// NO.20-048 | 10/2020

DISCUSSION PAPER

// VERA HUWE AND JOHANNES GESSNER

Are there Rebound Effects from Electric Vehicle Adoption? Evidence from German Household Data

Leibniz Association

ZEW

Are there Rebound Effects from Electric Vehicle Adoption? Evidence from German Household Data

Vera Huwe* and Johannes Gessner†

September 16, 2020

Abstract

Widespread electric vehicle adoption is considered a major policy goal in order to decarbonize the transport sector. However, potential rebound effects both in terms of vehicle ownership and distance traveled might nullify the environmental edge of electric vehicles. Using cross-sectional household-level microdata from Germany, we identify rebound effects of electric vehicle adoption on both margins for specific subgroups of electric vehicle owners. As our data is cross-sectional, we resort to data-driven methods which are not yet commonly used in the economic literature. For the identification of changes in the number of cars owned after electric vehicle adoption, we predict counterfactual car ownership using a supervised learning approach. Furthermore, we investigate the effect of electric vehicle adoption on household mileage based on a genetic matching of households owning electric vehicles to similar owners of conventional cars. For the selection of covariates for matching, we contrast ad hoc variable selection based on the available literature with a data-driven variable selection method (double LASSO). We cannot verify a significant increase in the number of cars owned for households with one electric and one conventional vehicle. For the subgroup of households who substitute the electric car for a conventional vehicle, electric vehicle ownership is associated with a significant reduction in annual mileage of -23% of the sample mean. The result indicates a strive for behavior consistent with the environmentally-friendly car choice rather than a rebound effect. Our results are subgroup-specific and may not generalize to the overall population. Methodologically, we find that data-driven variable selection identifies a refined set of covariates and changes the magnitude of the estimation results substantially. It may thus be considered a useful complement, especially in settings with limited theoretical or empirical knowledge established.

Keywords: Rebound Effect, Electric Vehicle Adoption, Variable Selection

JEL-Classification: R41, Q55

*ZEW – Leibniz Centre for European Economic Research, Environmental and Resource Economics, Environmental Management, P.O. Box 10 34 43, 68034 Mannheim, Germany, Email: vera.huwe@zew.de

†University of Mannheim, Department of Economics, Email: jgessner@mail.uni-mannheim.de

1 Introduction

In many countries in the Global North large-scale electric vehicle (EV) adoption has become a firm political objective in order to decarbonize the transport sector where greenhouse gas (GHG) emissions continue to rise steadily (EEA, 2019). However, it is a matter of controversial debate whether EV adoption reduces overall carbon emissions or not. Besides higher production emissions for EVs (Held et al., 2011) and CO_2 emissions caused by electricity generation, their decarbonization potential might be undermined by rebound effects. Rebound effects may occur on two different margins: On the *extensive margin*, households may buy an EV as an additional car in order to shift trips previously completed with zero- or low-carbon means of transport like walking, cycling or public transit or to undertake additional trips. On the *intensive margin*, EV adoption may lead to a rebound effect in terms of distance traveled by car, either as a result of the lower operating costs EVs reveal compared to conventional cars (Lévay et al., 2017) or due to moral licensing (Miller et al., 2010, for a review). Taking these potential rebound effects into account is crucial for the comprehensive assessment of the potential EVs have for transport decarbonization. Moreover, rebound effects would have important implications for transportation policy in more general. As EVs use road infrastructure and induce road-related externalities like congestion and accidents but do not pay fuel taxes, Davis et al. (2019) discuss whether they should therefore pay a mileage tax. Externalities might magnify if EVs were deployed as additional cars driving additional kilometers on a large scale.

