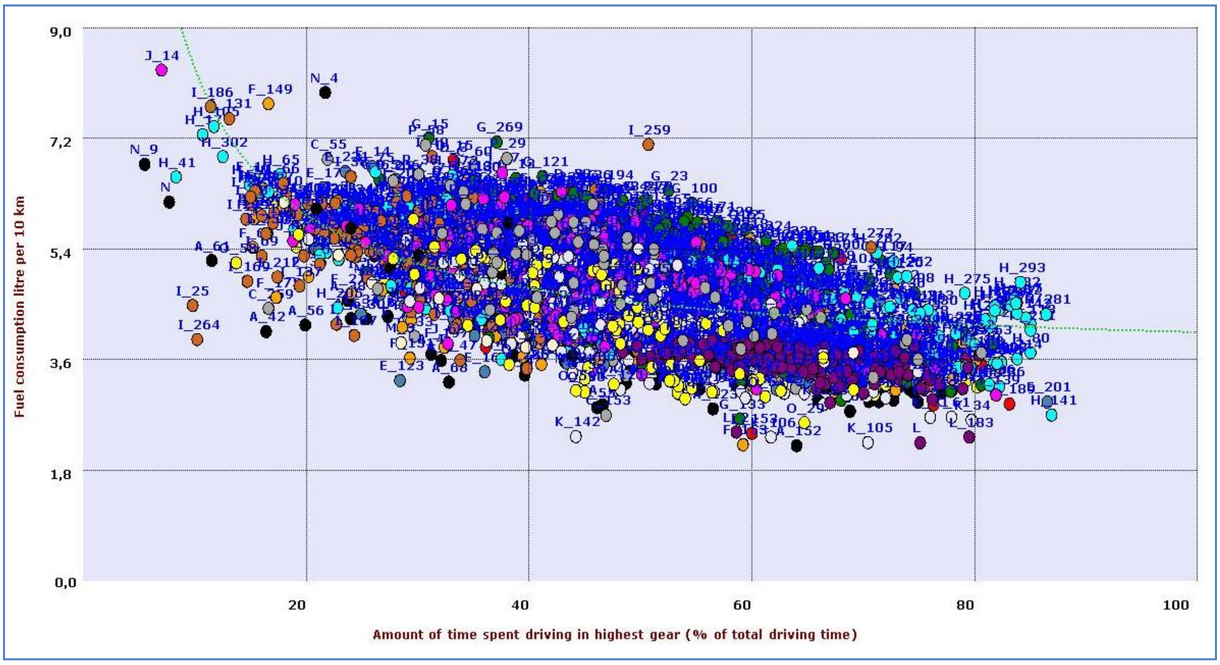

Figure 29 Fuel consumption by percentage of time spent in highest gear per driver

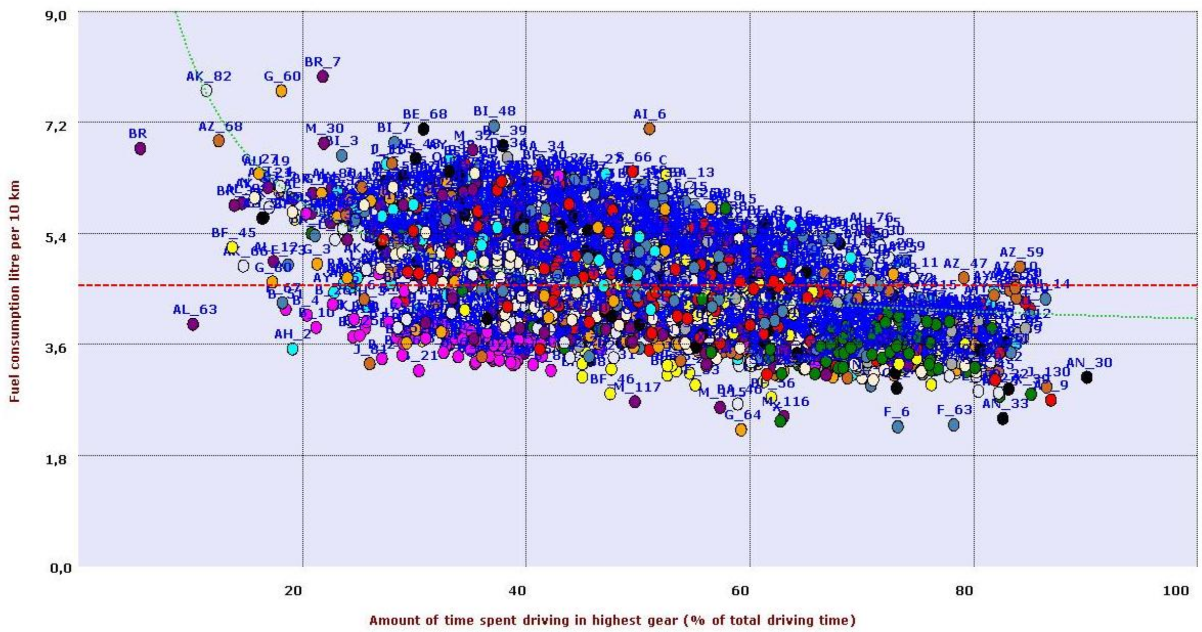

Both figures show the same trend, the longer time spent driving in highest gear, the lower is the fuel consumption. The effect of one extra percentage point of total driving time spent in highest gear is largest when that relative driving time in highest gear is low. The trucks use the lowest gears while climbing steep hills, therefore the engine load is high in lower gears. On the contrary, in the highest gear, the engine works with less load and therefore the fuel consumption pr distance will be lower. The effect of relative amount of time spent driving in highest gear on fuel consumption is significant, and the mean elasticity for the regression coefficient suggests that a ten percent increase in the relative amount of time spent driving in highest gear will reduce the fuel consumption by 1,2 percent [22].

We have used an inverse regression line in the figure above since the figure suggests that the effect of increasing the time driving in highest gear is largest when that relative time amount is low to start with. The more of the time already spent driving in the highest gear, the less is the effect of an additional amount of time spent driving in that gear. This effect is captured by the inverse regression line and is reported in Table 44 which is based on drivers' data.

Table 44 Effect on fuel consumption by increasing the relative amount of time spent in highest gear

Relative amount of time spent driving in highest gear	Average fuel consumption litre per 10 km	Effect of increasing driving time with 10 percentage point	Effect of increasing driving time relative to 10 percentage point in highest gear
10 %	8,3	0	0
20 %	5,9	-2,4	-2,4
30 %	5,1	-0,8	-3,1
40 %	4,7	-0,4	-3,5
50 %	4,5	-0,2	-3,8
60 %	4,3	-0,2	-3,9
70 %	4,2	-0,1	-4,0
80 %	4,1	-0,1	-4,1

The table shows the expected fuel consumption for different amounts of relative time spent driving in highest gear. The effect of increasing this relative time from 10 to 20 percent is 12 times higher higher than increasing it from 50 to 60 percent. Therefore, when a relative small amount of driving time is spent in highest gear, increasing this amount will have the largest effect. Again, this is the bivariate total effect of driving in highest gear. It may be that the direct independent effect from the multivariate regression model is smaller after controlling for indirect and spurious effects.

Table 45 Amount of time spent driving in highest gear for participants in driving behaviour course

	Before	After	Diff	t	Degrees-of-freedom	p-values	Signifi-cance	Fuel saving
AK	37,6	40,8	3,2	1,30	82	0,1972		-0,1
AY	52,3	47,9	-4,4	-1,11	58	0,2705		0,4
BA	58,7	62,4	3,7	1,77	72	0,0804		-0,4
E	50,3	54,9	4,6	2,97	109	0,0036	*	-0,3
L	43,8	48,5	4,7	0,90	24	0,3778		-0,1
M	48,9	51,7	2,9	1,84	168	0,0682		-0,4

[22] Note that with elasticities, we assess the increase in percent rather than percentage points and the effect is in *percent* change in fuel consumption per 10 km rather than in *litre* per 10 km.

Table 45 shows the difference in use of highest gear for participants in the driving behaviour course summer 2011. Most participants have increased the relative amount of time spent driving in highest gear, but we find a statistically significant effect for only one participant. The other differences are no larger than what can be expected from pure chance alone. The participant with the significant effect has an increase of 4,6 percentage points in the amount of driving time spent in highest gear, a very substantial increase given that the mean value is about 50%.

Rolling without engine load

Table 46 shows descriptive statistics for amount of driving time in percent spent rolling without any engine load. The table show data distribution pr vehicle and pr driver. The mean is 9,8% of driving time pr vehicle while the mean pr driver is slightly lower. The maximum amount of time spent rolling was almost 29,5% for the distribution per vehicle while its is practically the same per driver.

The variation in Table 46 suggests that there may be a potential for increasing the amount of driving time spent rolling without engine load, given that each truck each day have the same probability of using the same terrain and the same route with the same cargo. The lowest 10% of vehicles spend 2,9% of driving time rolling without engine load while the highest 10% of observations spend 17%. The values are about the same for drivers. Half of all vehicles and drivers each day spent more than roughly 9% of driving time without any engine load at all.

Table 46 Amount of driving time spent rolling without engine load

	Pr vehicle	Pr driver
N	3308	2702
Mean	9,8%	9,2%
Standard error	0,0949	0,0949
CV	1,0	1,0
Min (P0)	0,3%	0,3%
P10	2,9%	2,8%
P25	6,2%	5,8%
Median (P50)	9,2%	8,8%
P75	13,2%	12,6%
P90	17,0%	15,9%
Max	29,5%	29,1%

The percentiles suggest a skewed distribution since the distance from maximum to mean is larger than the distance from minimum to mean. This is confirmed by the histogram in Figure 28. There are some vehicles some days that spend a high amount of driving time rolling without engine load. The histogram also shows that the distribution has two peaks, there is one peak for very low amount of driving time spent rolling and one peak for the mean and median value of about 8-9% of driving time. This again may indicate that there is a potential for more rolling without using engine load. Figure 29 shows the same histogram for drivers. As already mentioned, the distributions are practically identical.

Figure 30 Histogram amount of driving time spent rolling without engine load pr vehicle

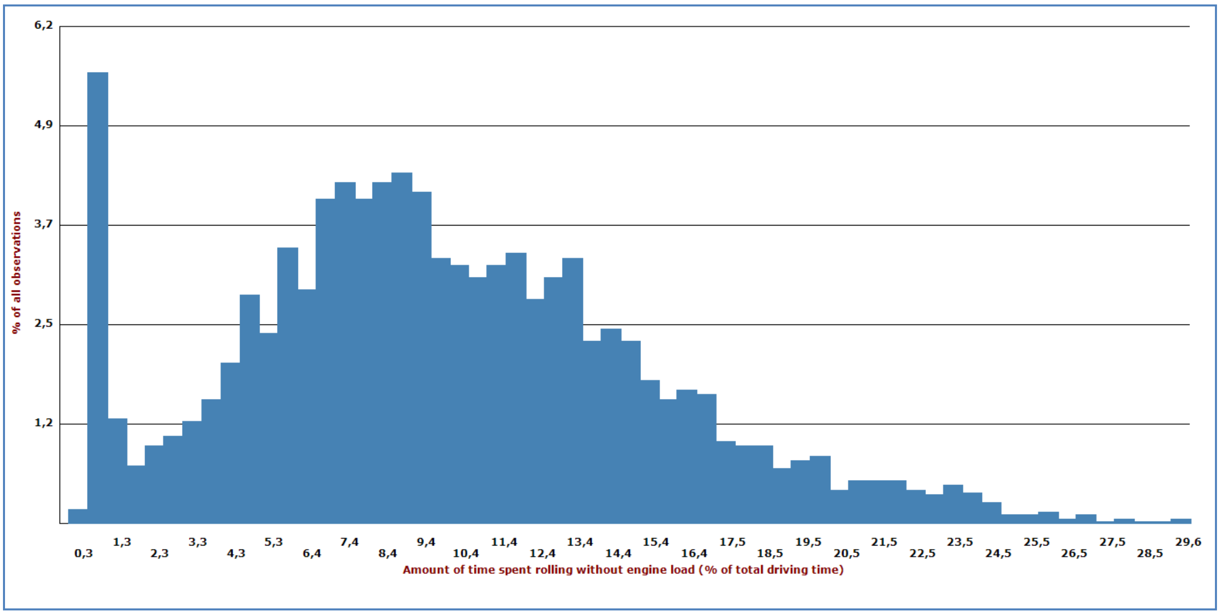

Figure 31 Histogram amount of driving time spent rolling without engine load pr driver

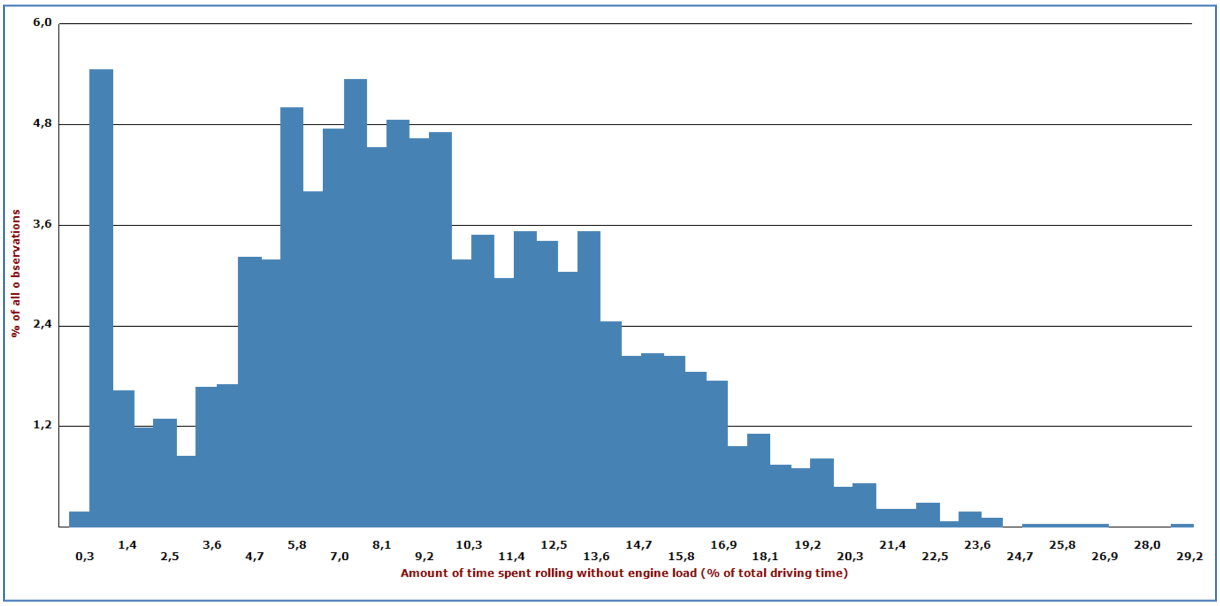

Figure 32 shows the relationship between amount of time spent rolling without engine load and fuel consumption pr 10 km per vehicle. Figure 33 shows the same relationship pr driver. The two figures basically show the same relationship. Based on vehicle data, the effect is slightly negative, if the amount of driving time spent rolling without an engine load is increased by 10 percentage points the fuel consumption will *decrease* by 0,062 litre. The effect is significant. Using drivers' data, the effect is the opposite, a similar increase will *increase* fuel consumption by 0,8 litre. This effect is also significant. In both models, the amount of variations in fuel consumption explained by rolling without engine load (R^2) is about zero. When we have as many data observations as we have both for vehicles and drivers, almost any effect is statistically significant even if the model explains nothing of variations in fuel consumption. Rolling without engine load also varies with the average speed since speed is required for rolling to take place. Rolling takes place downhill so the speed is obtained by going uphill with low average speed. The explanation for the low effect of rolling on fuel consumption may be that the amount of fuel saved rolling downhill is balanced by the increase in fuel consumption going uphill measured on a daily basis.

