On the distribution of individual daily driving distances

Patrick Plötz [a],[*], Niklas Jakobsson [b], Frances Sprei [b]

[a] *Fraunhofer Institute for Systems and Innovation Research ISI, Breslauer Str. 48, 76139 Karlsruhe, Germany*
[b] *Chalmers University of Technology, Energy and Environment, 412 96 Göteborg, Sweden*

A B S T R A C T

Plug-in electric vehicles (PEV) can reduce greenhouse gas emissions. However, the utility of PEVs, as well as reduction of emissions is highly dependent on daily vehicle kilometres travelled (VKT). Further, the daily VKT by individual passenger cars vary strongly between days. A common method to analyse individual daily VKT is to fit distribution functions and to further analyse these fits. However, several distributions for individual daily VKT have been discussed in the literature without conclusive decision on the best distribution. Here we analyse three two-parameter distribution functions for the variation in daily VKT with four sets of travel data covering a total of 190,000 driving days and 9.5 million VKT. Specifically, we look at overall performance of the distributions on the data using four goodness of fit measures, as well as the consequence of choosing one distribution over the others for two common PEV applications: the days requiring adaptation for battery electric vehicles and the utility factor for plug-in hybrid electric vehicles. We find the Weibull distribution to fit most vehicles well but not all and at the same time yielding good predictions for PEV related attributes. Furthermore, the choice of distribution impacts PEV usage factors. Here, the Weibull distribution yields reliable estimates for electric vehicle applications whereas the log-normal distribution yields more conservative estimates for PEV usage factors. Our results help to guide the choice of distribution for a specific research question utilising driving data and provide a methodological advancement in the application of distribution functions to longitudinal driving data.

© 2017 The Authors. Published by Elsevier Ltd.
This is an open access article under the CC BY-NC-ND license.
(http://creativecommons.org/licenses/by-nc-nd/4.0/)

1. Introduction

Plug-in electric vehicles (PEVs) charged with renewable electricity are a possible way of reducing greenhouse gas emissions from the transport sector without abandoning individual car-based mobility (Chan, 2007). But the limited electric driving range of battery electric vehicles is a major hurdle for many consumers and the electric range of hybrid PEVs strongly impacts the PEVs utility (Plötz et al., 2014). This limited electric driving range has brought more attention to the distribution of individual vehicle kilometres travelled (VKT) (Greene, 1985; Pearre et al., 2011; Smith et al., 2011; Lin et al., 2012; Tamor et al., 2013). Several studies choose specific distribution functions for analysing and modelling driving vehicle usage, but the choice of a distribution and its consequences have not yet been fully understood nor systematically analysed.

Table 1
Overview of data sets.

Name of data set	German Mobility Panel	Swedish data	Winnipeg data	Settle data
Location	Germany	Sweden	Canada	USA
Collection method	Questionnaire	GPS	GPS	GPS
Sample size	6339	429	72	420
Avg. observation period	7 days	58 days	216 days	251 days

Several studies provide evidence that individual-longitudinal and cross-sectional VKT distributions are peaked and right skewed such as the Weibull, log-normal and Gamma distribution. Greene (1985) and Lin et al. (2012) analyse two sets of data and argue that the Gamma distribution is most suitable (Greene, 1985; Lin et al., 2012). However, Blum (2014) and Plötz et al. (2012) argue that the log-normal distribution provides the best fit for most drivers and for all daily VKT. Pearre et al. (2011) empirically analyse days with long-distance driving and find the travel pattern of individual vehicles does not resemble the average fleet travel pattern. Smith et al. (2011) use individual vehicle travel data to construct an average commuter driving cycle. Finally, Tamor et al. (2013) use a mixture of a normal and exponential distributions with five free parameters to model vehicle specific daily VKT distributions.

Overall, the evidence for the best two-parameter distribution for daily VKT is not conclusive. Furthermore, the word 'best' is in this context highly application-dependent. A certain distribution may be performing overall better than another according to some goodness of fit measures, but be worse when it comes to predicting short or long daily driving distances. This is especially important in the context of PEVs where research often focuses on the utility factor (UF) of PHEVs, see e.g. Gonder et al. (2007), Millo et al. (2014), Silva et al. (2009), Smart et al. (2014), or the days with long-distance travel, i.e. days requiring adaptation (DRA), for BEVs, e.g. Jakobsson et al. (2016), Tamor and Milacic (2015), Greene (1985), Pearre et al. (2011), Smith et al. (2011), Lin et al. (2012), and Tamor et al. (2013). The UF of PHEVs depends on the short daily driving distances, while the DRA depends on the long daily driving distances; in both of these cases, the choice of distribution impacts the results obtained.

The aim of the present paper is to provide a systematic comparison of the choice of distribution function for individual daily VKT with respect to (1) goodness of fit and (2) predictive power of DRA for BEVs and UF for PHEVs. We use four different data sets, with various complementary properties, to analyse the three two-parameter probability distributions with respect to daily driving data that received most attention in the literature. The three distributions are log-normal, Weibull and Gamma. We use three GPS measured data sets from Western Sweden; Winnipeg, Canada; and Seattle, USA; respectively. The fourth data set is survey based and from Germany. The data sets differ in sample size and measurement length, which provide robustness to our results. Furthermore, we analyse the effect of measurement length on the stability of goodness of fit, and overall best distribution.

The present paper differs from previous work in several aspects. First, we test the assumption of independently and identical distributed (iid) observations underlying the use of distribution functions for daily mileages. Second, we perform a systematic comparison of several data sets. Third, we analyse the effect of observation period on the goodness of fit for the distribution functions. Fourth, we compare the consequences of distribution choice in applications related to PEV.

2. Data and methods

2.1. Data

We use four data sets to analyse the goodness of fit of different distributions. The data sets comprise vehicle motion from Germany (MOP, 2010), Sweden (Karlsson, 2013), Canada (Smith et al., 2011) and the USA (PSRC, 2008; Transportation Secure Data Center, 2015). The average observation periods range from 7 to more than 200 days. The different data sets are summarised in Table 1. A description of the data sets follows below. More detailed summary statistics are given in Table 2.

The German Mobility Panel (MOP, 2010) is one of two national household travel surveys in Germany. Since MOP is a household travel survey which focuses on people and their trips, we assigned trips to vehicles if unambiguously possible (see Kley, 2011 and Plötz et al., 2014 for details). By using all data from 1994 until 2010, we obtain 6339 vehicle driving profiles with 172,978 trips in total. Apart from driving, the profiles contain socio-economic information about the driver (e.g. age, sex, occupation, household income, education) and the vehicle (e.g. size, owner, garage availability). This data set is representative for German driving in terms of daily and annual mileage, vehicle size and garage ownership (Gnann, 2015).

The Swedish Car Movement Data (SCMD) consists of GPS measurements of more than 700 privately driven cars in the provinces of Västra Götaland and Kungsbacka in Western Sweden. Of these, we have selected 429 cars that have at least 30 days of good GPS measurements, whereas the rest have less than 30 days and are not included in the analysis (for details see Björnsson and Karlsson (2015)). Measurements were evenly distributed over the years 2010–2012. The cars were randomly sampled from the Swedish vehicle registry with an age restriction on the car of maximum 8 years, the positive response rate of the selected households was 5%. Western Sweden is representative for Sweden in terms of urban and rural areas, city sizes and population density. The sample is representative in terms of car size and car fuel type. There is a slight overrepresentation of measured cars having higher annual VKT cars in the households compared to the national average due

Table 2
Summary statistics of the data sets.