Yet empirical evidence on the existence and magnitude of potential rebound effects is scarce. Most existing studies in the context of EV adoption make uniform assumptions about counterfactual mileage for all EVs without empirical backing. For instance Holland et al. (2016) and Davis et al. (2019) simply assume 15,000 miles/year for all EVs considered. Demand models as in Xing et al. (2019) explicitly do not account for the case of electric vehicles replacing non-car means of transport and may hence overestimate the implied emission reduction. A notable exception is Klöckner et al. (2013) who documents both extensive and intensive margin rebound effects among electric vehicle owners in Norway based on an online survey. Yet the response rate is rather low and Norway constitutes a special case due to a particularly high share of renewable energy sources in the power mix (NCM, 2014) and EVs in the Norwegian car fleet. By contrast, we focus on Germany which has a substantially lower share of electric vehicles in the overall fleet. On top of that, the share of renewable energy in the power mix in Germany has increased but is also substantially lower than in Norway (Agora Energiewende, 2018).

We use comprehensive household-level microdata from the 2017 German national travel survey *Mobilität in Deutschland* (MiD) for an empirical assessment of the potential rebound effects. Households report engine type and vehicle kilometers traveled (VKT) for one year at the vehicle level as well as an extensive set of sociodemographic characteristics. Our treatment of interest is EV adoption. As treatment effects may vary conditional on the household's car portfolio, we restrict the analysis to a well-defined subsample. We investigate households with one EV and one internal combustion engine vehicle (ICEV) as our treatment group of interest as it is the largest subgroup among the EV-owning households in the sample. We consider two possible cases: Either the EV may represent an ICEV substitute or it may be bought as an additional car. According to this definition, we observe 183 treated households. We first predict the treated households' counterfactual number of cars by means of a supervised learning approach. 23.5% of treatment households are predicted to have bought the EV as an additional car. We develop a test for significance and assess general model performance measures. The evidence on the presence of an extensive margin rebound effect is inconclusive. Building on the extensive margin analysis, we restrict the analysis to 140 ICEV substituters in order to clearly disentangle the intensive margin effect of EV adoption from the effect of buying an additional car. We investigate the effect of EV adoption on household mileage by means of variable ratio genetic matching. For variable selection, we contrast a selection based on a literature review referred to as *ad hoc* with a data-driven variable selection method (double LASSO). As ad hoc variable selection has been criticized for a lack of clear guidance on how to select the relevant variables (as well as the functional form), data-driven variable selection is suggested as complement for settings where identification relies on identifying the correct set of relevant observables (Belloni et al., 2013). The direct comparison is also informative for identifying weaknesses of the existing literature and for verifying effect sizes. We find a significantly negative effect of EV adoption on aggregate household mileage in both the LASSO-based and the ad hoc specification. Consequently, EV adoption seems to induce a strive for behavior consistent with the environmentally-friendly car choice rather than an intensive margin rebound effect among ICEV substituters. Evidently, lower annual mileage relates to lower GHG emissions, irrespective of how the mileage is split among the household's cars. Lastly, we find that data-driven variable selection provides valuable insights. It does not only identify a refined set of relevant confounders but also reveals that the ad hoc specification underestimates the effect by a sizeable 30%.

This paper's contribution is twofold. While the largest part of the existing literature operates in hypothetical settings, our analysis is among the first to provide comprehensive empirical evidence among actual EV owners on both extensive and intensive margin rebound effects. Evidently, these potential effects are directly linked to the potential EVs offer for decarbonizing the transport sector. In contrast to the previous literature, we explicitly allow for EV adoption affecting car ownership. Secondly, the paper makes a methodological contribution by contrasting a data-driven variable selection method with ad hoc variable selection which provides novel insights into the potential of data-driven methods for verifying variable selection and estimation precision.

The remainder of this paper is structured as follows. In section 2, we provide background information about the data. Section 3 describes the empirical strategy and section 4 presents our results. Section 5 concludes.