Figure 32 Amount of time spent rolling without engine load vs fuel consumption per vehicle

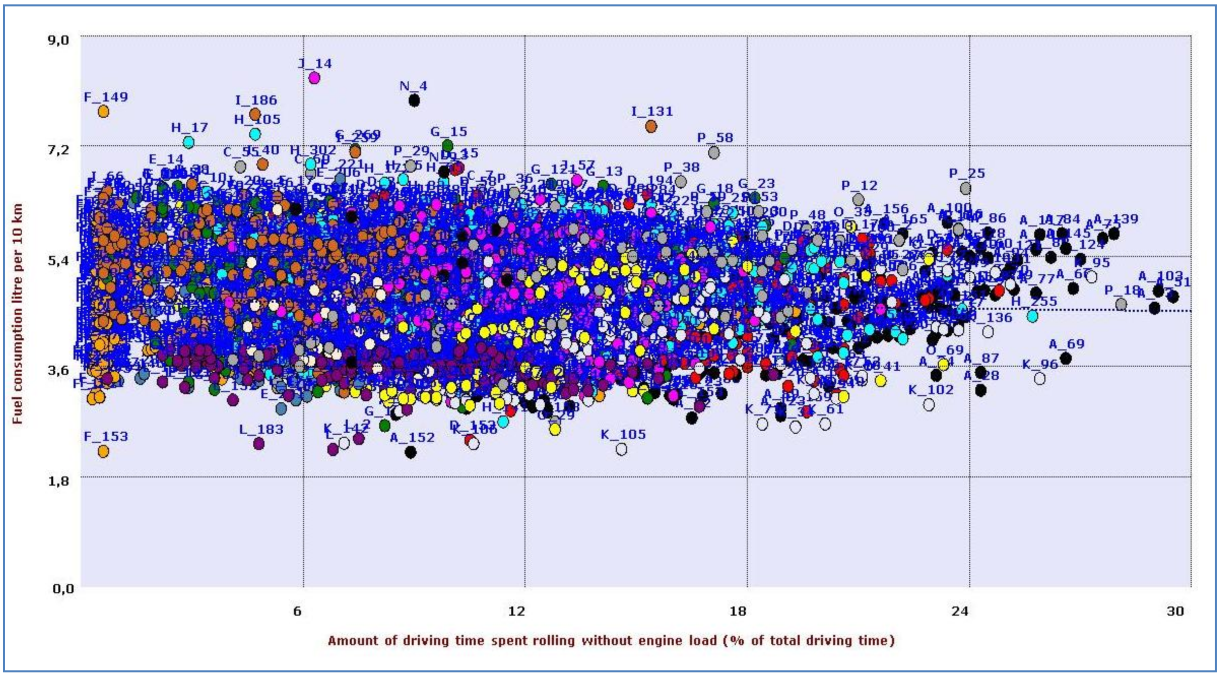

Figure 33 Amount of time spent rolling without engine load vs fuel consumption per driver

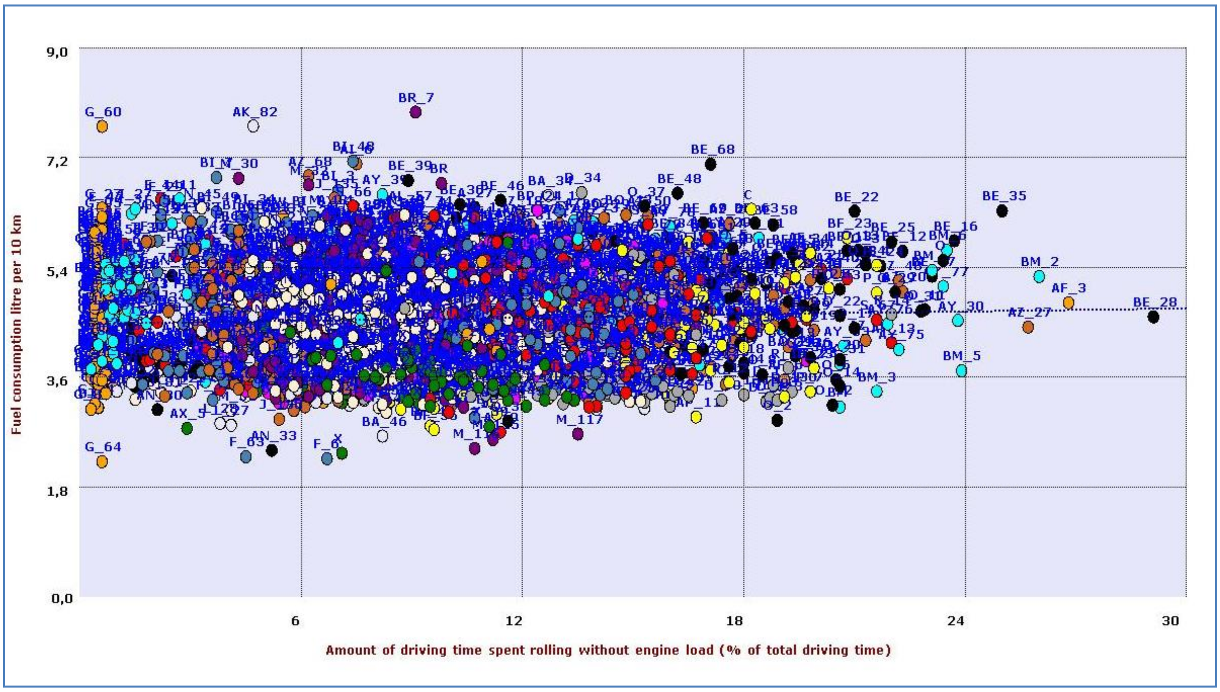

Table 47 shows the change in amount of driving time spent rolling without engine load for participants in the driving behaviour course. Five out of six drivers spend a significantly higher amount of time rolling without engine load after participating in the course.

Table 47 Amount of time spent rolling without engine load for participants in driving behaviour course

	Before	After	Difference	t-values	Degrees of freedom	p-values	Signifi-cance	Fuel saving
AK	7,3	8,2	0,9	2,00	87	0,0483	*	-0,1
AY	10,7	14,4	3,7	4,37	66	0,0000	*	0,4
BA	11,5	13,3	1,8	2,54	68	0,0133	*	-0,4
E	6,4	10,3	3,9	5,83	122	0,0000	*	-0,3
L	6,6	7,3	0,7	1,25	21	0,2243		-0,1
M	7,1	8,9	1,8	7,69	133	0,0000	*	-0,4

Weight distribution

Dynafleet also registers how much of the driving time that is spent with different weight loads. There are three loads, low, medium and high. A low weight is defined as weight of vehicle plus freight load up until 13 tonnes. A medium weight is between 13 and 28 tonnes and a high weight load is a weight of vehicle and freight load of more than 28 tonnes. We will look at the amount of time spent driving with high weight load. The first observations of weight loads were registered in Dynafleet at March 10th 2011, so there are considerably fewer observations among vehicles for this distribution.

Table 48 Descriptive statistics for amount of time in percent spent driving with a weight load of more than 28 tonnes

	Pr vehicle	Pr driver
N	2267	2507
Mean	64,3%	60,6%
Standard error	0,7829	0,7757
CV	1,2	1,3
Min (P0)	0,0%	0,0%
P10	0,6%	0,0%
P25	32,7%	18,6%
Median (P50)	80,3%	76,5%
P75	97,2%	97,0%
P90	100,0%	100,0%
Max	100,0%	100,0%

Table 48 shows descriptive statistics for relative amount of time spent driving with a high weight load of more than 28 tonnes (including the vehicle). Since minimum is 0% and maximum is 100% there is a considerable variation. On average, 64,3% of driving time for vehicles is spent with a high weight load while the same number for drivers is 3,7 percentage point lower . The median suggests that half of all vehicles at any day spend more than 80% of driving time with a high weight load. For drivers, the median is nearly 4 percentage points lower than for vehicles.

There are some striking differences between drivers and vehicles. If we look at the lower quartile (the value which cuts off the lower 25% of the distribution) we see that the proportion which drives with a high weight load is about 1,8 times higher for vehicles than for drivers. It is hard to interpret these differences since drivers do not choose the weight load, it is rather an effect of the market situation.

Figure 34 Histogram amount of driving time spent with high weight load pr vehicle

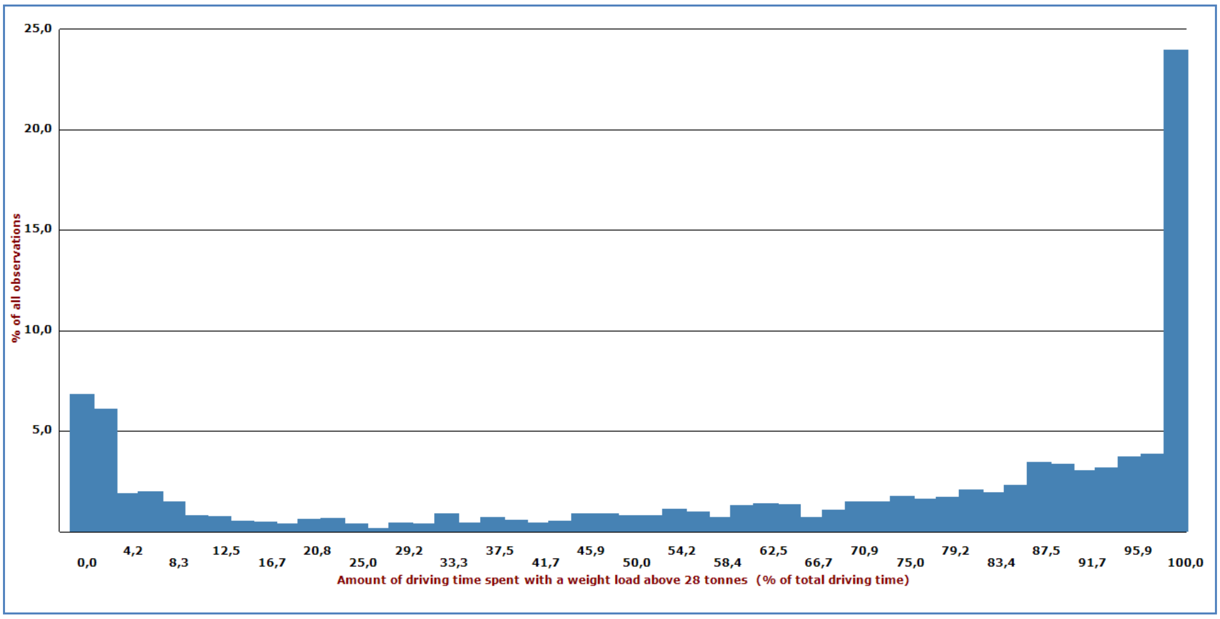

The histograms in Figure 34 and Figure 1 Figure 35 show that the distributions both for vehicles and drivers are heavily skewed to the right The upper quartile (P75 value) suggests that the 25% of all vehicles and drivers with the highest values spend more than 97% of their driving time with a high weight load.

Figure 35 Histogram amount of driving time spent with high weight load pr driver

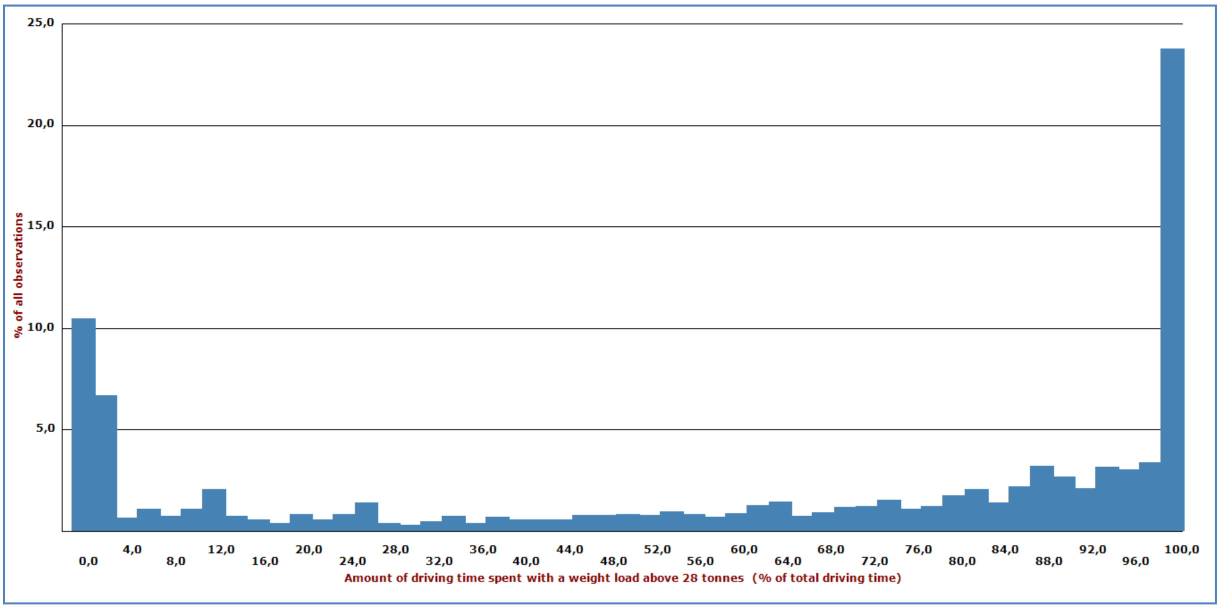

Figure 37 shows the relationship between driving time spent with a high weight load and fuel consumption per 10 km per vehicle. The linear effect of weight load on fuel consumption is obvious and significant - the more weight the higher consumption.