	0.25-quantile	Median	Mean	0.75-quantile
Swedish data (N = 429)				
Observation period [days]	51	59	58	64
Share of driving days	0.67	0.83	0.8	0.96
Average daily VKT [km]	38.4	51.9	57.1	72.3
German Mobility Panel data (N = 6339)				
Observation period [days]	Seven for all vehicles by design			
Share of driving days	6/7	1	0.92	1
Average daily VKT [km]	22	28.3	50.6	65
Winnipeg data (N = 75)				
Observation period [days]	108	238	216	325
Share of driving days	0.58	0.67	0.67	0.81
Average daily VKT [km]	22.5	28.9	33.2	39.3
Seattle data (N = 420)				
Observation period [days]	270	276	260	276
Share of driving days	0.73	0.85	0.80	0.92
Average daily VKT [km]	35.4	48.2	50.3	61.8

to the age criterion in the sampling. With regard to the driver's age, there is a slight overrepresentation of senior citizens. A full description of the data including pre-processing is available in Karlsson (2013).

The PSRC traffic choice study data contains measurements of 484 cars for up to 18 months, the cars come from households in the Seattle Metropolitan Area and the data was collected between 2004 and 2006 (PSRC, 2008). The data was originally collected as part of a study on congestion charges where part of the measurement period included an artificial toll for driving on certain roads. The first 7 months of measurement is the pre-toll period, whereas the last 11 months is the toll-period of the measurement. The application of the artificial toll changes the driving behaviour of the measured cars, and to avoid artefacts from the implementation of the toll, we have chosen to only include the last 11 months (the toll-period) in our analysis. These 11 months contain measurements of 437 cars. A further filtering where we only include cars with at least 30 measurement days and a removal of one car with unreasonably low annual VKT yield a total of 420 cars included in our analysis.

The Winnipeg data set contains seventy-five vehicles' GPS data recorded from volunteers in or around Winnipeg, Canada. Smith et al. (2011) collected and pre-processed the data as well as made it available for download. All GPS data were recorded to the split second and data gaps were filled by Smith et al. (2011). Individual trips were joined "if the gap between ending and starting times was less than or equal to 120 s" (see Smith et al., 2011 for further details). In addition, the individual trips were aggregated to daily VKT for the present paper.

The four data sets complement each other; specifically, the German data has a very large sample and short observation period while the Winnipeg data has a long measurement period and small sample size, the Swedish data is in-between in both categories, and the Seattle data has a moderate sample size and long measurement period. Furthermore, the data sets are from different countries and geographical settings, as well as having measurements spread out over the year. In total, the four data sets cover a total of 190,000 driving days and 9.5 million VKT.

2.2. Methods

We analyse three right-skewed two-parameter distributions. Following Lin et al. (2012), our chosen distributions are the log-normal $f(r) = \exp[-(\ln(r) - \mu)^2 / (2\sigma^2)] / (r\sqrt{2\pi}\sigma)$, Gamma $f(r) = r^{k-1}\exp[-r/\theta] / (\Gamma(k)\theta^k)$, and Weibull $f(r) = (k/\lambda)(r/\lambda)^{k-1}\exp[-(r/\lambda)^k]$ distribution where $f(r)$ denotes the probability density function (PDF). They are widely used in modelling right-skewed data (Hahn and Shapiro, 1967; Kundu et al., 2005). All three distribution functions assign zero probability to a trip of zero km length and fall off slower than a normal distribution for long distances.

Many statistical methods rely on the assumption of independently and identical distributed (iid) observations. High autocorrelation between the mileages on several days of an individual could bias the results of standard statistical methods if they rely on the iid assumption. We use Ljung–Box tests to analyse the autocorrelation and use linear regression of lagged daily mileages on daily mileages to study the extent of autocorrelation present in the individual driving data (Brockwell and Davis, 2006).

The statistical tests of the zero autocorrelation assumption are performed for a large number of drivers with the threat of a noteworthy number of false positives. The number of false positives grows with the number of tests and the individual p-values for the null hypothesis have to be analysed with care (Storey and Tibshirani, 2003). Statistical methods have been developed to estimate the fraction π_0 of hypothesis that are null, i.e. in our case where the null hypotheses of zero autocorrelation cannot be rejected. We apply the methodology of Storey (2002) that estimates the share of false discovery rates ("the expected proportion of false positive findings among all the rejected hypotheses" (Storey, 2002)) below a given thresh-

Table 3
Autocorrelation of daily mileages in the four data sets.

	Winnipeg	Mobility Panel	Sweden	Seattle
Average coefficient of determination $\overline{R^2}$	0.069	0.167	0.045	0.04
Share of drivers with $R^2 > 0.3$	3%	19%	2%	1%
Average ACF(1)	0.14	-0.07	0.06	0.12
Fraction of tests where the null hypotheses of zero autocorrelation cannot be rejected π_0	58%	90%	92%	39%

old. As a result of this method, we obtain this share π_0 of null hypothesis that cannot be rejected. A large π_0 means that zero autocorrelation is a good approximation for a given data set. We use the q-value package for the R statistical software (Bass et al., 2015).

The match between the observed and expected distribution of daily mileages is assessed by four goodness of fit measures:

(1) the Akaike information criterion (AIC) is a penalised log-likelihood $AIC = -2\,LL + 2\,(p+1)$, where p is the number of the model parameters and LL the log-likelihood;
(2) the root mean squared error $RMSE = \sum_i (y_i - \hat{y}_i)^2 \ /n$,
(3) the mean average percentage error $MAPE = \sum_i |(y_i - \hat{y}_i)/\hat{y}_i| \ /n$, and
(4) the $\chi 2$ statistic $\chi^2 = \sum_i (y_i - \hat{y}_i)^2 \ /\hat{y}_i$

where n is the number of driving days, y_i the empirical cumulative density function (CDF) and $\hat{y}_i$ the predicted CDF value at r_i. We calculate the goodness of fit according to these measures for each distribution, driver and data set.

In application of distribution functions with respect to PEV, not the distribution function as such is of interest but its ability to predict relevant quantities. For example, an understanding of the distribution of daily VKT allows us to estimate the probability of rare long-distance car travel (Plötz, 2014). If $f(r)$ is the individual PDF of daily VKT, the probability of driving more than L km on a driving day is given by $\int_L^\infty f(s)ds = 1 - F(L)$ where $F(r)$ is the CDF of $f(r)$. A simple measure for the reduction in utility of a PEV is given by the number of days per year $D(L)$ with more than L km of driving: $D(L) = 365\,(n/N)\,[1 - F(r)]$ if the vehicle is used on n days out of N observation days. Thus, $D(L)$ is the number of days requiring some form of adaptation due to the range limitation for a potential BEV user. Please note that the different distribution functions have noteworthy consequences for the probability of rare long-distance car travel: The log-normal distribution falls off slower than the Weibull and Gamma distribution for large distances and makes many days with long car trips more likely. For PHEV, the UF measures the share of kilometres driven electrically in the total distance driven. We calculate the DRA and the UF using the three distributions for each vehicle in each data set for three different range limitations, and compare these results with each other, as well as to those of an extrapolation of the number of DRA and the UF factor.