2 Data

We draw on recent household-level micro data from the German national travel survey *Mobilität in Deutschland 2017* (MiD). The MiD is a large-scale household survey commissioned on behalf of the German Federal Ministry for Transport and Digital Infrastructure (BMVI) every few years. The 2017 edition is the first edition with a sizeable number of 410 EV households. Our unit of analysis is the household. At the car level, households report the number of cars owned, their respective engine type which provides information on whether a vehicle is conventional or electric as well as annual vehicle kilometers traveled (VKT) for up to three household cars.[1]Household VKT is obtained by aggregating VKT across all household cars. We consider the car-owning households who report complete car-level information only.[2] Among those households, ICEV-EV ownership is distributed as presented in Table 1. 0.29% of all households own at least one EV. While still low, this

no. ICEV/no. EV	0	1	2	3+	Sum
0	-	0.04	0.01	0.00	0.05
1	55.20	0.15	0.01	0.00	55.36
2	35.23	0.08	0.00	0.00	35.31
3 +	9.29	0.00	0.00	0.00	9.29
Sum	99.72	0.27	0.02	0.00	100.01

Table 1: Distribution of car portfolios among car-owning households in the MiD 2017

share exceeds the 0.1% EV market penetration in the overall German car market in 2017 (Kraftfahrtbundesamt, 2017) by factor 3. Among the EV-owning households, the largest share owns one EV and one ICEV (52%).

A rich set of socio-demographic characteristics is available at both the individual and the household level. For variables elicited at the individual level, we aggregate the information to the household level.[3] We also summarize several reply options for dimension reduction

[1]The data set consists of several sub-datasets which report information at different levels. The survey methodology is described in detail in Nobis et al. (2018). For our analysis, we draw on the sub-datasets provided at the car level, the household level and the individual level. Some variables are included in more than one sub-dataset. As some of the variables mismatch in different datasets and the set of households who reported information varies between the different levels, we consider the car-level data as the main data source. Consequently, we only cover households who report at the car level. For these households, we then merge additional household- and individual-level information from the other sub-datasets.

[2]More precisely, we exclude all households who do not report engine type or valid VKT for any of their cars reported in the car-level data. We further exclude households with hybrid or non-specified engine types as they are not clearly identifiable as EV or ICEVs. If in a two-vehicle household one car is found to have reported zero VKT, we interpret this as a car which is not registered or broken and thus not available for use. In this case, we correct the number of cars to one car which is available for use.

[3]We define an indicator for the existence of at least one household member with (one or more) car sharing membership(s). We elicit the maximum education level reported among the members of a household. We compute the share of adults who regularly use environmentally-friendly means of transport (bike, public transit), referred to as ENVI, where regular is defined as at least once a week. We compute the share of adults, the share of adult females and the share of unemployed household members to reflect

for some variables.[4] A detailed list of the extended variable set we use with all (summarized) reply options can be found in Appendix A. As information on sociodemographics is not perfectly complete for the whole sample, we restrict the sample to complete cases if the variable is missing for below 5% of households considered and impute missing values otherwise.[5]

We are interested in the effects of EV adoption on household mobility patterns including car ownership and mileage traveled by car. We regard EV adoption as the treatment and hence use the terms treated household and EV household interchangeably. In general, notice that the effects of EV adoption may vary depending on the household's car portfolio. Particularly, in multi-vehicle households EV adoption might be less likely to affect total mileage as the EV's limited range can be compensated by between-vehicle substitution. In order to investigate a well-defined treatment effect, we restrict the analysis to specific subgroups of EV-owning households. To compare households with identical substitution possibilities, we restrict the set of treated households to EV households owning one EV and one ICEV. In this restricted treatment group, we observe 183 EV households. For the effect of EV adoption on car ownership, we allow for two plausible cases: Firstly, the EV could have replaced a previously owned ICEV. In this case, the household had been a two-ICEV household prior to EV adoption and no extensive margin rebound occured. We label this case *substitution.* Secondly, if the household had previously owned only one ICEV, the EV represents an *additional* car and induces an extensive margin rebound effect. Importantly, for the EV households considered this implies that treated households had been ICEV-only households prior to EV adoption. Thus, we consider the effects of EV adoption in ICEV-only households.[6]

We further specify the definition of treatment for the intensive margin and extensive margin analysis separately. On the extensive margin, as some households may replace their previously owned ICEV while others buy the EV as additional car, the effect of EV adoption on car ownership is potentially heterogeneous. Treatment on the extensive margin is thus defined as the decision to adopt an EV. As a result, the subset of ICEV-only counterparts in the data is the universe of one- and two-ICEV households. Henceforth, we refer to this set of households as ICEV households. We observe 66,830 households

the household's structure. We count the number of adults with a valid driving license as these are the ones that can use the household's car(s) independently of each other. We also count the number of adult employees as a proxy for the number of household members most likely to be commuters. Based on personal characteristics, we exclude minors-only households and households which do not report age for all household members as we cannot compute all aggregate characteristics.