Figure 36 Amount of time spent driving with a high weight load vs fuel consumption per 10 km per vehicle

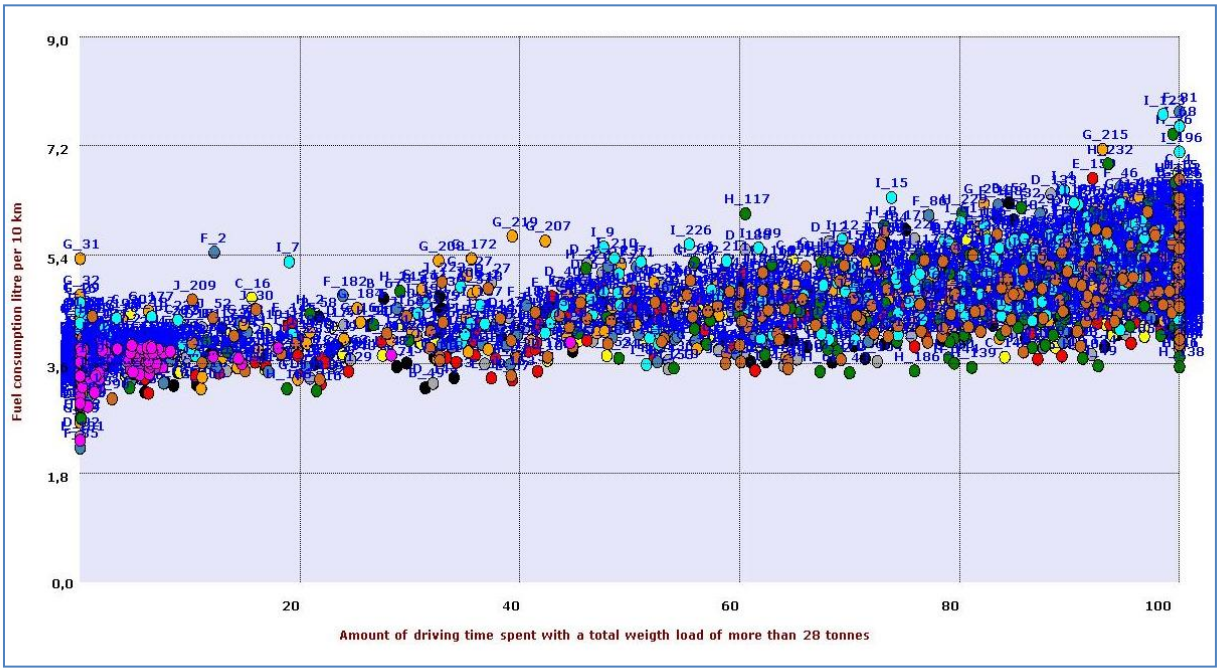

Figure 37 Amount of time spent driving with a high weight load vs fuel consumption per 10 km per driver

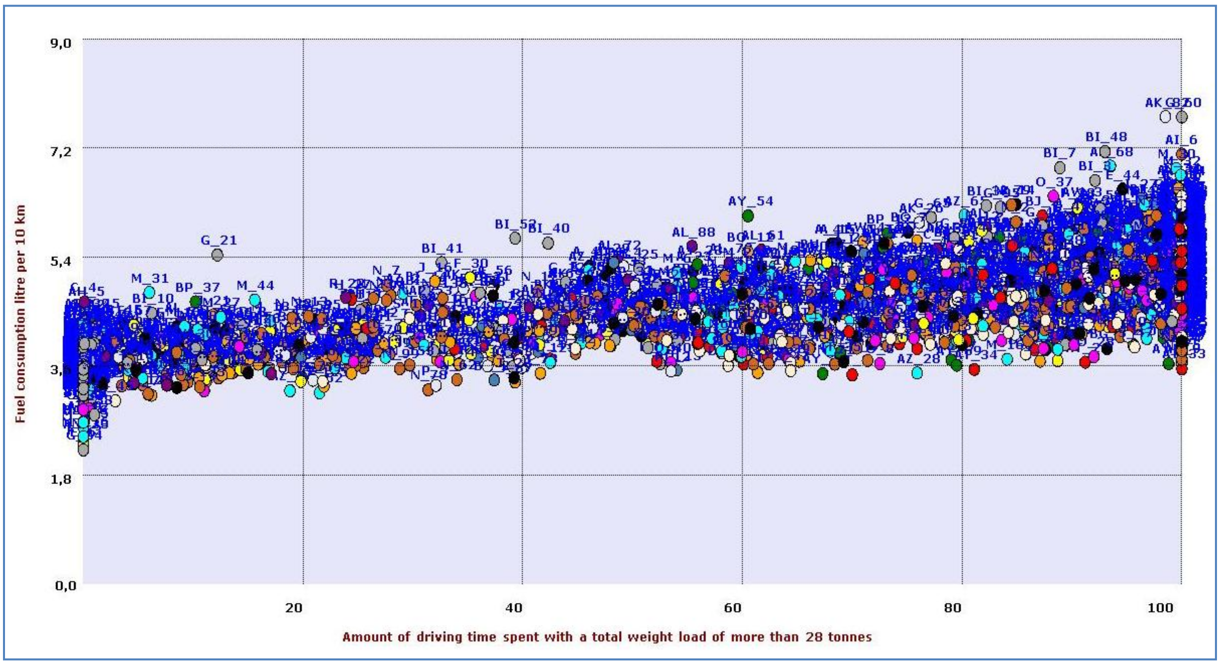

Figure 37 shows the same relationship for drivers. The linear trend is almost exactly the same but the variation around the regression line is slightly smaller for drivers since weight load explains 57% of the variation in fuel consumption for vehicles and 60% for drivers.

The magnitude of the effect may be smaller than expected. The estimated linear regression line from the drivers' data suggests that an increase of 10 percentage points in driving time with high weight load increases the fuel consumption by 0,16 litre per 10 km. A driver driving 100 000 km a year will increase the fuel consumption by 1600 litre diesel if the amount of driving time spent with a high weight load increases by 10 percentage points. This implies an increase in emission of CO_2 by almost 4,3 tonnes.

Analysis presented earlier indicate that increased fuel consumption as a result of increasing weight load can easily be offset by increased use of cruise control, more time spent driving in high gear, more use of automatic gear shift, higher average speed, more rolling without engine load and less amount of time driving with a high engine load. More optimal driving behaviour and better infrastructure would seem a reasonable price to pay in order for society to transport more goods per vehicle without increasing energy consumption or CO_2-emissions.

Weight load check

We have obtained data from the freight company which show the weight each vehicle transports on a specific task assigned to the vehicle on a specific date. A task, such as transporting goods from the factory in Sogndal to a shop in Oslo, can last several days if the destination is far away in Sweden or if the vehicle has to make a detour in winter because mountain passes are closed. Our analysis so far uses data registered on a vehicle or a driver per day. Since tasks and days are not directly comparable, analyzing fuel consumption and weight distribution pr task is a useful check on our analysis per day. In Dynafleet, information about driving behaviour and the amount of time spent driving with a high weight load is only available on a daily basis.

The data from the freight company shows information about the task, the vehicle assigned to the task, where the recipient of the transport is located and how much weight is assigned to the task. By matching with vehicle data in tracking reports from Dynafleet we can obtain information about distance travelled and fuel consumed for completing the specific task. All in all, the data obtained from the freight company covers 88 different tasks all in December 2011. It should be noted that December is a month with very varying driving and weather conditions. Some of the spread in fuel consumption is therefore attributable to conditions that we cannot control for since we do not know the exact driving condition on each route that was used in order to complete the tasks. Appendix B list all data items used in this weight analysis.

Every transport task is a delivery of freight from Lerum Fabrikker to a customer. Every transport task starts at the freight company address in Kaupanger, Western Norway. The freight company is situated next to Sognefjorden, the second longest fjord in the world. The freight company is therefore located at approximately sea level. In order to travel to Eastern Norway and Southern Norway the vehicles have to pass one mountain pass, usually Hemsedalsfjellet, with a peak altitude of 1100 meter. To travel to Trøndelag in summer the vehicles have to pass two mountain passes, Sognefjellet (1400 meter above sea level) and Dovrefjell (1100). In winter the vehicles pass Innvikfjellet (630) and Strynefjellet (943) before passing Dovrefjell in order to get to Trøndelag. The vehicles climb Utvikfjell, Strynefjell, Hemsedalsfjell and Sognefjell mountain passes from sea level. To travel to Møre og Romsdal (also in Western Norway) the vehicles travel narrow, winding roads but no mountain passes.

In Dynafleet, the weight indicator is not weight directly but rather how much of the driving time in percent that was spent with a weight load above 28 tonnes including the vehicle's own weight. A vehicle with trailer can weigh up to 15-20 tonnes. A fully loaded vehicle has about 30 tonnes of freight. If the vehicle weighs i.e. 15 tonnes it does not take more than 13 tonnes of freight weight for that vehicle to be counted as driving with a high weight load in the Dynafleet system. Therefore, the category high weight load in Dynafleet is not precise and do not discriminate enough between vehicles with different freight loads.

This analysis of freight data was therefore carried out as a check of the weight effect found in data from Dynafleet. Figure 38 shows freight weight and fuel consumption for the 88 tasks completed in December 2011. Each vehicle is identified with a letter code and with a number identifying the registration for that vehicle. So G_4 means the fourth registration for vehicle G. The colour codes are different geographical freight destinations. We differ between Eastern Norway, Southern Norway, Trøndelag, Western Norway and Møre og Romsdal. Figure 39 shows the colour codes used in Figure 38.

Figure 38 Freight weight vs fuel consumption

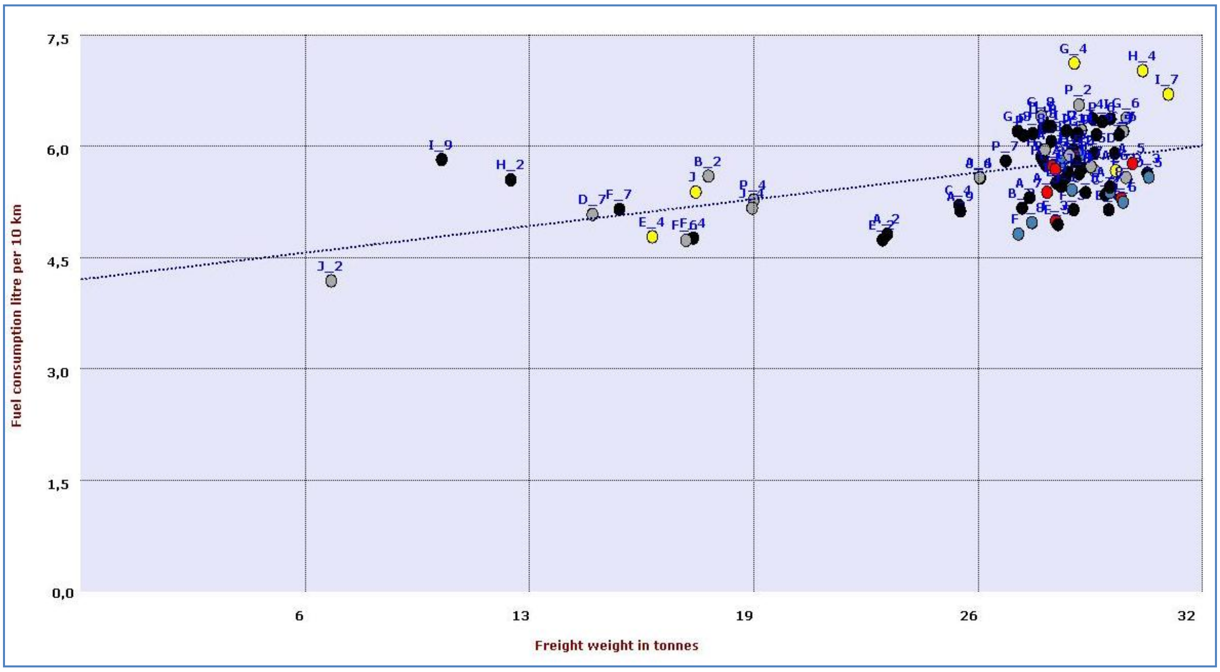

There is clear linear trend in Figure 38. The more freight weight, the more fuel consumption, as expected. Freight weight alone explain about 26% of variations in fuel consumption. For each extra tonne of freight carried the fuel consumption increases on average by 0,056 litre. Therefore, if the freight load increases with 10 tonnes the expected increase in fuel consumption is about 0,56 litre per 10 km.

It is difficult to detect a pattern for the colour codes used in the figure. The colour codes seem to be scattered randomly round the regression line. For height freight loads though the yellow points seem to be concentrated well above the predicted regression line. It may be that the rather poor infrastructure from Kaupanger to Møre og Romsdal with narrow winding roads increases fuel consumption compared to other destinations with better infrastructure. It could also be that the weather and driving conditions in these part of Norway in December is more demanding than weather conditions in Eastern and Southern Norway. There is much more precipitation and wind in Western Norway (including Møre og Romsdal) in December than in eastern and southern parts of Norway.

Figure 39 Colour code for freight deliveries

We performed an analysis of variance for differences in fuel consumption for different routes. Table 49 shows the result of the model estimation. The routes are categorized using the colour codes in Figure 39. The analysis was performed using a regression model with weight data and a dummy for each destination except Western Norway. The effect of the last category was measured by the constant term. By including the freight weight load in the analysis we compare vehicles travelling to different destinations with the same freight weight. In this manner we can isolate the differences between destinations from differences in freight load. The destinations are defined by the location of the first delivery for the freight transport.

Table 49 Regression analysis of freight weight and destination of first freight delivery

	Regression-coefficients	*Standard error*	*t-Stat*	*P-value*
Constant term	4,06408753	0,26575415	15,2926588	9,862E-26
Freight weight	0,06635472	0,00998984	6,64221784	3,1405E-09
Eastern Norway	-0,14289347	0,11813569	-1,20957073	0,22991937
Møre og Romsdal	0,33708988	0,20293741	1,66105345	0,1005218
Trøndelag	-0,4320809	0,19480094	-2,21806371	0,02931591
Southern Norway	-0,50223128	0,17886169	-2,80793095	0,00622952

When we include destination in the analysis the effect of a 10 tonne increase in freight weight increases from 0,56 litre per 10 km in the bivariate case cited above to 0,66 litre per 10 km.