All goodness of fit test and simulation have been implemented in the Matlab software. Tests for zero autocorrelation and for the identification of potential false positives were performed with the R statistical software (R Core Team, 2016; Bass et al., 2015; Wickham, 2009).

3. Results

3.1. Independence of daily mileages

The usage of distribution function for daily VKT as found in the literature assumes that observations of daily VKT are identically and independently distributed (iid). This assumption has to the best of our knowledge not been tested so far in the literature. In the following we test this assumption by quantification of the autocorrelation in daily VKT of many drivers.

We perform a linear regression of the driving distances of day t on day $t-1$. If daily mileages on subsequent days were strongly correlated one should observe high R^2 coefficients of determination for many users. Fig. 1 shows the histograms of R^2 from the linear regression of lagged daily VKT on daily VKT for all four data sets. We observe very low correlation between the daily mileages on subsequent days for most drivers in all data sets. Indeed, the R^2 of this linear regression is larger than 0.3 for only 3% of the Winnipeg drivers, 19% of the German drivers, 2% of the Seattle and 1% of the Swedish drivers. Furthermore, the average R^2 for the same data sets read 0.069, 0.167, 0.045 and 0.04 respectively. The larger values for the German Mobility Panel data set are a consequence of the few observation days (for only six lags a linear fit can reduce the variance substantially which is not to easily possible for larger samples) and the rounding errors from reporting. These results are summarised in Table 3.

The autocorrelation function $ACF(l) = Cov(x_t, x_{t-l})/Var(x_t)$ at lag $l = 1$ quantifies the autocorrelation between daily mileages on subsequent days. Fig. 2 shows the histograms of ACFs for lag one for all four data sets. Consistent with the linear regression results, most ACFs are close to zero. The absolute autocorrelation deviates from zero more strongly for the Mobility Panel data. However, this is a consequence of rounding errors since daily mileages are reported by users in the Mobility Panel data as opposed to measured daily mileages in the other data sets and also enhanced by the short observation period.

The limited observation periods raise the question whether the observed autocorrelation is significantly different from zero. We perform Ljung–Box tests for the null hypothesis of no autocorrelation ACF(1) between daily VKT for each individual
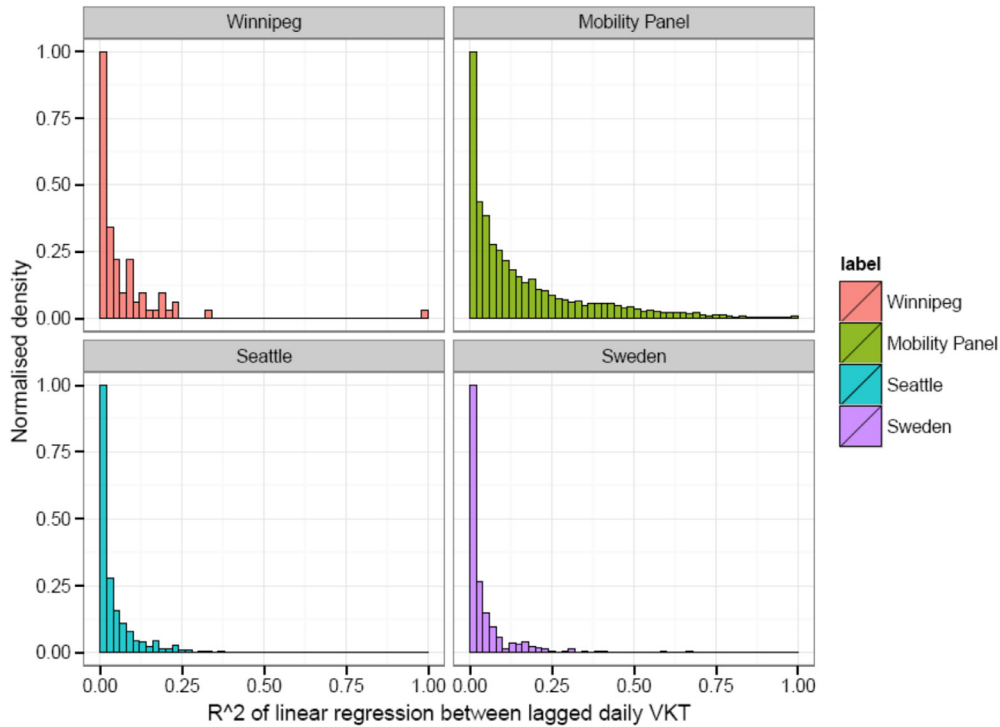
Fig. 1. Correlation between daily mileages. Shown are histograms (normalised to maximal count) of the R^2 measures for regressions of lagged daily mileage $\mathbf{VKT}_{t-1}$ on current daily mileage $\mathbf{VKT}_t$ for the different data sets.

Table 4
Summary of goodness of fit statistics. The best distribution for most users in bold face.

Goodness of fit	German Mobility Panel			Winnipeg			Sweden			Seattle		
	ln N	Weib.	Γ	ln N	Weib.	Γ	ln N	Weib.	Γ	ln N	Weib.	Γ
AIC	32,3%	**59,9%**	7,8%	**40%**	35%	25%	34,7%	**39,2%**	26,1%	14,8%	**72,9%**	12,4%
RMSE	**74,0%**	12,5%	13,5%	**36%**	35%	29%	**43,6%**	37,3%	19,1%	35,7%	**36,4%**	27,9%
χ^2	**88,2%**	9,1%	2,7%	17%	**41%**	**41%**	29,6%	33,8%	**36,6%**	**68,3%**	18,3%	13,3%
MAPE	**75,2%**	21,9%	2,9%	9%	**51%**	40%	19,8%	**44,5%**	35,7%	**64,3%**	24,3%	11,4%

driver. Since the number of drivers in the different samples varies between 75 and 6339 (see Section 2.1), this is a multiple testing problem. The fraction of tests where the null hypotheses cannot be rejected, i.e. where zero autocorrelation is a good approximation, is 58% for the Winnipeg data, 90% for the Mobility Panel, 92% for the Swedish data and 39% for the Seattle data. These results are summarised in Table 3.

We quantified the autocorrelation between daily mileages on subsequent days. The independence of daily mileages on individual days is a pre-requisite for applying statistical distribution functions and has to our knowledge not been tested previously. We find the autocorrelation to be not significantly different from zero in many cases and—when quantified— to be limited in most cases. Thus, the daily mileages of individual users can be treated as iid to a good approximation.

3.2. Best overall distribution

3.2.1. Goodness of fit statistics

For each individual vehicle of the different data sets all three distribution functions have been fitted for the daily VKT using maximum likelihood estimates (MLEs). Table 4 summarises the goodness of fit results by indicating what share of individual driving patterns were best according to the corresponding goodness of fit measure. This table shows the relative overall performance of the three distributions with respect to each other.