[4] We summarize home ownership into home owners and renters. We create an indicator of at least one child below 18 (instead of using the number of kids) in the household and place of residence in East Germany (instead of using the precise state). We also categorize employment status into employed full-time, employed other than full-time, unemployed, retired and in education (traineeship, school, university).

[5] We use nearest neighbor imputation for the variables highest vehicle segment and share of regular bike/public transit users.

[6] Hence we exclude the possible but in our opinion unlikely cases that an EV replaces more than one car or that EV adoption leads to the adoption of an additional ICEV in households that did not own any cars before treatment.

	extensive margin	**intensive margin**
treatment group	households with 1 ICEV and 1 EV	households with 1 ICEV and 1 EV where the EV is an ICEV substitute
n	183	140
control group	households with 1 or 2 ICEVs	households with 2 ICEVs
n	109.728	42.898

Table 2: Definition of treatment and control groups

owning one ICEV and 42,898 owning two ICEVs. Whether the EV represents an ICEV substitute or an additional car determines the number of household cars post-treatment. For households buying the EV as additional car, mileage effects may both arise from the effect of having an additional car at disposal and from the car being an EV. In our setting, these two effects cannot be disentangled. In order to identify a well-defined treatment effect at the intensive margin, we further restrict the treatment group to households for which the EV is predicted to be an ICEV substitute in the extensive margin analysis. This reduces our treatment group to 140 EV households. On the intensive margin, the treatment considered is the substitution of one ICEV by an EV. The observed group is restricted to households owning two ICEVs pre-treatment. As substitution among vehicles may occur within-household upon EV adoption, our variable of interest is the annual household VKT which we refer to as household mileage in the following. Yet notice that a rebound effect in household mileage does not necessarily imply a rebound effect in GHG emissions as the effect on GHG emissions depends on how mileage is split between the electric and the conventional car.[7]. Table 2 summarizes the relevant information for our definition of treatment for both the extensive and intensive margin analysis.

3 Empirical Strategy

As we work with cross-sectional non-experimental survey data, we require some assumptions to reintroduce the temporal structure required to conduct a treatment study. Particularly, we do not know the temporal order of car adoption. Therefore, we make a first crucial assumption:

Assumption 1: The EV is the car bought most recently.

Otherwise the car portfolio could have changed since EV adoption and we would not be able to attribute the effects observed to EV adoption. Moreover, any information we observe reflects the moment of survey completion and is thus post-treatment. Yet for the analysis to be informative, we need the covariates controlled for to reflect pre-treatment

[7]A rebound effect on GHG emissions may not occur if the EV is used to cover a larger share of total mileage, even if total mileage increases.

information which is only true if they have not changed since EV adoption. Therefore, we require the following second assumption:

Assumption 2: The observable covariates and proxy variables for unobservables used in the intensive and extensive margin analysis have not changed since EV adoption.

Importantly, assumption 2 includes that the relevant covariates have not been affected by the treatment itself. We believe that this assumption is innocuous for the covariates we use. The full set of covariates used in both models can be reviewed in Table 9. Assumptions 1 and 2 allow us to consider EV adoption as treatment, the covariates as pre-treatment information and only the outcomes car endowment and household VKT as post-treatment information. We discuss our assumptions in depth in section 5.