The analysis shows that a vehicle travelling in Western Norway with a freight weight of 28 tonnes has an expected fuel consumption of 5,9 litre per 10 km. If the same vehicle with the same freight load travelled to Southern Norway the expected fuel consumption would be 0,5 litre lower pr 10 km. If the vehicle travelled to Trøndelag the fuel consumption would be 0,43 litre lower per 10 km. These effects are statistically significant. If the vehicle travelled to Eastern Norway the fuel consumption would be 0,14 litre per 10 km lower, but this effect is not statistically significant. On the other hand, if the vehicle travelled to Møre and Romsdal the expected fuel consumption would be 0,33 litre *higher* per 10 km, but this effect is also not statistically significant. It may that non-significant effects would be significant if we could collect more data for this analysis. The regression model explained 39% of the variation in fuel consumption.

Brake counter

Table 50 shows descriptive statistics for number of times the brakes have been applied per 100 km for both vehicles and drivers. Each time the brakes are applied, the Dynafleet computer register it. At the end of the day, the number of times the brakes have been applied is divided by distance travelled in 100 km.

Table 50 Descriptive statistics for brake counter per 100 km

	Pr vehicle	Pr driver
N	3308	2702
Mean	72,2	76,3
Standard error	0,5999	0,8912
CV	0,8	1,2
Min (P0)	0	0
P10	32	30
P25	48	47
Median (P50)	69,5	69
P75	92	94
P90	117	128
Max	281	392

Breaks are always applied by drivers, but since many vehicles are driven by several drivers the effects of drivers can be modified by the vehicles they drive. If vehicles data show less difference than drivers data it is reason to believe that vehicle properties cancel out driving behaviour, while if the opposite is the case it may be that some vehicles' attributes make it more difficult to reduce brake application.

On average, the brakes are applied 72 times per 100 km for vehicles and 76 for drivers (rounded numbers). The maximum is 281 number of times for a vehicle. The P90 value tells us that the 10% of vehicles with the highest values have applied the brakes more than 117 times per 100 km while the P10 value tells us that the 10% of observations with the lowest values have applied the brakes less than or equal to 32 times per 100 km. The median (69,5 times) is less than the mean (72) which suggests that he distribution is skewed to the right. Some vehicles at some days use the brakes a lot while most trucks at most days have a much lower application of brakes per 100 km.

The distribution for drivers are almost identical, except for the P90 and the maximum value. The maximum value is much higher for drivers than for vehicles. This may suggest that driving behaviour vary more than variations in vehicle attributes should indicate and that this behaviour is the key factor for reducing brake applications and obtain a smoother driving pattern. In other words, the potential for reduction of brake applications is more dependant on driving behaviour than on vehicles' attributes.

Table 51 Histogram brake counter per 100 km per vehicle

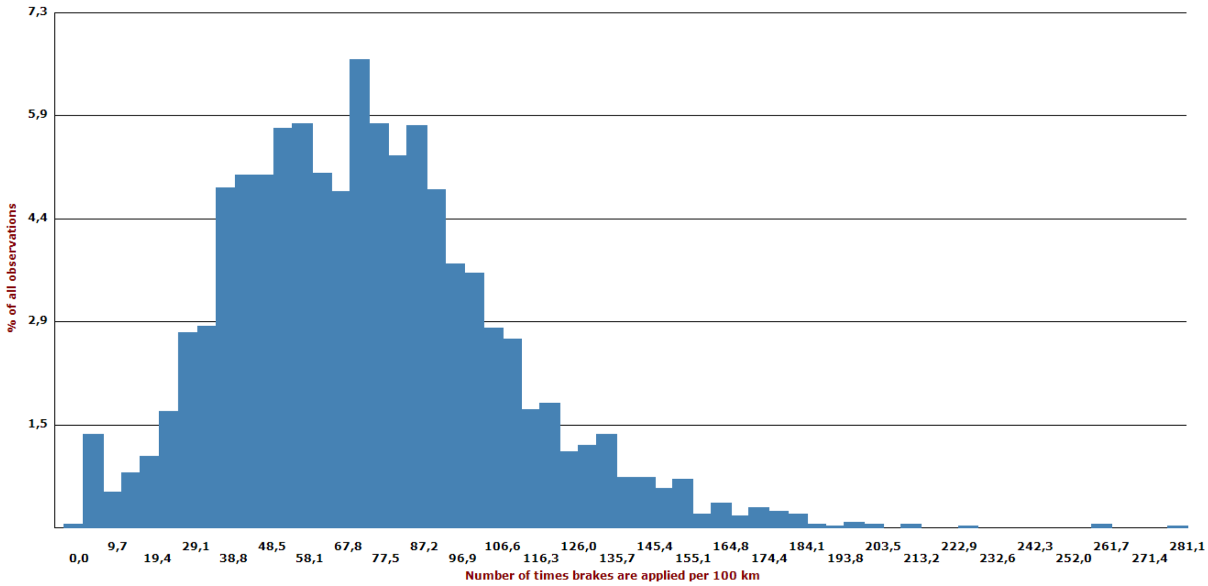

The histogram of the distribution for vehicles confirm that some vehicles at some few days use the brakes vary much compared to most other vehicles most other days. This is probably due to driving conditions on these days, especially in winter months.

Figure 40 shows the same histogram for drivers. The figure shows that this distribution is even more right-skewed and less symmetrical which, as discussed, suggests that driving behaviour probably play a role in determining the amount of times brakes are applied per 100 km.

Figure 40 Histogram brake counter per 100 km per driver

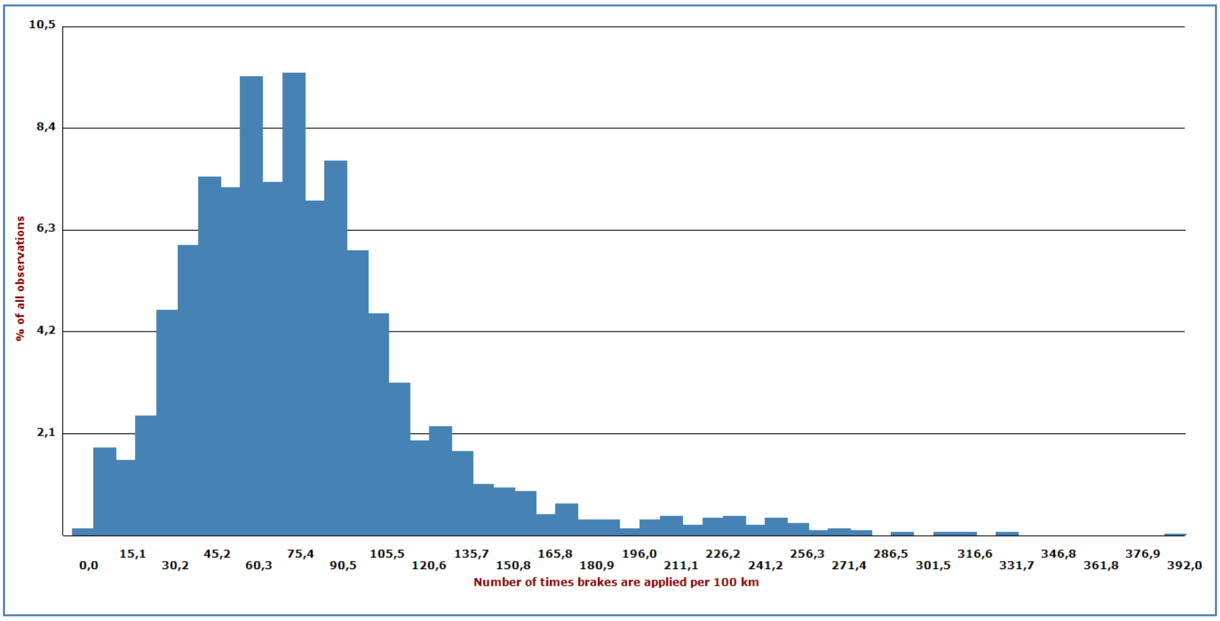

Figure 41 shows the relationship between brake application per 100 km and fuel consumption in litre per 10 km per vehicle. Figure 42 shows the same relationship for drivers. There is a positive significant linear trend for vehicles but not for drivers. For vehicles, if brakes are applied 100 times more per 100 km the expected increase in fuel consumption is 0,43 litre per 10 km. The model for vehicles explains only 3% percent of total fuel variations. The reason we don't find any significant effect for drivers is probably because some drivers some days have more extreme values for application of brakes.

Figure 41 Brake application per 100 km vs fuel consumption per 10 km per vehicle

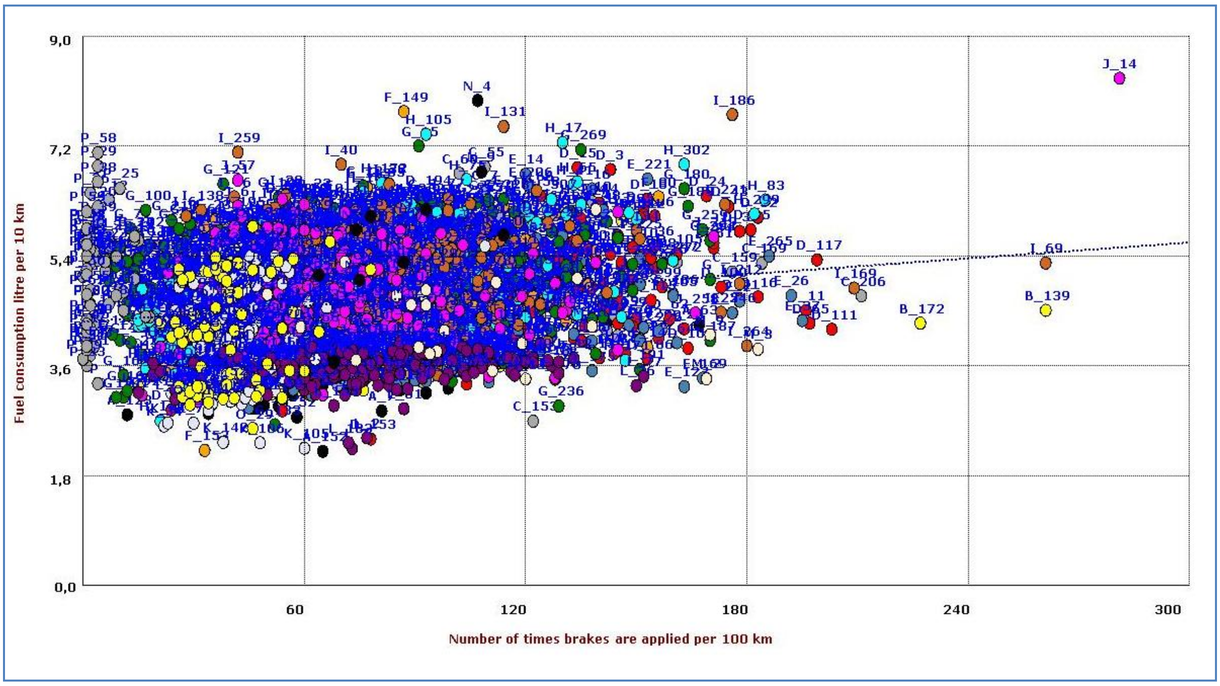

Figure 42 Brake application per 100 km vs fuel consumption per 10 km per driver

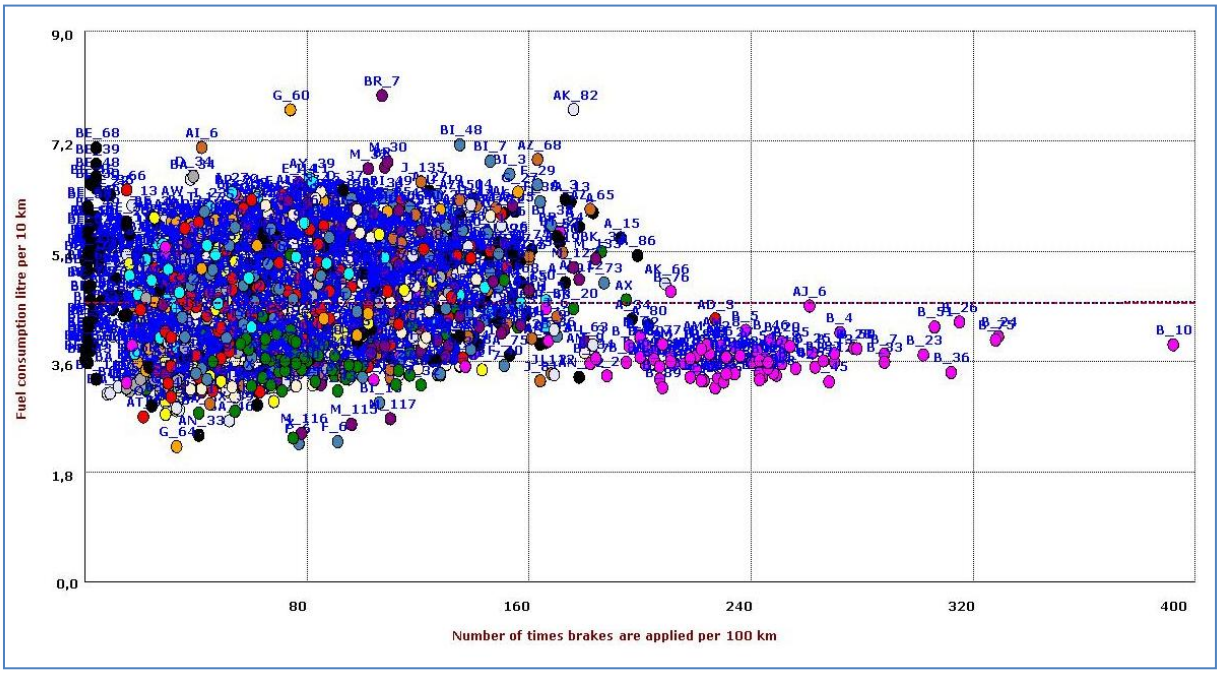

Table 52 Number of time brakes are applied for participants in driving behaviour course

	Before	After	Difference	t-values	Degrees-of-freedom	p-values	Signifi-cance	Fuel saving
AK	80,4	87,4	7,0	1,24	86	0,2171		-0,1
AY	50,3	70,1	19,9	3,21	75	0,0020	*	0,4
BA	26,0	24,2	-1,8	-0,47	71	0,6364		-0,4
E	61,7	68,7	7,0	1,69	113	0,0937		-0,3
L	116,4	80,0	-36,4	-3,00	17	0,0081	*	-0,1
M	80,8	77,9	-2,9	-0,64	148	0,5246		-0,4

Table 52 shows the distribution for number of brake applications per 100 km for participants in the driving behaviour course summer 2011. The table shows that after the course, one participant *increased* the use of brake applications per 100 km significantly while one driver significantly reduced it. It may be an indication that driving conditions and weather are more important factors for application of brakes than driving behaviour. On the other hand the variation among drivers is large. Also, the driver that reduced the application of brakes had the highest value before the driving course took place.