For the Mobility Panel data with only seven days of observation, the log-normal distribution fits most daily VKT best according to three out of four goodness of fit measures. The picture is less clear for the Swedish, Seattle, and Winnipeg
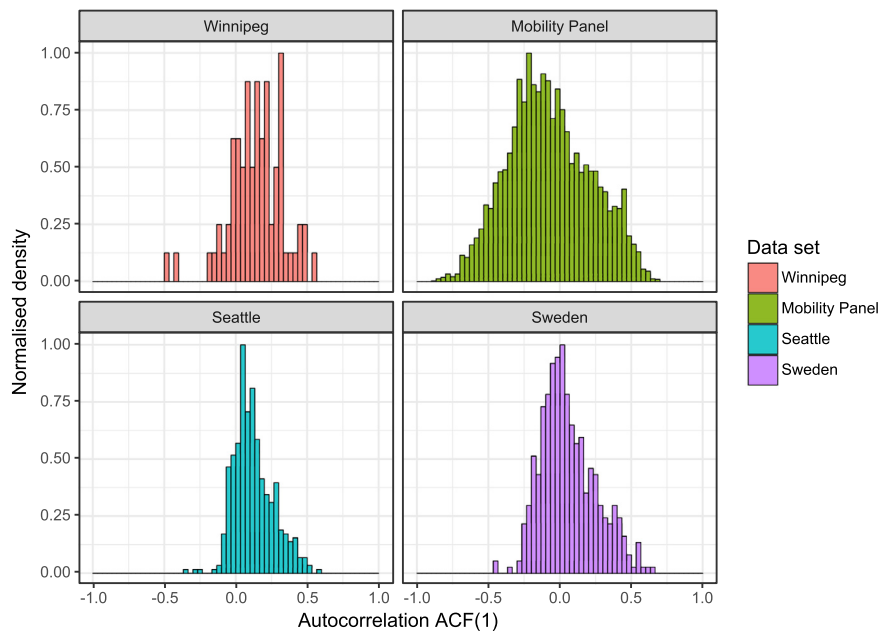
Fig. 2. Autocorrelation between mileages on subsequent days. Shown are histograms (normalised to maximal count) of the ACF(1) of daily mileages for the four data sets.

data: Each distribution is best for most of the driving profiles in at least one measure and data set; however, the Gamma distribution performs worst, as it is the best choice in only 2 out of 16 cases. Contrary to earlier literature, we cannot see that one distribution stands out as clearly better than the others, though both log-normal and Weibull show an overall better fit than Gamma.

There can be many reasons why different distributions fit the different data sets better as shown in Table 4. The driving patterns can differ between the different measurement locations, there can be nuances in data collection that causes differences in the maximum likelihood estimates for the data sets, and so on. One can also note that the observation length varies between the data sets, and it is possible that this influence the results. It turns out that this is the case, as we shall see in the next section.

3.2.2. What influences goodness of fit statistics: varying observation period

By varying the number of driving days included in the maximum likelihood estimate for the longer data sets we directly observe the effect of measurement period on goodness of fit for the three distributions. Fig. 3 shows the share of individual drivers for which each respective distribution performs best according to the individual goodness of fit measure for the Swedish data. Fig. 4 shows the same for the Seattle data. A bootstrapping algorithm has been used, where the driving days included in the MLE are randomly selected from the measured driving days, which has been averaged over 120 bootstrap samples. We focus on these two data sets in the present section since the observation period is too small in the Mobility Panel data and varies too strongly in the Winnipeg data.

Both data sets are plotted up to the maximum number of days measured (note that the bootstrapping procedure can no longer smooth out the estimates for a large number of days, as this results in that the same days are included in each iteration). In six out of the eight cases, the Gamma distribution performs less well than both the other distributions, in no case does it perform best for a substantial part of the observation period spectrum. With regards to the log-normal and Weibull distributions, which is deemed best depends on both the goodness of fit measure and observation length. For short observation periods, log-normal performs better than Weibull in seven out of eight cases, however for longer observation periods, the picture is less clear. In the Swedish data, log-normal has declining performance, in favour of Weibull, with increasing observation time for all the goodness of fit measures. With the even longer observation period of the Seattle data log-normal achieves increasing performance for most of the observation period spectrum for MAPE, and consistently better performance for RMSD and Chi-squared. For AIC, Weibull performs best.

The best distribution according to the goodness of fit measures changes with observation period both for Swedish and Seattle data. Noteworthy is the change of monotonicity with observation period, e.g. for MAPE and Chi-squared for the Seattle data, where one would mainly expect convergence with increasing observation period but not change of monotonicity. A likely explanation for this change for MAPE and Chi-squared lies in the fact that both measures divide the deviation between observed and expected value by the expected value. For larger number of days, the distributions are sampled in more
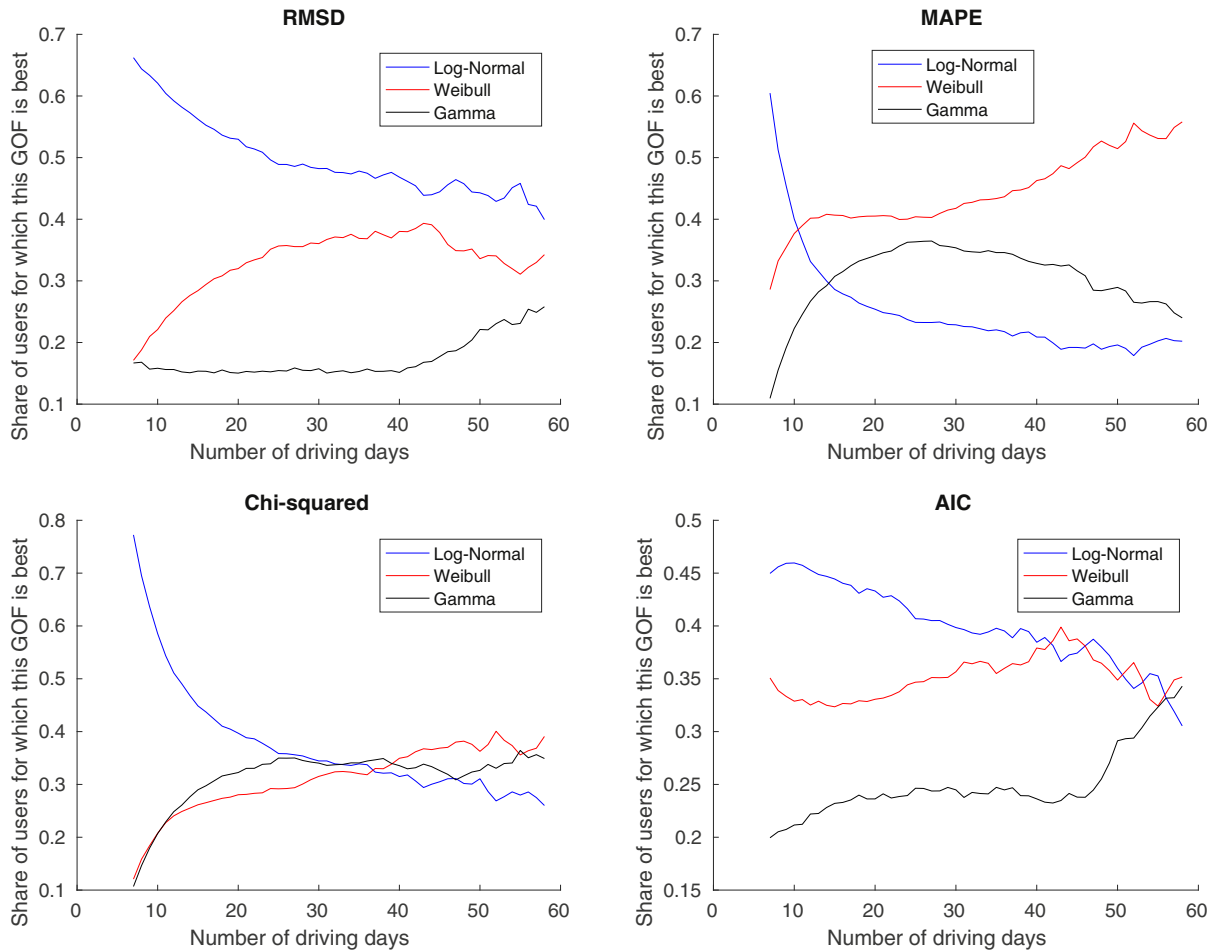
Fig. 3. Share of individual drivers for which the respective distributions perform best with respect to observation length for the Swedish data according to RMSD (top left panel), MAPE (top right), Chi-squared (bottom left), and AIC (bottom right).