Stop counter

Dynafleet also registers how often the vehicle has stopped per 100 km. Table 53 shows the distribution for number of stops per 100 km for vehicles and drivers. As for the brake counter, the maximum and 90 percentile values are higher for drivers than for vehicles but the differences are very small.

Table 53 Descriptive statistics for number of stops per 100 km

	Pr vehicle	Pr driver
N	3308	2702
Mean	10,2	11,0
Standard error	0,1087	0,1550
CV	1,1	1,4
Min (P0)	1	1
P10	4	4
P25	6	6
Median (P50)	9	9
P75	13	13
P90	18	20
Max	60	62

Table 53 shows that vehicles on average stop about 10 times per 100 km while the corresponding number from the drivers' data is roughly one more stop per 100 km. The maximum value for vehicles is 60 and the top 10% of the distribution have more than 18 stops per 100 km.

Figure 43 Histogram number of stops per 100 km per vehicle

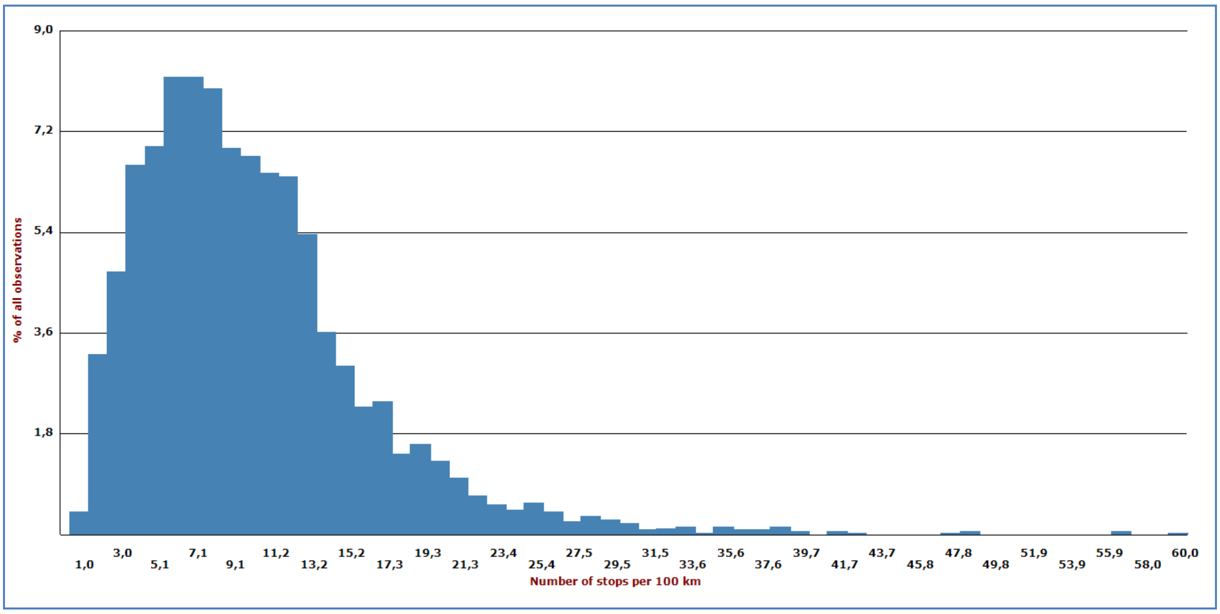

Looking at drivers' data, the 10% of them with the highest number of stops have more than 20 number of stops per 100 km. The upper quartile shows that 25% of all vehicles and drivers stop more than 13 times per 100 km while the lower quartile shows that the lowest 25% of both vehicles and drivers have 6 or less stops per 100 km. The median is slightly lower than the mean which suggests a right-skewed distribution, some few vehicles stop a lot on some few days. This is confirmed by the histogram in Figure 43 which shows the vehicles' distribution.

Figure 44 shows the same histogram for drivers. The figure shows that the distribution is right skewed as already discussed. The figure also shows one striking peak around 8 stops which is not present in the vehicles' distribution. We interpret this discrepancy between the two distributions as an indication that driving behaviour have an impact on number of stops per 100 km.

Figure 44 Histogram number of stops per 100 km per driver

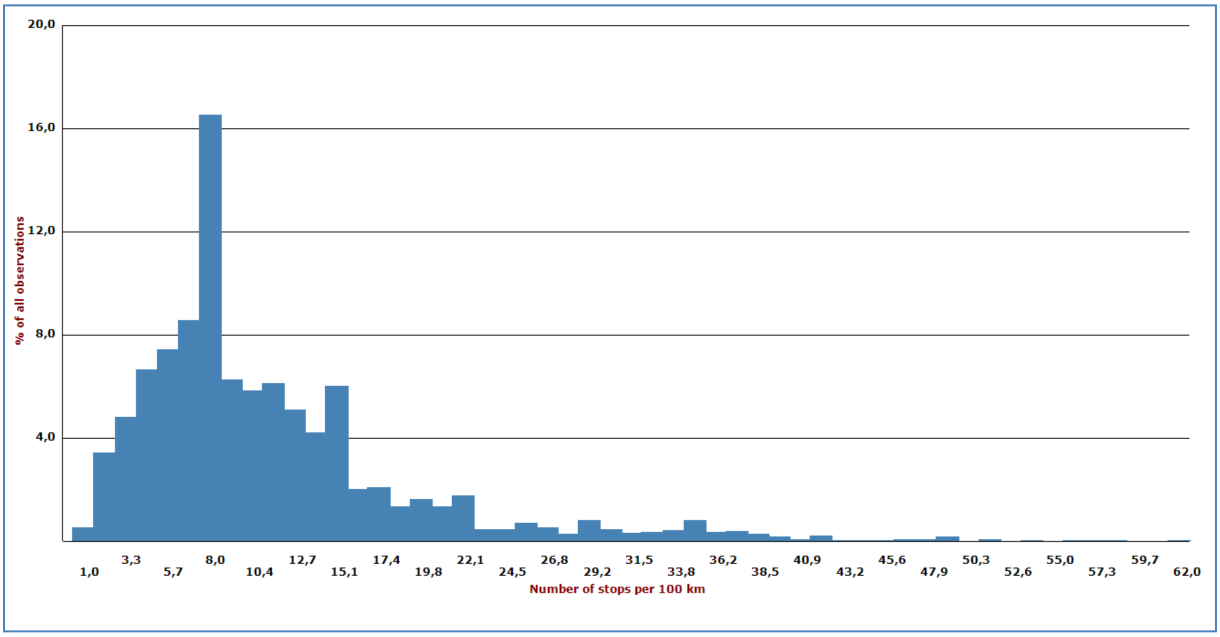

Figure 45 Number of stops per 100 km vs fuel consumption per 10 km per vehicle

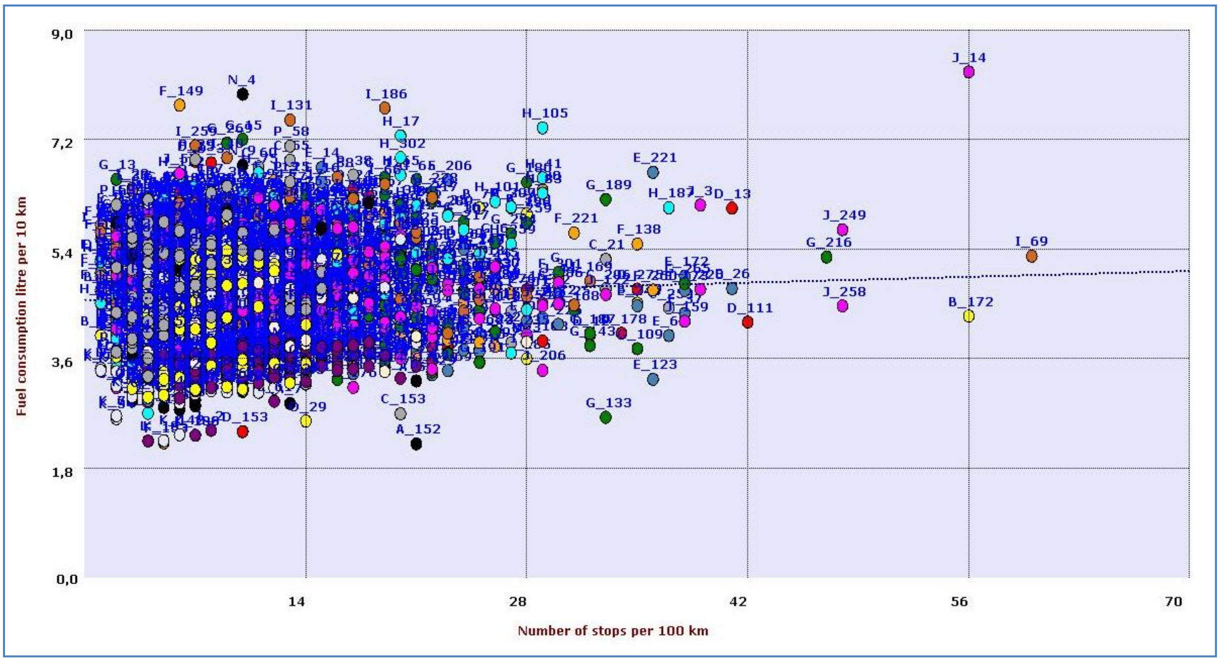

Figure 45 shows number of stops per 100 km vs fuel consumption per 10 km per vehicle. Number of stops by itself explains practically nothing of the variation in fuel consumption.

Figure 46 Number of stops per 100 km vs fuel consumption per 10 km per driver

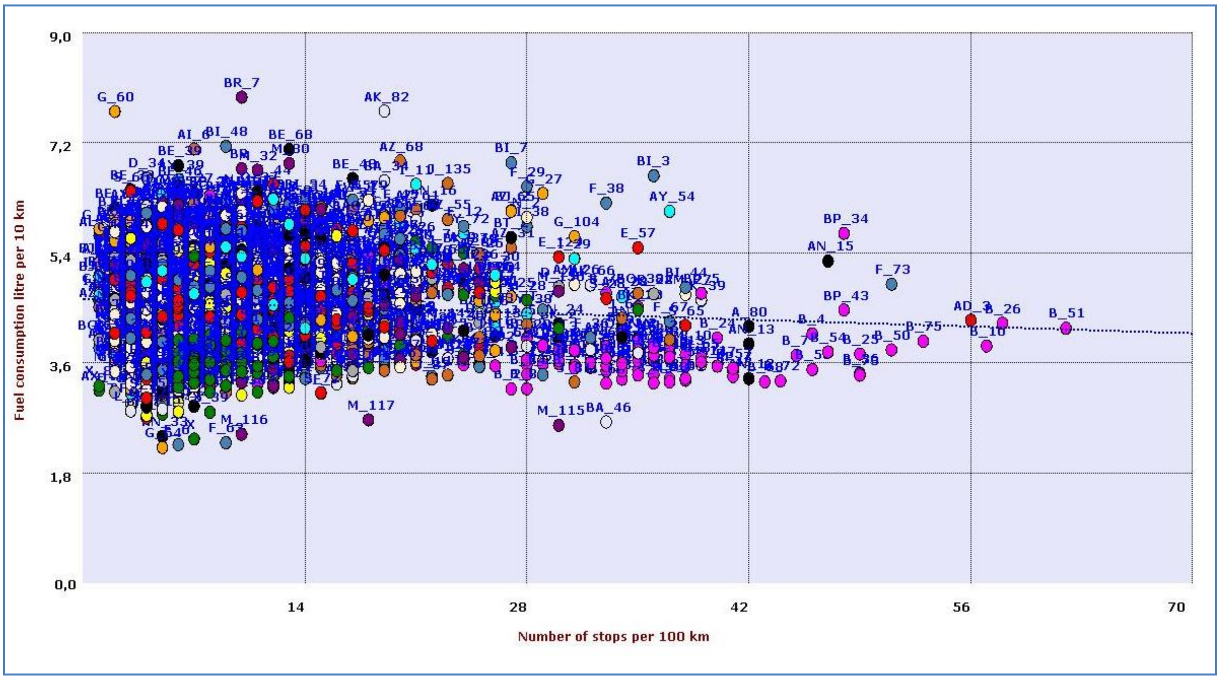

Figure 46 shows the same figure for drivers. The most interesting feature is that the sign of the regression coefficient changes from vehicles' data to drivers' data. Among vehicles, an increase in number of stops increase fuel consumption while for the drivers the effect is the opposite. This seems to be a result of a few drivers which have a large number of stops relative to other drivers. That feature is not so striking among vehicles. Again, it may indicate the impact of driving behaviour.

As for vehicles, number of stops alone explains practically nothing of the variation in fuel consumption between drivers. The regression effect itself is still significant, but this rather reflects a large number of observations than a substantial explanation of fuel consumption. The effect is weak, 10 less stops per 10 km will reduce the fuel consumption about 0,08 litre per 10 km for drivers.

Table 54 Number of stops for participants in driving behaviour course

	Before	After	Difference	t-values	Degrees-of-freedom	p-values	Signifi-cance	Fuel saving
AK	8,1	8,8	0,7	0,62	87	0,5353		-0,1
AY	7,1	9,5	2,4	1,72	76	0,0898		0,4
BA	11,1	12,0	0,9	0,50	71	0,6155		-0,4
E	9,9	10,0	0,1	0,08	108	0,9351		-0,3
L	9,5	8,2	-1,2	-0,62	19	0,5415		-0,1
M	9,4	11,1	1,8	2,09	148	0,0383	*	-0,4

Table 52 shows the distribution for number of stops per 100 km for participants in the driving behaviour course summer 2011. There is only one significant effect and that driver has *increased* number of stops. While driving behaviour may influence number of stops, it could also be that driving conditions and route attributes are having an even greater impact. We have no data to back up further speculations on this point.

Multivariate regression analysis

Finally, we can perform a multivariate regression analysis. We will use drivers' data to estimate the multivariate regression model. The units for the analysis are daily trips made from January 2011 to January 2012.

The dependent variable in the model is fuel consumption in litre per 10 km per day. The drivers' data and vehicles' data are joined by assuming that trips made on the same day, with the same time span, with the same distance travelled and with the same fuel consumption is the same trip in the two tables for drivers and vehicles. Consequently, when these criteria are satisfied, we know which vehicle was driven by which driver on that day. We need to couple drivers and vehicles since some independent variables measure vehicles' attributes while most independent variables are related to drivers.

We will use the independent variables discussed in the bivariate models or scatterplots above. We exclude number of stops per 100 km since it is correlated with number of brake applications per 100 km [23]. When number of brake applications and number of stops per 100 km are both included it is hard to separate the effect of one from the other since they are highly correlated. They both measure the same driving behaviour, therefore we use only one of them in the analysis.

In addition to the variables discussed above we introduce three additional independent variables. These are two dummy variables for two of the three engine types included in the model and a dummy-variable for winter months. The purpose of the winter dummy-variable is to control for different driving conditions in different parts of the year. Each daily trip is identified by a date and the dummy is set to 1 for days in December, January, February and March.

The vehicles included in the analysis are either 2010 models or 2011 models. We assume that models made in 2010 and 2010 do not differ much from each other in terms of engines and energy-efficient solutions. Consequently, we have not included model year as an independent variable.

The engine type also indicate the number of horsepower for that vehicle. Three engine types are included in the estimation of the multivariate regression model, these are

- D13C500,
- D13C540,
- D16G700.

We will use two dummy-variables for the last two engine types. A dummy variable has only two values, one if the vehicle in question is of the relevant engine type and zero if it is not. The effect of the first engine type is measured by the constant term. The last three digit in the engine type is the number of horsepower for each engine type. The effect of the dummy variable for engine type

[23] The linear correlation coefficient between number of stops and number of brake applications per 100 km is 0,64 for the data material applied in the analysis.

D13C540 therefore measures the difference in mean fuel consumption between a vehicle with 500 horsepower and a vehicle with 540 horsepower, assuming values for all other independent variables are held constant. And, of course, the effect of the dummy for engine type D16G700 measures the difference in mean fuel consumption between vehicles with 500 horsepower and vehicles with 700 horsepower, again under the assumption of constant values for all other independent variables.

We also control for whether the vehicle is a semi-trailer or a vehicle with a separate trailer. A semitrailer has 16 wheels that have a width of 385 mm, a height that is 55% of that width and where the rim has a diameter of 22,5 inches. A vehicle with a trailer has 22 wheels and the wheels have a width of 265 mm, the height is 65% of that width and the rim diameter is 19,5 inches. It may be that different number and types of wheels will give different rolling drag or rolling resistance which again will have an impact on fuel consumption. A semitrailer dummy will control for this.

The independent variables in the multivariate regression model are :

- average driving speed (idle running excepted),
- relative amount of driving time per day the vehicle is running idle,
- relative amount of driving time per day the vehicle uses cruise control,
- relative amount of driving time per day the vehicle is driven by automatic gear shift,
- relative amount of driving time per day the vehicle is driven with an engine load above 90% of maximum torque,
- relative amount of driving time per day the vehicle is driven with the highest gear shift,
- relative amount of driving time per day the vehicle rolls without using engine power,
- relative amount of driving time per day spent driving with high weight load,
- number of brake applications per 100 km per day,
- a dummy variable for engine type D13C540,
- a dummy variable for engine type D16G700,
- winter months (1 for December, January, February and March, 0 for every other month)
- a dummy for semitrailer (1 for semitrailer, 0 for vehicles with separate trailers).

The point of the multivariate regression analysis is to determine the independent, separate effect of each single independent variable. When several variables vary concurrently and when there is a relationship between them and fuel consumption, it will be difficult to assess the independent effect of one of them only by looking at its bivariate relationship with fuel consumption. This is simply because it is hard to say that any effect we find really comes from one specific variable when the others are varying at the same time. Average speed, for instance, vary in much the same way as driving in high gear or driving with a high engine load. If they all vary much the same way, how do we know what is the real effect of just one of them on fuel consumption?

This is the problem of estimating direct and indirect effects. The bivariate effects we have found above is the total effect of each independent variable. This total effect includes the direct effect on fuel consumption and the indirect effects that go through other independent variables. If average speed has an effect of the amount of time driven in highest gear there are two effects of average speed on fuel consumption, one direct and that is mediated through the independent variable driving in highest gear. If there were only these two effects from average speed the bivariate effect would be equal to the sum of them. In reality there will be many other indirect effects on fuel consumption from average speed. We do not know each of them but we will estimate the direct effect of average speed in this section. We also know the bivariate effect of average driving speed from the discussion above. Therefore, even if we do not know the specification of each indirect effect, we knot the sum of these indirect effects since this sum is the difference between the bivariate effect and the direct one.

For the independent variables that mediate indirect effects of other independent variables, the difference between the bivariate and direct effect are not indirect effects but spurious effects. In the bivariate case, these independent variables will be assigned an effect that is too high since some of that total effect really is the effect of some other independent variables that the mediating independent variable is partly an effect of.

To sum up: An indirect effect is an effect of one independent variable on the dependent one through an intermediate variable. A spurious effect is an effect of one independent variable on the dependent one that is actually an intermediate effect of *another* independent variable.

To correct for this we apply a multivariate regression model where we introduce all the independent variables in one single model. In a multivariate model we control for the effect of other independent variables. Controlling means letting only one independent variable vary while assuming that the value of the others are constant. With such a model, we can answer questions like: What is the isolated, controlled effect of using more automatic gear shifts when all other independent variables are assumed to be constant? What is the difference between two vehicles where one uses cruise control and the other does not when they have the same values on all other independent variables?

Let us illustrate what we mean by direct, indirect and total bivariate effects by using a hypothetic causal model with three of the independent variables from the multivariate model. Figure 47 shows the model. We assume that average speed influence the potential for use of cruise control and automatic gear shift. Therefore, these independent variables, cruise control and time spent driving in highest gear, are intermediate variables between average speed and fuel consumption which is dependent variable in the model.

Figure 47 Hypothetic causal model

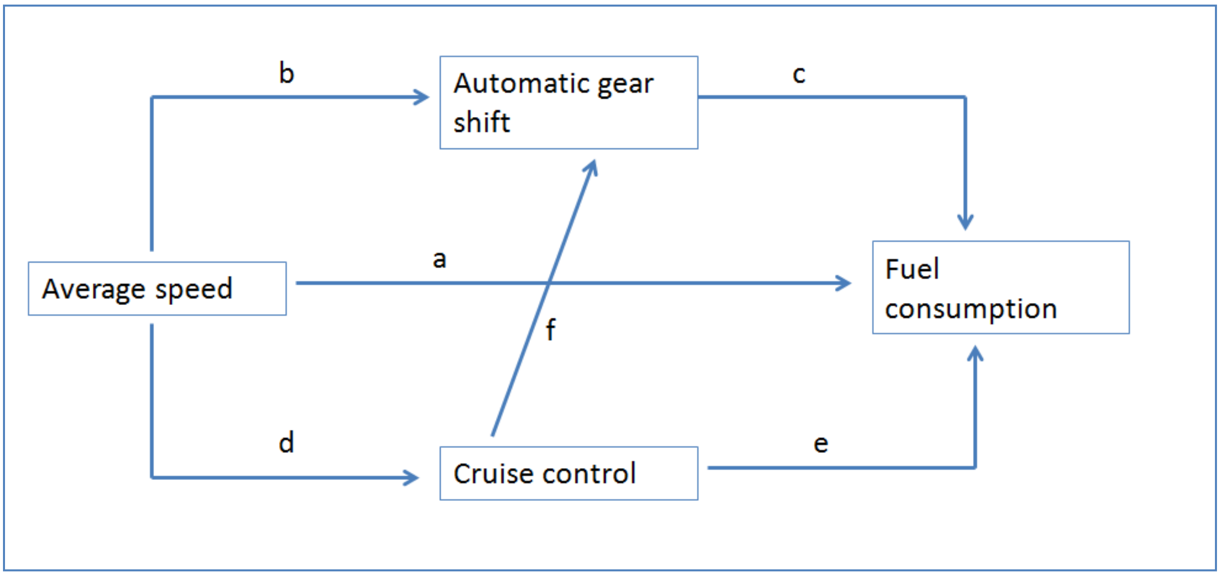

We also assume that cruise control influence the potential for use of automatic gear shift. The more use of cruise control, the more even speed which will allow for more use of automatic gear shift. Time spent driving in highest gear therefore is a intermediary variable between cruise control and fuel consumption. For the two independent variables use of cruise control and driving in highest gear the effect from average speed is a spurious one. If average speed was not included in the model, we would assign effects to the other two independent variables that really would be an effect of average speed. This would have been a spurious, uncontrolled effect. In the bivariate models where fuel consumption is only related to one independent variable, say use of cruise control, the effect of cruise control will be too high since we do not control for spurious effects, effects that have an impact on *both* use of cruise control and fuel consumption.

Table 55 Total effect of independent variable average speed

Effect	Value
Direct	a
Indirect	b*c
Indirect	d*e
Indirect	d*f*c
Total	a+(b*c)+(d*e)+(d*f*c)

Table 55 shows the total effect of independent variable average speed if we assume the causal model in Figure 47 is a correct representation of reality. We can see from the table that there are three indirect effects from average speed. One goes through cruise control, one goes through automatic gear shift and one goes through both. Since changes in average speed will cause changes in use of cruise control and time spent driving in highest gear, part of the effects going through these last two independent variables will have their origin in average speed.

We have not formulated a total causal model where all indirect and spurious effects are specified. In our context it is sufficient to know that the difference between the total and direct effect of say average speed is the sum of all the indirect effects.

What effect is the most interesting one, the total effect or the direct one? When say average speed changes, so do a lot of other indicators too. The most reasonable effect could therefore be said to be the total one. When we interpret the direct effect of average speed we assume constant values for all other independent variables. We might say then that if average speed is changing by one km per hour, fuel consumption changes by y amount assuming all other independent variables are constant. But this assumption is really an unrealistic one. It is difficult, if not impossible, to change only average speed without changing the values of many other independent variables. Consequently, the direct effect is useful mostly for estimating the *relative* effect of each independent variable in order to determine their causal contribution. If, on the other hand, the purpose is to give a realistic picture of what an increase in average speed may mean for fuel consumption, the bivariate effect may be the better choice.

The independent variables have different measurement units which makes it difficult to assess which of them has the highest effect. Using elasticities will enable us to compare the effects with each other since elasticity are unit free, they measure the percentwise effect of one percentage change in the independent variable on the dependent. Equation 16 shows the formula for calculating elasticity. In the equation, b_i is the regression coefficient for the i'th independent variable, $\overline{X}_i$ is the mean for the same i'th independent variable, $\overline{Y}$ is mean for the dependent variable, fuel consumption, and E_i is the elasticity for the i'th independent variable.

Equation 16 Calculation of elasticity in a multivariate regression model

$$E_i = b_i * \frac{\overline{X}_i}{\overline{Y}}$$

Any elasticity must be calculated at a specific point for the independent and dependent variable. For all elasticities, we use mean values for the independent and dependent variable as the point where the percentwise effect on the dependent variable from one percent increase in the independent is measured.

In the multivariate model, the effects of all independent variables are assumed to be linear except for average speed and relative amount of time spent driving in highest gear. These two independent variables are assumed to have an inverse effect on fuel consumption. This is the functional form or trend we found when interpreting the effect of these variables in the bivariate case. This trend is captured by taking 1 over the independent variable and include it in the model. The inverse effect is asymptotically decreasing, the effect of an increase in the independent variable is larger when the value of that independent variable is small. For all other variables, we use linear effects which means the effect of an increase in the independent variable is the same regardless of the level of that independent variable.

There are three possible weight variables in Dynafleet. There is one variable for the amount of time driving with a weight load less than 13 tonnes, one for the amount of time driving with a weight load between 13 and 28 tonnes and one variable for amount of driving time with a weight load of more than 28 tonnes. If we know the values of two of the variables the value of the third is given as a residue of the others. Therefore, If we include all three variables in the model we will have a high degree of correlation between them. This is not a desired property [24]. We have solved this by only entering one variable, the amount of time spent driving with a high weight load. These weights include vehicle's own weight. Above, in the section *Weight load check,* we discussed the appropriateness of using this variable as a weight indicator.

We assume that the independent variables engine load, average speed and amount of time spent driving in high gear capture the effect of infrastructure, landscape and terrain. Vehicles spend a lot of time travelling on routes with steep hills and winding roads. This restrict the speed the vehicles can travel with and how often they can drive in the highest gear. Steep hill climbing also implies a high engine load. Therefore, these variables are assumed to capture structural effects not primarily related to driving behaviour.

On the other hand, variables like use of cruise control, automatic gear shift, rolling without engine load and running idle are supposed to capture differences in driving behaviour. We assume that the

---

[24] http://en.wikipedia.org/wiki/Multicollinearity

value of these independent variables are more determined by drivers' choice than by infrastructure and terrain.

We do not control for drivers in the multivariate model. The idea is that when all relevant variables are controlled for, the unexplained variation that is left is randomly distributed between drivers. Drivers do have an impact on fuel consumption though their driving behaviour. Drivers choose when to use cruise control, they choose when to use automatic gear shift, they choose when to let the engine run idle and so on. When all these variables are included in the model, variation in their values will reflect variation in driving behaviour. The assumption is therefore that drivers as such do not have a significant impact on fuel consumption, their impact is mediated through their values on different driving behavioural indicators included in the model.

Earlier we analyzed the difference between drivers using analysis of variance (ANOVA). The purpose of that analysis was to detect whether different drivers had significant different fuel consumption. We found that of five drivers, one had significant different fuel consumption than all the others but the other four had no significant differences between them. In that analysis we used drivers as indicators of different driving behaviour. In the multivariate regression model differences between drivers are measured explicitly by differences in driving indicators. Therefore, the multivariate analysis is a useful check of earlier analysis.

A residual plot will reveal whether rest variation from the multivariate regression model is randomly distributed between drivers. If our assumption is true, residuals will show no trend and different drivers will not have systematically different residual values.