extreme regions where the Gamma and Weibull distribution fall off quickly. Thus, an actual observation at large distance is divided by a very small expected probability at large distances leading to an increase in the error measure and preference for the log-normal distribution (numerically, this can lead to drastic instability for more than 150 observation days when the denominator gets close to zero and cannot be distinguished from 10^{-16}).

Variability in performance for the distributions can be due to either the distributions, or the goodness of fit measure. To separate the two effects, we created artificial data from a known probability distribution. In Fig. 5 we show artificially created log-normal data (left panel) and artificially created Weibull data (right panel), and measured the fraction of cases where log-normal performs better than Weibull for the four goodness of fit cases with respect to observation length. The artificial log-normal and Weibull data has been generated from the average Mobility Panel log-normal ($\mu = 3.34$, $\sigma = 0.86$) and Weibull $\tau = 42.6$, $\beta = 1.53$ parameters and the results have been averaged over 1,000 simulation runs to reduce statistical noise. Of the four goodness of fit measures, AIC manages to achieve a high rate of correct predictions of log-normal (left panel) while maintaining a low rate of incorrect predictions of log-normal (right panel). Similar simulations have been performed for a comparison between log-normal and Gamma as well as Gamma and Weibull distribution with similar outcomes. Thus, AIC is the preferable measure with respect to changing observation period.

In summary, AIC provides a safe measure to distinguish between the different distribution function after approximately 60 days of observation. Furthermore, no distribution clearly outperforms the others: For the Swedish data, each distribution is better than the others for some users, whereas for the Seattle data, the Weibull distribution is best for most users. In the light of the effect from finite observation period, Table 4 indicates that log-normal and Weibull are good approximations for some users. Possibly, a distribution function with more than three parameters that contains the three two-parameter distribution functions discussed here as limiting cases, could lead to a coherent description of all daily mileages but is beyond the scope of the present paper.

Fig. 4. Share of individual drivers for which the respective distributions perform best w.r.t. observation length for the Seattle data.

3.3. Consequences of distribution choice in prediction of PEV utility

3.3.1. Days with long-distance travel

When considering BEVs, the days within a year with daily mileage larger than the all-electric range, which will be referred to as "days requiring adaptation" (DRA) in the following, is a key measure. In Tables 5 and 6 we present estimates of DRA in the four data sets for the different distributions. For the distributions, the DRA has been calculated according to the method in Section 2.2. For the extrapolated values in Table 6, the DRA has been calculated over the measurement period and extrapolated to 1 year (for normalisation). Table 5 shows the percentage of users with number of DRA < 1, DRA < 12 and DRA > 52 for three range limitations (100, 150, 200 km) as obtained from using the distribution functions and from direct extrapolation of the data (column "extrapolation" – this column is missing for the German Mobility Panel data since the low number of seven observation days does not allow a distinction at the required level of accuracy). Similarly, Table 6 shows the mean and median percentage of DRA for the same range limitations (the extrapolation median has been omitted for the German Mobility Panel data since the extrapolation from seven days of driving to 1 year leads individual DRA of 0, 52, 104, etc. without sufficient resolution). Confidence intervals (95%) have been calculated as Clopper–Pearson approximation in all cases except the median, where it is calculated from BCa bootstrap (with 1000 bootstrap samples).

We note some differences among the four data sets. Firstly, the Winnipeg data has far fewer DRAs than the other data sets, this may be due to Winnipeg's geographical isolation, where a user would have to travel very far to go elsewhere, and this happens seldom. Furthermore, none of the Winnipeg drivers had more than 45 DRA per year for 150 km range. Secondly, though the Swedish and German data have similar average share of DRA, they have them distributed differently. This may be due to the limited observation period of the German data, where the included days could either have a large fraction of long-distance driving or large fraction of short-distance driving. Thirdly, the Swedish and German data have an overall higher share of DRA than both the North American data sets.

Though there are differences among all three distributions, it is clear that log-normal differs more in prediction of share DRA from Weibull and Gamma, than Weibull and Gamma do from each other (especially for mean and median). What should especially be noted, is that log-normal estimates a higher fraction of DRAs than Weibull and Gamma. Consider, for example, the share of users with DRA < 1 for a range of 150 km in the Swedish data. Log-normal predicts 7.9% of the users

Table 5
Share of users in percent with adaptation needs below 1 DRA, 12 DRA and above 52 DRA per year, for different the distributions, ranges and data sets and compared to direct extrapolation from the data.

	L [km]	Germany: Mobility Panel			Sweden				Winnipeg				Seattle Puget sound			
		ln N	Weib.	Γ	ln N	Weib.	Γ	Extrapolation	ln N	Weib.	Γ	Extrapolation	ln N	Weib.	Γ	Extrapolation
DRA< 1 [%]	100	17.3 ± 0.9	36.9 ± 1.2	31.3 ± 1.1	3 ± 1.6	9.1 ± 2.7	7.2 ± 2.5	11.7 ± 3.0	3 ± 2	36 ± 5	32 ± 5	28 ± 5	0.7 ± 0.8	8.1 ± 2.6	6.9 ± 2.4	5.5 ± 2.2
	150	28.5 ± 1.0	53.0 ± 1.2	48.0 ± 1.2	7.9 ± 2.6	21.2 ± 3.9	17.5 ± 3.6	25.9 ± 4.1	19 ± 4	61 ± 6	61 ± 6	47 ± 6	3.6 ± 1.8	28.1 ± 4.3	22.4 ± 4	15.5 ± 3.5
	200	37.0 ± 1.2	64.2 ± 1.2	62.3 ± 1.2	11.4 ± 3	34.7 ± 4.5	30.8 ± 4.4	42.2 ± 4.7	36 ± 6	81 ± 4	83 ± 4	64 ± 6	10.5 ± 2.9	51.2 ± 4.8	46.9 ± 4.8	26.7 ± 4.2
DRA< 12 [%]	100	37.4 ± 1.2	51.7 ± 1.2	48.8 ± 1.2	15.6 ± 3.4	24 ± 4	22.4 ± 3.9	24.7 ± 4.1	59 ± 6	72 ± 5	73 ± 5	72 ± 5	16.2 ± 3.8	30.0 ± 5.0	28.8 ± 5.0	35 ± 4.6
	150	53.6 ± 1.2	68.9 ± 1.1	67.0 ± 1.1	30.1 ± 4.3	51.1 ± 4.7	49.9 ± 4.7	49.2 ± 4.7	87 ± 4	92 ± 3	92 ± 3	89 ± 4	39.1 ± 5.9	63.1 ± 6.4	65.0 ± 6.8	68.6 ± 4.4
	200	65.3 ± 1.2	78.8 ± 1.0	77.5 ± 1.0	40.1 ± 4.6	69.5 ± 4.4	69.5 ± 4.4	64.3 ± 4.5	89 ± 4	99 ± 1	97 ± 2	95 ± 2	61.4 ± 7.1	85.7 ± 6.2	85.7 ± 6.5	82.9 ± 3.6
DRA> 52 [%]	100	32.1 ± 1.1	29.4 ± 1.1	30.1 ± 1.1	45.9 ± 4.7	38.5 ± 4.6	39.2 ± 4.6	31.2 ± 4.4	4 ± 2	4 ± 2	4 ± 2	5 ± 2	31.0 ± 5.3	27.4 ± 4.8	26.2 ± 4.7	15 ± 3.4
	150	17.9 ± 0.9	15.6 ± 0.9	16.1 ± 0.9	18.4 ± 3.7	13.3 ± 3.2	13.3 ± 3.2	11.9 ± 3.1	0	0	0	0	8.8 ± 2.8	5.0 ± 1.8	4.1 ± 1.7	2.9 ± 1.6
	200	10.4 ± 0.8	8.6 ± 0.8	9.2 ± 0.8	11.4 ± 3	4.9 ± 2	5.3 ± 2.1	5.1 ± 2.1	0	0	0	0	1.7 ± 1.2	1.0 ± 0.7	1.0 ± 0.7	0.7 ± 0.8

Table 6

Mean and median adaptation needs according to different distributions, ranges and data sets and compared to direct extrapolation from the data.

	L [km]	Mobility Panel				Sweden				Winnipeg				Puget sound			
		ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation
Mean DRA	100	45.8 ± 1.4	43.1 ± 1.5	44.3 ± 1.5	42.6 ± 1.6	52.7 ± 3.8	48.4 ± 4.3	48.9 ± 4.3	45.3 ± 4.7	16.1 ± 4.2	11.3 ± 4.8	11.1 ± 4.5	12 ± 5	41.6 ± 2.9	37.0 ± 3.4	36.4 ± 3.3	28.9 ± 3.1
	150	25.5 ± 0.9	21.1 ± 1.0	22.2 ± 1.0	22.0 ± 1.1	32.4 ± 2.7	22.6 ± 2.7	23.2 ± 2.7	22.4 ± 2.7	7.2 ± 2.1	3.2 ± 1.8	3.2 ± 1.8	5 ± 2	21.4 ± 1.8	13.2 ± 1.8	13.1 ± 1.7	12.0 ± 1.7
	200	16.7 ± 0.7	12.3 ± 0.7	13.1 ± 0.7	13.7 ± 0.8	22.4 ± 2.1	11.9 ± 1.8	12.1 ± 1.8	12.8 ± 1.8	3.8 ± 1.2	1.0 ± 0.7	1.1 ± 0.7	2 ± 1	12.7 ± 1.3	5.2 ± 0.9	5.3 ± 0.9	6.6 ± 0.8
Median DRA	100	26.6 ± 1.6	10.2 ± 1.3	13.1 ± 1.5	–	49.6 ± 4.2	35.4 ± 6.1	37.5 ± 6.8	29.0 ± 3.3	10.3 ± 3.4	4.0 ± 2.3	4.5 ± 2.5	4 ± 3	36.4 ± 3.9	27.3 ± 4.6	27.1 ± 3.8	18.0 ± 2,6
	150	9.4 ± 0.7	0.4 ± 0.1	1.3 ± 0.2	–	27.3 ± 3.5	11.2 ± 2.7	12.2 ± 2.7	12.2 ± 2.1	4.2 ± 1.6	0.2 ± 0.4	0.5 ± 0.5	1 ± 2	16.5 ± 2.1	4.8 ± 1.4	5.9 ± 1.3	6.6 ± 1.3
	200	4.1 ± 0.4	0.01 ± 0.01	0.1 ± 0.04	–	17.4 ± 2.0	3.8 ± 1.5	4.1 ± 1.4	6.0 ± 1.0	1.9 ± 0.9	0.01 ± 0.05	0.04 ± 0.05	0 ± 0	8.6 ± 1.4	0.9 ± 0.5	1.2 ± 0.4	4.0 ± 0.7

Fig. 5. Goodness of fit for artificial log-normal (top panel) and Weibull random data (bottom panel) as a function of observation period.

to have so few DRA, while Weibull and Gamma predict 21.2% and 17.5% respectively. Thus the choice of distribution has a large impact on results when considering DRA. If one wishes to have a conservative estimate of the number of users who would fulfil their driving with a BEV, one might choose to model driving data with the log-normal distribution. However, since the other distributions and the empirical calculation give similar results there is an indication that these distributions may give a more accurate estimate of what the real number of DRA is for a user.

In summary, the extrapolated values for mean and median of the DRA in Table 6 agree quite well with the estimate from the Weibull and Gamma distributions. In comparison, the log-normal distribution is more conservative and estimates higher average DRA.

3.3.2. PHEV utility factor

The utility factor (UF), i.e. the share of kilometres driven electrically, is an important measure in assessment of PHEVs. It measures the distance driven on electricity divided by the total distance driven by the car. Here we have calculated the UF by simulating 50,000 driving days from each driver based on the distribution parameters obtained from the MLE described in the Methods section. These distribution based UFs are than compared to a direct UF calculation from the daily VKT (column "extrapolation"). Summary statistics of all UFs are given in Tables 7 and 8. Confidence intervals (95%) are calculated as the Clopper–Pearson intervals in Table 7 and as BCa bootstrap for Table 8. Table 7 shows the share of users with UF above 50% and 80% for three all-electric ranges (25, 50, 75 km), while Table 8 shows the mean and median UF for the same range limitations.

The different data sets give different results on UF where Winnipeg stands out with a higher UF than the other data sets, the overall lowest UF appears in Sweden, which may be due to an inclusion criterion for cars in the Swedish data. In the Swedish data only cars younger than 8 years of age are measured and these cars drive more on an annual basis compared to the average Swedish car. Furthermore, the Seattle data and the German data show quite similar UF in-between the low Swedish UF and high Winnipeg UF.

Table 7

Share of users with UF above 50% and 80% according to different distributions, ranges and data sets and compared to direct extrapolation from the data.