Ideally, we would also control for the exact weight load and the route travelled by each driver on each day included in the model. The weight variable we use has some deficiencies as already discussed. Also, the route travelled will have an impact if some drivers consistently travel on some routes that have different road quality or less steep hill climbing than other routes. In Dynafleet driving indicators are measured on a daily basis. Drivers typically perform tasks, they transport goods from Lerum factory in Sogndal to some customers in other parts of Norway. There can be more than one task on one day or a task can last several days. There is consequently a discrepancy between the distance travelled or the fuel consumed by a driver per day and per task. Since driving indicators are only measured per day we use day as the entity for our analysis. The model then assume that drivers are randomly distributed on different routes and that there are no significant differences in freight load between them. We consider this to be a reasonable assumption but we will point it out explicitly so that considerations can be taken when interpreting the model's results.

Table 56 shows the result of estimating the multivariate regression model. A total of 1823 observations are used in the estimation.

Table 56 Multivariate regression model. Dependant variable fuel consumption in litre per 10 km.

	Regression coefficients	*Standard error*	*t-Stat*	*P-value*
Constant term §	2,61678592	0,1104821	23,6851567	2,975E-108
Average speed #	50,9502687	7,00405173	7,27439926	5,1583E-13
Running idle *	0,01246221	0,00099649	12,5060918	1,7605E-34
Use of cruise control *	-0,00337865	0,00063882	-5,28888832	1,38E-07

Use of automatic gear shift *	-0,00720667	0,00068239	-10,5609861	2,4259E-25
Above 90% of maximum torque *	0,0860449	0,00135609	63,4506055	0
Use of highest gear *#	20,6243657	1,61463528	12,7733897	7,7451E-36
Rolling without engine load*	-0,02152745	0,00128984	-16,689953	2,8029E-58
High weight load *	0,00233401	0,00024593	9,49059427	6,9627E-21
Brake application	0,00085312	0,00022676	3,76215466	0,00017382
Engine type: D13C540 ¤	0,21517204	0,03010862	7,14652517	1,2852E-12
Engine type: D16G700 ¤	0,65940031	0,02161996	30,4996076	3,233E-165
Winter season++	0,13603108	0,01267032	10,7361953	4,1282E-26
Semitrailer %	-0,11230945	0,02620989	-4,28500298	1,9234E-05

* Measured as amount of driving time in percent with use of the specified property # Inverse effects § Constant term is engine type D13C500 ++1 for winter months December, January, February, March, 0 otherwise ¤ 1 if the vehicle has the relevant engine type, 0 otherwise %1 if the vehicle is a semitrailer, 0 otherwise

Table 56 shows the results of estimating the multivariate regression model. Every independent variable has a significant effect on fuel consumption.

The vehicles use more fuel in winter months. Controlling for other attributes, the fuel consumption is about 0,14 litre higher per 10 km in winter months than in the rest of the year. If we only use winter months as independent variable we find that on average the fuel consumption in winter months is 0,22 litre per 10 km. Above, in the bivariate ANOVA-analysis of winter months and fuel consumption, we found that the fuel consumption is 0,4 litre higher in winter. This last analysis included all data from the drivers database. The bivariate effect discussed here only include the drivers data that can be matched with vehicles data and therefore are used in the multivariate regression model. The sum of the indirect effects of winter season on fuel consumption is therefore estimated to be 0,08 litre per 10 km, these are the effects of winter season that are mediated through other independent variables.

The effect of engine types must be evaluated relative to the engine type with 500 horsepower. The effect of this engine type is the constant term. If we assume mean values on all independent variables the expected fuel consumption for vehicles with 500 horsepower is 4,5 litre per 10 km. The engine type with 540 horsepower is expected to have an average fuel consumption that is 0,22 litre higher per 10 km. Vehicles with 700 horsepower will have a fuel consumption that is 0,66 litre higher per 10 km than the vehicles with 500 horsepower, assuming the two types of vehicles have the same use of cruise control, highest gear, the same speed, both travel in winter months and so on.

If we look at the bivariate effect of the vehicle types, the vehicles with 700 horsepower will on average use 0,35 litre more per 10 km more than vehicle types with 500 horsepower. The bivariate, total effect is less than the direct one because there are indirect effects from horsepower that have the opposite effect on fuel consumption as the direct effect. More horsepower can allow the vehicle to travel with higher speed in difficult terrain, more horsepower can allow for more use of cruise control, more time spent in highest gear and more use of automatic gear shift. All these independent variables have a negative effect on fuel consumption, i.e. more use of cruise control will lower fuel consumption. Therefore the indirect effect of horsepower on fuel consumption through i.e. use of cruise control is negative since a positive effect on cruise control is multiplied by a negative effect of cruise control on fuel consumption.

Semitrailers have less fuel consumption per 10 km then vehicles with separate trailers. As already mentioned, this may be because the rolling drag is different for the vehicle types. On average, a semitrailer will use 0,11 litre less per 10 km assuming constant values on the other independent variables.

If average speed is increased by 10 km an hour from 40 to 50 km per hour, the estimated fuel consumption in litre per 10 km is reduced by 0,25 litre, if we assume constant mean values for all other independent variables. A vehicle travelling 100 000 km a year would then save 2 500 litre if average speed could be increased from 40 to 50 km per hour. If the average speed was increased from 30 to 40 km an hour the expected fuel savings would be 0,42 litre per 10 km all other things being equal. If a vehicle travels 100 000 km, this would imply fuel savings of 4 200 litre per year and a reduction in CO2-emissions of 11 tonnes. For average speed, the effect is highest when the average speed is low since we have used an inverse functional form for that independent variable in the multivariate regression model.

Table 57 Estimated effect of speed increase on fuel consumption from multivariate regression model

Average speed in km pr hour	Estimated fuel consumption in litre pr 10 km	Savings relative to previous speed	Savings relative to 30 km an hour	Total bivariate effect
30	5,5	0,00	0,00	8,5
40	5,1	-0,42	-0,42	6,7
50	4,9	-0,25	-0,68	5,5
60	4,7	-0,17	-0,85	4,8
70	4,6	-0,12	-0,97	4,3
80	4,5	-0,09	-1,06	3,9

The effect of speed in the multivariate regression model is summed up in Table 57 which can be compared to the bivariate effects of average speed in Table 37. The effect of driving speed from the multivariate model is the direct effect while the effect from the bivariate model is the total effect which also include indirect effects. The bivariate effects in the rightmost column in Table 57 are calculated with the same data sample as the multivariate model. The bivariate effect calculated earlier was based on drivers' data alone, while the multivariate model is based on a joining of data from both drivers and vehicles. When average speed varies, so does use of cruise control, use of automatic gear shift, rolling without engine load and so on. In the multivariate case we need to make assumptions about the values of the other independent variables in order to produce an outcome from the regression model. We have assumed mean values for all other independent variables. It is very probable that vehicles with low average speed will have a quite different use of cruise control than the mean value. This is captured in the bivariate case since it measures the total effect of average speed.

The indirect effects of average speed are larger than the direct ones when average speed is low. When average speed increases, the indirect effects become negative so that the total effect is lower than the direct one for higher average speed values. This confirms our assumption that increased average speed facilitates use of other driving indicators that are beneficial for lower fuel consumption such as more use of cruise control, more time spent in highest gear, more rolling without engine load and more use of automatic gear shits. The indirect effects of average speed mediated through these other indicators reinforces the negative effect of increased driving speed on fuel consumption.

An increase of 10 percentage points in driving time spent running idle will increase fuel consumption by 0,12 litre per 10 km all other things being equal. A similar increase in the amount of time spent with an engine load of more than 90% of maximum torque would increase the fuel consumption by 0,86 litre.

An increase in the amount of driving time spent in highest gear shift from 30 to 40 percentage points would reduce the fuel consumption with 0,17 litre. A similar increase from 40 to 50 percentage points would reduce fuel consumption by 0,10 litre per 10 km. Again, all other variables are assumed to have constant mean values.

Table 58 Estimated effect of increase in relative amount of time spent driving in highest gear on fuel consumption from multivariate regression model

Amount of time spent driving in highest gear	Estimated fuel consumption in litre per 10 km	Savings relative to previous speed	Savings relative to 10% of driving time	Total bivariate effect
10 %	6,3	0,00	0,00	9,3
20 %	5,3	-1,03	-1,03	6,4
30 %	4,9	-0,34	-1,37	5,4
40 %	4,7	-0,17	-1,55	4,9
50 %	4,6	-0,10	-1,65	4,6
60 %	4,6	-0,07	-1,72	4,4
70 %	4,5	-0,05	-1,77	4,3
80 %	4,5	-0,04	-1,80	4,2
90 %	4,5	-0,03	-1,83	4,1

Table 58 sums up the effect of an increase in the relative amount of time driving in highest gear. The effect of an increase is highest when use of highest gear is low to start with. This is captured by the inverse regression effect. The effect of an increase of 10 percentage point is lower in the multivariate regression model than in the bivariate regression model. This is because the indirect effects of driving in highest gear have the same sign as the direct one so that the total effect of driving in highest gear is larger than the direct one. The total, bivariate effect of in the rightmost column in Table 58 is based on the same data sample as the multivariate model.

Again, we see the same tendency for amount of driving time spent in highest gear as we saw for average speed. When the amount of driving time in highest gear is high the sum of the indirect effects from that variable is negative so that the total effect is less than the direct one. This indicates that values of other independent variables that have a beneficial effect on fuel consumption also increases as a consequence of more driving time in highest gear.

An increase of 10 percentage points in the amount of driving time spent rolling without engine load would reduce the fuel consumption by 0,22 litre per 10 km. A similar 10 percentage point increase in the amount of time driving with high weight load will increase the fuel consumption by 0,02 litre per 10 km. This suggests that weight load itself does not have a decisive effect on fuel consumption. As discussed above, we consider the validity of variable weight load in Dynafleet to be weak since the operationalization of high weight load is set to low so that many vehicles with high own weight will automatically fall into this category.

The effect of brake applications is weak, but significant. If brakes are applied ten times more per 100 kilometre the increase in fuel consumption is expected to be 0,01 litre per 10 km. The effect of using cruise control is also negative and statistically significant, a increase of 10 percentage point in the relative amount of driving time using cruise control will lower fuel consumption by 0,03 litre per 10 km. A similar increase in driving time spent using automatic gear shift will lower the fuel consumption by 0,07 litre per 10 km.

Table 59 shows elasticities for the regression coefficients. The table shows that the four most important independent variables in descending order are amount of time spent driving with engine load of more than 90% of maximum torque, average speed, use of automatic gear shift and amount of time spent driving in highest gear.

Table 59 Elasticities

	Mean	Elasticity
Average speed	0,016	0,175
Running idle	12,277	0,033
Use of cruise control	5,370	-0,004
Use of automatic gear shift	93,523	-0,145
Above 90% of maximum torque	14,395	0,266
Use of highest gear	0,021	0,093
Rolling without engine load	9,238	-0,043
High weight load	68,086	0,034
Brake application	73,598	0,013
Engine type: D13C540	0,111	0,005
Engine type D16G700	0,162	0,023
Winter season	0,265	0,008
Semitrailer	0,839	-0,020

Mean value for dependant variable=4,620 litre per 10 km.

Figure 48 shows the residual plot for the multivariate regression model presented in Table 56. The X-axis in the plot is the predicted response (YHat) for each data item. The Y-axis is the residual for each data item, the difference between the actual (Y) and predicted value (YHat) for each data item used in the model. The residuals should be randomly distributed without any obvious trend or pattern in their variation. Such a trend can indicate that an important explanatory variable is left out of the model. Also, the residuals should not be distributed in series, which indicate that they are not independent of each other as they are assumed to be.

As already discussed, the effect of drivers when all driving behavioural indicators are controlled for are assumed to be randomly distributed. Each driver is identified with a colour code in the plot. If our assumptions are correct, the colour codes should not display any pattern or trend but appear randomly distributed. Some drivers share colour code but the visual inspection should still detect any pattern if it is present in the plot. Each driver is also identified with a number so that G_55 is the 55th registration for driver G.

There are no obvious trends or patterns in Figure 48. The different colours also seem to be randomly distributed. A visual inspection therefore seems to confirm our assumptions. There might be a slight bow-shaped tendency in the distribution of residuals, this may suggest that a second-degree polynomial should be applied for one or more of the independent variables. We have not consider this tendency to be strong enough to invalidate our assumptions.

There is one item that has a specific high absolute residual value, this is the 28th registration for driver AL. This driver had a low fuel consumption on December 8th 2011 with 3,9 litre per 10 km. The driver's predicted value from the model given values for the independent variables is 6,4 litre per 10 km. The driver had a daily average speed of just over 48 km per day which is low given that the P10 value for average speed is 42,2 km per hour per day. The driver spent 38,6% of driving time running idle which is high given that the P90 value is 24,5%.   The driver also spent just a little over 10% of the driving time driving in highest gear, that is very low given that the P10 value is 33,7% of driving time.

All in all, the multivariate regression model explains 91 % of variation in fuel consumption [25]. This is a very large proportion which indicates that we have a model with good explanatory power.