UF	L [km]	Germany Mobility Panel				Sweden				Winnipeg				Seattle Puget sound			
		ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation
UF> 50% [%]	25	42.9 ± 1.0	51.2 ± 0.8	51.0 ± 0.8	29.7 ± 1.1	18.4 ± 0.2	24.9 ± 0.2	24.9 ± 0.2	26.8 ± 0.2	63 ± 7	76 ± 4	79 ± 4	77 ± 9	22.6 ± 0.2	32.1 ± 0.1	33.6 ± 0.1	36.0 ± 0.1
	50	70.9 ± 0.4	80.6 ± 0.3	80.5 ± 0.3	50.6 ± 1.2	46.6 ± 0.0	76.0 ± 0.2	77.6 ± 0.2	76.5 ± 0.2	91 ± 2	99 ± 1	99 ± 1	97 ± 3	72.4 ± 0.2	90.1 ± 0.3	91.7 ± 0.3	91.0 ± 0.3
	75	81.4 ± 0.2	90.3 ± 0.1	90.1 ± 0.1	57.5 ± 1.2	64.3 ± 0.1	92.3 ± 0.3	92.5 ± 0.3	91.1 ± 0.3	97 ± 1	99 ± 1	100 ± 0	99 ± 3	89.1 ± 0.3	98.3 ± 0.4	98.6 ± 0.4	98.3 ± 0.4
UF> 80% [%]	25	14.5 ± 1.3	20.7 ± 1.3	20.6 ± 1.3	10.1 ± 0.7	28.0 ± 0.3	4.6 ± 0.3	4.0 ± 0.3	4.2 ± 0.3	1 ± 9	21 ± 13	21 ± 13	23 ± 9	0.5 ± 0.4	3.3 ± 0.4	3.3 ± 0.4	4.3 ± 0.3
	50	37.2 ± 1.1	50.5 ± 0.9	49.9 ± 0.9	29.2 ± 1.1	12.1 ± 0.3	21.0 ± 0.2	20.3 ± 0.2	20.3 ± 0.2	37 ± 11	67 ± 6	71 ± 5	67 ± 11	13.8 ± 0.3	26.9 ± 0.2	27.9 ± 0.2	30.7 ± 0.1
	75	52.6 ± 0.8	68.9 ± 0.5	68.2 ± 0.5	41.5 ± 1.2	24.0 ± 0.2	48.0 ± 0.1	48.5 ± 0.0	44.1 ± 0.04	70 ± 5	89 ± 2	89 ± 2	81 ± 9	33.3 ± 0.1	64.3 ± 0.1	66.9 ± 0.1	61.7 ± 0.1

Table 8
Mean and median UF according to different distributions, ranges and data sets and compared to direct extrapolation from the data.

UF	L [km]	Mobility Panel				Sweden				Winnipeg				Seattle Puget sound			
		ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation	ln N	Weib.	Γ	Extra-polation
Mean [%]	25	48.0±0.6	54.0±0.6	53.8±0.6	50.2±0.6	33.2±1.9	41.9±1.6	42.2±1.6	42.1±1.6	54±4	62±4	62±4	63±4	39.8±1.4	45.8±1.4	46.3±1.4	47.0±1.5
	50	66.0±0.6	74.6±0.6	74.3±0.6	71.3±0.6	49.3±2.3	64.3±1.7	64.5±1.7	63.6±1.7	74±4	84±3	84±3	83±3	60.0±1.6	70.5±1.4	71.0±1.4	71.3±1.4
	75	74.7±0.6	83.9±0.5	83.7±0.5	81.3±0.6	58.6±2.4	76.9±1.5	77.1±1.5	75.5±1.6	83±3	92±2	93±2	90±2	71.0±1.6	83.5±1.2	83.8±1.2	82.7±1.2
Median [%]	25	44.4±0.8	50.8±1.0	50.9±1.1	46.5±1.0	29.4±0.1	39.1±1.0	39.7±0.3	39.4±0.03	54±4	62±5	63±4	63±5	38.6±0.0	43.8±0.4	44.1±0.3	44.6±0.0
	50	69.2±1.1	80.4±1.0	80.0±1.0	74.7±1.4	47.9±0.4	63.3±2.0	64.0±0.1	62.4±0.2	76±2	87±4	87±3	85±6	60.9±0.0	70.8±0.1	71.7±0.8	72.0±0.6
	75	81.9±0.9	94.0±0.6	93.0±0.6	91.6±1.0	59.6±0.3	78.9±2.4	79.0±0.4	76.9±0.3	87±2	96±2	96±2	94±4	73.3±0.2	86.3±0.3	86.0±0.1	84.4±0.5

Again we see that log-normal differs more from the other results than Weibull, Gamma and the extrapolated calculations differ from each other. Log-normal consistently estimates lower UF than the other distributions and the extrapolation. This means that a researcher interested in a conservative estimate of the UF might wish to choose the log-normal distribution over the others. However, the similarity between the other distributions and the empirical calculation hints at that these distributions may give a more accurate estimate of what the UF would be for these users, if they were provided with a PHEV.

4. Discussion

Though using four data sets makes our results robust, it should be remembered that all data has limitations, in our case, no data set has been measured for a full year, which may influence the variance in driving observed in the data. Furthermore, very long observation times would be needed to accurately estimate tail probabilities for the distributions. Here, we only make statements about the consequence of choosing one distribution over another for tail probabilities. However, to make a statement of precision for the tail, e.g. by using extreme value distributions, even larger data series are needed. Given the continued growth of data availability, this is a possible avenue for future research. With even larger data sets, it would also be possible to investigate distributions with more parameters, or specific distributions for different groups of individuals. Such specific groups could be urban and rural households, or one-car households and multi-car households. The base of our analysis is four data sets of individual daily mileage over different periods of time. The limited observation time of seven days for the German mobility panel data, and the fact that the daily mileages are stated instead of measured sets limits to the conclusions that can be drawn from this data set. Yet the large sample size and the state-of-the-art method of collecting these data (MOP, 2010) make it valuable for a comparison. For the Winnipeg data, the sample size is limited as well as the average driving distances, yet the long observation period allows us to cover long periods of driving.

Further, it should be noted that even though using four data sets from different cities in different countries is a strength that makes our results robust, four data sets is still a small sample for drawing the conclusion that a particular probability distribution is always the best one. We want to emphasise that with respect to the PEV performance measures (DRA and UF), all four data sets agree that Weibull is a good predictor of said measures, but that for goodness of fit, the picture is less clear, even though our results point towards Weibull as well. And thus, for purposes of using overall daily driving distributions, care should be taken to pick a suitable distribution.

We found small or negligible correlation between driving on subsequent days when analysing the independence of daily mileage on subsequent days. Of course higher order autocorrelation, e.g. similar trips every Friday or Sunday as would be present in ACF (7) are not detected by this method. Yet, the Ljung–Box tests we applied show that this does not seem to be an important contribution in the overall autocorrelation in the driving data and further tests of seven-day autocorrelation with the Winnipeg data did not show noteworthy differences from the results stated in Section 3.1.

We found that the results from different goodness of fit measures change with the observation period and they should thus be applied with care. We found that the log-likelihood based AIC converges monotonically for increasing observation period as one would expect from a reliable goodness of fit measure. This goodness of fit measure thus appears preferable.

In the application of our findings to PEV, the comparison to empirical findings is limited by the fact that PEV have only been simulated from the driving of conventional vehicles. Yet, users can be expected to wish to be able to use PEV as they would use conventional vehicles and the comparison made is thereby still relevant.

5. Summary

We analysed goodness of fit statistics for three distributions to describe daily driving distances of individual car users with driving data from different countries and of varying measurements length. Different distribution functions have been used in the literature but with no comprehensive comparison to empirical data so far. We contribute to closing this methodological research gap in several ways: by analysing goodness of fit and the effect of distribution choice on PEV assessment measures, by analysing the independence of daily driving distances, thus enabling analysis by distribution functions, and by analysing the stability of goodness of fit measures with respect to varying measurement period lengths.