Figure 48 Residual plot for multivariate regression model

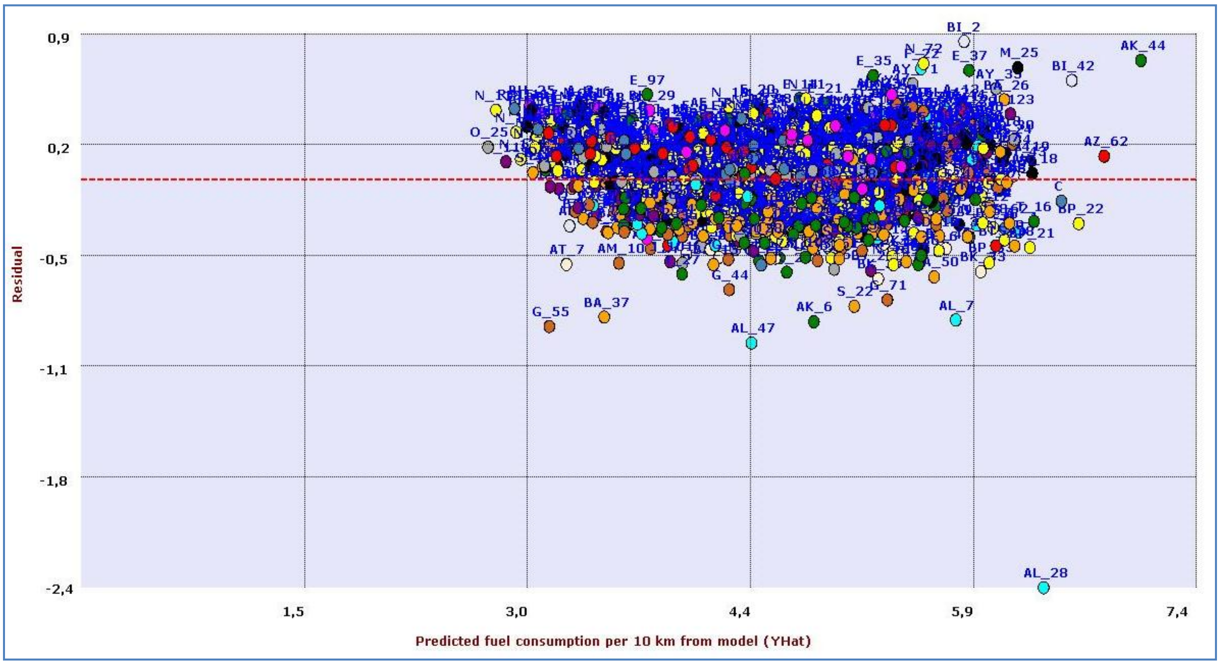

Finally, Table 60 shows the total and direct effects for each independent variable except average speed and relative amount of driving time spent in highest gear which we have presented earlier. The total, bivariate effects are estimated with the same data sample which is used in estimation of

[25] Measured by Adjusted R^2

the multivariate model. The table shows the biggest differences in effects for engine types. The effect of 700 horsepower as opposed to 500 horsepower has a direct effect of 0,66 litre per 10 km and a total effect of 0,35 litre per 10 km. The effect of 540 horsepower has a direct effect of 0,22 litre per 10 km and a total effect of 0,16 litre per 10 km. The indirect effects from engine types are mediated through other independent variables that varies between different engine types. Use of cruise control, use of automatic gear shifts, power outtake, average speed etc can be different in different engine types. We do not know exactly how these indirect effects are mediated, but since the direct effect is higher than the total the sum of all the indirect effects must be negative.

For the semitrailer variable the situation is quite the opposite. In absolute terms, the total effect is larger than the direct one. This means that some of the effect we assign to the semitrailer variable in the bivariate case is mediated through other independent variables. These other independent variables must be intermediary since no other independent variable can cause a change in a vehicle's type. Therefore, the difference between the total and direct effect of the variable semitrailer cannot consist of spurious effects, these effects must be indirect ones.

Table 60 Total and direct effects

	Direct effect	Total effect	Difference
Running idle	0,012	0,016	0,004
Use of cruise control	-0,003	-0,023	-0,019
Use of automatic gear shift	-0,007	-0,024	-0,017
Above 90% of maximum torque	0,086	0,080	-0,007
Rolling without engine load	-0,022	0,001	0,022
High weight load	0,002	0,016	0,014
Brake application	0,001	0,005	0,004
Engine type: D13C540	0,215	0,164	-0,051
Engine type D16G700	0,659	0,347	-0,312
Winter season	0,136	0,221	0,085
Semitrailer	-0,112	-0,339	-0,227

The difference in direct and total effect for rolling without engine load in Table 60 is also interesting. The total effect of this variable is positive but not statistically significant in the bivariate case. In the multivariate model the effect of rolling without engine load is negative and statistically significant, more rolling without engine load decreases fuel consumption. We believe there are spurious effects that cause the bivariate effect of rolling without engine load to be weak and positive. Variables like use of engine load higher than 90% of maximum torque has an impact on rolling without engine load. Rolling without engine load is therefore an intermediary variable between power outtake and fuel consumption. When maximum power outtake is high the vehicle is driving uphill on poor infrastructure with little potential for rolling without engine load. Therefore, power outtake increases fuel consumption and lower the potential for rolling without engine load. Since the effect of power outtake on fuel consumption and the effect of rolling without engine load on fuel consumption have different signs, power outtake is shadowing the effect of rolling without engine load when we do not control for it in the bivariate case. The effect of rolling without engine load will be less in the bivariate than in the multivariate case or the effect may even change sign.

Similar reasoning can be done with the variable average speed. Rolling without engine load is also a intermediary variable between average speed and fuel consumption. The effect of speed on rolling without engine load and fuel consumption has opposite signs, when speed increase fuel consumption goes down while the potential for rolling without engine load increases. Therefore the effect of average speed is shadowing the effect of rolling without engine load. More rolling without engine load takes place when average speed increases so we cannot tell them apart. The effect of rolling without engine load is therefore assigned to average speed. When we compare different levels of rolling without engine load for drivers with *identical* speed the effect becomes clear, more rolling without engine load decreases fuel consumption.

In this case the effect of rolling without engine load is positive in the bivariate case and negative in the multivariate. Measured in absolute values, the negative value is also 22 times higher. From the discussion above we assume this is because the relative amount of time spent with an engine load of more than 90% of maximum torque is shadowing the effect of rolling without engine load in the bivariate case.

Does rolling without engine load has an effect or not? What is the correct effect of rolling without engine load? We have established a fact that every truck driver know: Rolling without engine load is beneficial for fuel consumption when the conditions for using it is present. The correct conditions for use of using rolling without engine load are given by certain values in the other independent variables. Therefore the direct effect from the multivariate model is the correct effect. We can by use of that model quantify the effect of rolling without engine load: About 0,2 litre per 10 km if the amount of driving time spent rolling without engine load increases by 10 percentage points.

Conclusions

We have divided the independent variables into two main groups. The first group consists of what we call structural variables, they measure the effect of infrastructure, terrain and landscape with steep hill climbing and narrow, winding roads. This group consists of the independent variables relative amount of time spent driving with an engine load above 90% of maximum torque, average speed and relative amount of time spent driving in the highest gear. Rolling without any engine load can also be said to be part of this group. We argue that drivers cannot choose values for these independent variables voluntarily, choices are restricted by external structural factors as described above.

The second group consists of independent variables which measure driving behaviour. These are relative amount of time running idle, relative amount of time spent using cruise control, use of automatic gear shifts and number of brake applications per 100 km.

We believe that the effects in the multivariate models show that structural variables have a greater impact on fuel consumption than indicators for driving behaviour. Two of the three most important independent variables (measured by elasticities) come from the group of structural variables. We therefore conclude that the structural variables are the most important ones for increased fuel savings and reduced emissions of CO_2. The most important measures for reducing fuel consumption and emissions from goods transport with large vehicles would therefore be better infrastructure and reduction of steep hill climbing on narrow, winding roads.

On the other hand, use of automatic gear shift have a large impact on fuel consumption. The effect is larger than the one of average speed. Therefore, driving behaviour is not irrelevant. Fuel consumption is not indifferent to different driving behaviours.

The model has some deficiencies. The weight load indicator include the weight of the vehicle. Ideally this indicator should only measure the weight of the freight load. We also only have the amount of time spent driving with what is defined as a high weight load including the vehicles' own weight. It would be better to have freight weight in tonnes. We did a separate analysis of based on freight weight given by the freight company where we also controlled for the route vehicles are travelling. This analysis show a much higher effect of freight weight. We therefore believe that the weight effect based on data from Dynafleet is underestimated.

Literature

Newsom, J.T.: *Post Hoc Tests*, http://www.upa.pdx.edu/IOA/newsom/da1/ho_posthoc.pdf

Stevens, J.J.: *Post Hoc Tests in Anova*, http://pages.uoregon.edu/stevensj/posthoc.pdf

Toutain, J.E.W, Taarneby, G., Selvig, E.: *Energiforbruk og utslipp til luft fra innenlandsk transport*, Statistisk Sentralbyrå, Rapport 2008/49, http://www.ssb.no/emner/01/03/10/rapp_200849/rapp_200849.pdf,

United States Environment Protection Agency: *Emission Facts: Average Carbon Dioxide Emissions Resulting From Gasoline and Diesel Fuel*, http://www.epa.gov/oms/climate/420f05001.htm

Appendix A: Deleted observations

The following observations in Dynafleet for vehicle A were deleted because of unreasonable data values:

Date	Time	Distance km	Fuel consumption litre	Fuel consumption pr 10 km	Average speed	Route
03.01.2012	00:21	397,8	3,5	0,1	1499,4	No tracking report
19.10.2011	08:56	1313,2	212,5	1,6	151,1	Lidköping-Kaupanger
15.09.2011	11:19	1163,9	233,0	2,0	109,5	Bjøberg-Lidköping
29.09.2011	03:24	415,1	104,5	2,5	127,2	Indre Arna - Gudvangen
28.07.2011	08:32	962,5	225,0	2,3	118,8	No tracking report
11.08.2011	06:21	638,7	180,0	2,8	106,1	No tracking report
01.12.2011	08:53	755,0	189,5	2,5	103,1	Kaupanger-Vang
17.11.2011	05:05	464,9	136,0	2,9	93,5	Kyrkjebø-Fodnes
04.01.2012	09:20	775,6	242,5	3,1	85,1	Ålesund-Bromma
10.05.2011	10:36	769,6	262,5	3,4	83,0	No tracking report

For the last two observations were deleted because we believe an average speed above 82 km pr hour is not achievable on roads in Western Norway. For the last observations, May 10th 2011, there is no tracking report on vehicle A so we have no driving route to assess the driving speed against. We chose by discretion to delete this observation.

Appendix B: Freight data

Date	Vehicle	To	Fuel consumption litre pr 10 km	Freight load in tonnes
02.12.2011	H	Eastern Norway	6,246	27,68
09.12.2011	H	Eastern Norway	5,53	12,26
16.12.2011	H	Eastern Norway	6,228	27,45
20.12.2011	H	Møre og Romsdal	6,996	30,29
21.12.2011	H	Eastern Norway	5,838	27,39
27.12.2011	H	Eastern Norway	6,326	27,46
30.12.2011	H	Eastern Norway	5,732	27,52
01.12.2011	E	Western Norway	5,654	28,92
02.12.2011	E	Eastern Norway	4,725	22,87
06.12.2011	E	Trøndelag	4,98	27,81
13.12.2011	E	Møre og Romsdal	4,766	16,3
14.12.2011	E	Eastern Norway	4,93	27,87
19.12.2011	E	Eastern Norway	5,127	29,32
23.12.2011	E	Trøndelag	5,287	29,68
30.12.2011	E	Southern Norway	4,957	27,13
02.12.2011	F	Southern Norway	4,804	26,74
06.12.2011	F	Western Norway	5,562	29,81
07.12.2011	F	Eastern Norway	5,128	28,32
13.12.2011	F	Eastern Norway	4,75	17,47
20.12.2011	F	Eastern Norway	5,888	28,9
27.12.2011	F	Western Norway	4,717	17,26
30.12.2011	F	Eastern Norway	5,136	15,36
01.12.2011	D	Eastern Norway	5,887	29,49
05.12.2011	D	Eastern Norway	5,443	27,96
07.12.2011	D	Eastern Norway	5,622	30,42
12.12.2011	D	Western Norway	6,196	28,54
14.12.2011	D	Southern Norway	5,565	30,46
21.12.2011	D	Southern Norway	5,231	29,73
23.12.2011	D	Western Norway	5,066	14,59
27.12.2011	D	Eastern Norway	5,553	28,08
02.12.2011	G	Southern Norway	6,117	28,4
07.12.2011	G	Southern Norway	5,873	28,51
09.12.2011	G	Eastern Norway	6,184	26,72
13.12.2011	G	Møre og Romsdal	7,098	28,34
14.12.2011	G	Eastern Norway	5,933	28,3
19.12.2011	G	Western Norway	6,363	29,83
20.12.2011	G	Western Norway	6,175	29,68
21.12.2011	G	Western Norway	6,381	27,4
28.12.2011	G	Southern Norway	5,395	28,27
30.12.2011	G	Trøndelag	5,853	28,34

07.12.2011	C	Southern Norway	5,78	27,99
09.12.2011	C	Eastern Norway	5,334	29,22
14.12.2011	C	Trøndelag	5,72	27,71
16.12.2011	C	Eastern Norway	5,189	25,05
20.12.2011	C	Eastern Norway	5,36	28,66
23.12.2011	C	Western Norway	6,183	29,74
28.12.2011	C	Southern Norway	5,36	29,35
09.12.2011	B	Eastern Norway	5,789	27,44
13.12.2011	B	Western Norway	5,58	17,9
15.12.2011	B	Eastern Norway	5,15	26,85
02.12.2011	I	Eastern Norway	6,35	29,36
06.12.2011	I	Eastern Norway	6,19	28,11
09.12.2011	I	Eastern Norway	6,048	27,68
13.12.2011	I	Eastern Norway	6,35	28,87
15.12.2011	I	Eastern Norway	5,634	28,15
16.12.2011	I	Eastern Norway	5,68	28,56
20.12.2011	I	Møre og Romsdal	6,68	31,02
21.12.2011	I	Eastern Norway	5,49	27,82
23.12.2011	I	Eastern Norway	5,803	10,29
28.12.2011	I	Trøndelag	5,668	27,8
30.12.2011	I	Eastern Norway	6,154	28,43
06.12.2011	P	Eastern Norway	6,125	26,89
13.12.2011	P	Western Norway	6,536	28,47
14.12.2011	P	Eastern Norway	6,136	28,97
16.12.2011	P	Western Norway	5,26	19,19
19.12.2011	P	Eastern Norway	6,137	29,62
21.12.2011	P	Eastern Norway	6,312	29,13
23.12.2011	P	Eastern Norway	5,783	26,38
28.12.2011	P	Eastern Norway	6,153	27,15
30.12.2011	P	Eastern Norway	6,141	28,41
01.12.2011	J	Møre og Romsdal	5,367	17,54
09.12.2011	J	Western Norway	4,174	7,14
13.12.2011	J	Eastern Norway	5,78	28,39
15.12.2011	J	Western Norway	5,151	19,15
19.12.2011	J	Western Norway	5,86	28,2
21.12.2011	J	Western Norway	5,556	25,65
29.12.2011	J	Western Norway	5,709	28,81
30.12.2011	J	Eastern Norway	5,612	28,48
05.12.2011	A	Western Norway	5,936	27,5
07.12.2011	A	Eastern Norway	4,802	23
09.12.2011	A	Trøndelag	5,363	27,56
14.12.2011	A	Western Norway	5,563	25,64
16.12.2011	A	Trøndelag	5,751	30
20.12.2011	A	Møre og Romsdal	5,653	29,52
22.12.2011	A	Eastern Norway	5,291	27,06
27.12.2011	A	Eastern Norway	5,434	29,36

30.12.2011	A	Eastern Norway	5,111	25,09