Furthermore, the choice of the best matching distribution function has implications for the utility of BEVs and PHEVs as measured by number of days requiring adaptation and utility factor. The log-normal distribution falls off much slower than the Weibull and Gamma distribution for large distances indicating that long-distance trips are more likely, and short distance trips less likely. Thus, the decision between different distributions of daily VKT has direct consequences for calculations of the utility of EVs. Furthermore, we have argued that the AIC measure is a more stable measure of goodness of fit with regards to daily driving data of different measurement period lengths. It should also be noted that this measure favours the Weibull distribution as an overall good distribution.

Summing up we find that for individual daily driving distance the iid assumption is a good approximation and thus it is plausible to apply a distribution function to the data. In contrast to Lin et al. (2012) no single distribution clearly outperforms all the others, though the log-normal and Weibull distributions most often perform better than Gamma. While not conclusive, the Weibull distribution is the distribution that both performs best, for a large number of vehicles and makes reliable average predictions for PEV assessment; while the log-normal estimates are more conservative compared to both the other distributions and the empirical estimates.

These results have implications for researchers interested in modelling their driving data with a specific distribution. Our results underscore the need to carefully judge the effect of choosing a specific distribution and thus choose a distribution that fits the data well. We found the Weibull distribution to fit most vehicles well but not all and at the same time yielding good predictions for PEV related attributes. It should also be noted that choice of distribution function is applied in research that might underlie policy decision, it is thus important to understand how this choice might influence the results.

Acknowledgments

References

Bass, A.J., Dabney, A. and Robinson, D. (2015). *qvalue: Q-value estimation for false discovery rate control*. R package version 2.4.2, http://github.com/jdstorey/qvalue.

Björnsson, L.-H., Karlsson, S., 2015. Plug-in hybrid electric vehicles: how individual movement patterns affect battery requirements, the potential to replace conventional fuels, and economic viability. Appl. Energy 143, 336–347. doi:10.1016/j.apenergy.2015.01.041.

Blum, A., 2014. Electro-mobility. Master Thesis. infernum, Wuppertal.

Brockwell, P.J., Davis, R.A., 2006. Introduction to Time Series and Forecasting. Springer Science & Business Media.

Chan, C., 2007. The state of the art of electric, hybrid, and fuel cell vehicles. Proc. IEEE 95, 704–718. https://www.R-project.org/.

Gnann, T., 2015. Market diffusion of plug-in electric vehicles and their charging infrastructure PhD-thesis. Fraunhofer Publishing, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany.

Gonder, J., Markel, T., Thornton, M., Simpson, A., 2007. Using global positioning system travel data to assess real-world energy use of plug-in hybrid electric vehicles. Transp. Res. Rec. 2017 (1), 26–32.

Greene, D.L., 1985. Estimating daily vehicle usage distributions and the implications for limited-range vehicles. Transp. Res. Part B: Methodol. 19, 347–358.

Hahn, G.J., Shapiro, S.S., 1967. Statistical Models in Engineering. Wiley, New York (1967).

Jakobsson, N., Gnann, T., Plötz, P., Sprei, F., Karlsson, S., 2016. Are multi-car households better suited for battery electric vehicles?—Driving patterns and economics in Sweden and Germany. Transp. Res. Part C: Emerg. Technol. 65, 1–15. doi:10.1016/j.trc.2016.01.018.

Karlsson, S., (2013). *The Swedish car movement data project Final report*. (Report). Chalmers University of Technology, Gothenburg, Sweden.

Kley, F., 2011. Ladeinfrastrukturen für Elektrofahrzeuge—Analyse und Bewertung einer Aufbaustrategie auf Basis des Fahrverhaltens PhD thesis. Fraunhofer Publishing, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany.

Kundu, D., Gupta, R.D., Manglick, A., 2005. Discriminating between the log-normal and generalized exponential distributions. J. Statist. Plann. Inference 127 (1), 213–227.

Lin, Z., Dong, J., Liu, C., Greene, D., 2012. PHEV energy use estimation: Validating the gamma distribution for representing the random daily driving distance. Transportation Research Board 2012 Annual Meeting.

Millo, F., Rolando, L., Fuso, R., Mallamo, F., 2014. Real CO_2 emissions benefits and end user's operating costs of a plug-in hybrid electric vehicle. Appl. Energy 114, 563–571.

MOP, 2010. Mobilitätspanel Deutschland 1994–2010. Institute for Transportation Reseach, Karlsruhe Institute of Technology www.clearingstelle-verkehr.de.

Pearre, N.S., Kempton, W., Guensler, R.L., Elango, V.V., 2011. Electric vehicles: how much range is required for a day's driving? Transp. Res. Part C: Emerg. Technol. 19, 1171–1184.

Plötz, P., 2014. How to estimate the probability of rare long-distance trips. *Fraunhofer ISI Working Paper on Sustainability and Innovation* No. S1/2014. Fraunhofer ISI, Karlsruhe.

Plötz, P., Gnann, T., Wietschel, M., 2012. Total ownership cost projection for the German electric vehicle market with implications for its future power and electricity demand. In: 7th Conference on Energy Economics and Technology Infrastructure for the Energy Transformation, 2012, vol. 27, p. 12.

Plötz, P., Gnann, T., Wietschel, M., 2014. Modelling market diffusion of electric vehicles with real world driving data—Part I: model structure and validation. Ecol. Econ. 107, 411–421.

PSRC (Puget Sound Regional Council), (2008). *Puget sound regional council traffic choices study—summary report*, http://psrc.org/assets/37/summaryreport.pdf. Data available from NREL Secure Transportation Data Project, http://www.nrel.gov/vehiclesandfuels/secure_transportation_data.html.

R Core Team, 2016. R*: A language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Silva, C., Ross, M., Farias, T., 2009. Evaluation of energy consumption, emissions and cost of plug-in hybrid vehicles. Energy Convers. Manage. 50 (7), 1635–1643.

Smart, J., Bradley, T., Salisbury, S., 2014. Actual versus estimated utility factor of a large set of privately owned Chevrolet volts. SAE Int. J. Altern. Powertrains 3, 30–35.

Smith, R., Shahidinejad, S., Blair, D., Bibeau, E., 2011. Characterization of urban commuter driving profiles to optimize battery size in light-duty plug-in electric vehicles. Transp. Res. Part D: Transp. Environ. 16, 218–224.

Storey, J.D., 2002. A direct approach to false discovery rates. J. R. Stat. Soc. Ser. B 64, 479 (2002).

Storey, J.D., Tibshirani, Robert, 2003. Statistical significance for genomewide studies. Proc. Natl. Acad. Sci. 100 (16), 9440–9445.

Tamor, M.A., Gearhart, C., Soto, C., 2013. A statistical approach to estimating acceptance of electric vehicles and electrification of personal transportation. Transp. Res. Part C: Emerg. Technol. 26, 125–134.

Tamor, M.A., Milačić, M., 2015. Electric vehicles in multi-vehicle households. Transp. Res. Part C Emerg. Technol. 56, 52–60. doi:10.1016/j.trc.2015.02.023.

Transportation Secure Data Center (2015). National Renewable Energy Laboratory. Accessed on 2015-03-20. www.nrel.gov/tsdc.

Wickham, H., 2009. ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag, New York, p. 2009.