3.1 Extensive margin

Given the treatment defined above, we investigate the effect of EV adoption on the number of household cars based on the potential outcome framework as first formulated by Rubin (1974). Households differ in terms of their treatment status $T_h \in \{0, 1\}$, with one indicating treatment and zero denoting the absence of treatment. We designate the set of EV households as H_{EV} and the set of ICEV households as H_{ICEV}. For household h in treatment status T_h we observe the number of cars $Y_{h,T}$. To identify extensive margin rebound effects, we need to evaluate if EV adoption has led to an increase in the number of cars owned for a significant share of EV households. Under the given assumptions, the existence of extensive margin rebound effects amounts to testing the following hypothesis:

Hypothesis 1. *An extensive margin rebound effect of EV adoption defined by the event* $A = \{Y_{h,0} = 1 \wedge Y_{h,1} = 2\}$ *occurrs with non-zero probability.*

If we had complete information on $Y_{h,0}$ and $Y_{h,1}$, a feasible hypothesis test for Hypothesis 1 would be

$$H_0\text{: } P(Y_{h,1} = 2 \wedge Y_{h,0} = 1) = 0 \text{ against } H_1\text{: } P(Y_{h,1} = 2 \wedge Y_{h,0} = 1) \neq 0.$$

Yet only one of the potential outcomes $Y_{h,T}$ with $T \in \{0, 1\}$ is observed. Particularly, $\forall h \in H_{EV}$ $Y_{h,1}$ is observed, while $\forall h \in H_{ICEV}$ we observe $Y_{h,0}$. Thus, in order to assess the extensive margin rebound effect, an estimate of the counterfactual number of cars an EV household would have owned in absence of treatment, i.e. $Y_{h,0}$ for $h \in H_{EV}$, is required. We address the missing information problem using a supervised learning approach. As we only allow the EV to represent an ICEV substitute or an additional car, the outcome is binary by construction, i.e. $Y_{h,0} \in \{1, 2\}$. Furthermore, households characteristics X_h represent pre-treatment information by assumption 2. We can thus predict the counterfactual number of household cars in EV households by exploiting correlation between household characteristics X_h and the number of household cars Y_h prevailing among ICEV households.

We do so in three steps. First, we estimate the relationship between household characteristics X_h and observed car ownership $Y_{h,0}$ among the set of ICEV households H_{ICEV} with a binary logit model. To be precise, we estimate a model classifying a household with one ICEV as the "event" and a household with two ICEVs as the "non-event". To be able to evaluate the models performance after model estimation, we split the ICEV households randomly into a train set and a test set comprising 80% and 20% of the ICEV households, respectively. The model is estimated on the train data only. Secondly, the model performance is measured on the test data by a confusion matrix, the model's sensitivity and specificity as well as the area under the curve (AUC).[8] To do so, we predict the number of household cars for the households in test data which had not been used for model calibration. The model estimation is supervised in the sense that we know the true number of cars ICEV households own which we then compare to the model predictions on test data for assessing model performance. As the set of ICEV households available comprises a very large number of households, we are able to estimate the model with high confidence. By construction, EV households would have been one- or two-ICEV households in absence of treatment. Thus, the model's predictive ability on test data can be considered an estimate for the predictive ability on EV households. In a third step, we predict the number of cars EV households would have owned in absence of treatment based on the model estimated in the first step. Using the estimate for the share of EVs predicted to be substitutes for a previously owned ICEV from the treatment sample and the estimate for the model's specificity on the test data, we develop a feasible way to test Hypothesis 1. Details on the derivation can be found in Appendix C.

More details on steps 1 and 3 are provided below. We have a large set of observables χ_h. In the first step, we use the training data subset of ICEV households to estimate the relationship between the observables χ_h and the number of cars $Y_{h,0}$ in a household without EVs in a binary logit model with $Y_{h,0} \in \{1, 2\}$. As not all observables χ_h may be relevant predictors of $Y_{h,0}$, we exploit approximately linear relationships between the relevant subset of observables X_h and the logit of the outcome $Y_{h,0} = y_{h,0}$.[9]

$$\ln\left(\frac{P(Y_{h,0} = y_{h,0})}{1 - P(Y_{h,0} = y_{h,0})}\right) = X_h'\beta + \varepsilon_h$$

[8] The confusion matrix lists the two dimensions predicted versus actual number of cars for the set of classes, in our case one or two vehicles. Sensitivity and specificity report the true positive rate and true negative rate, respectively. The AUC summarizes the area under the Receiver Operator Curve (ROC) which plots the share of true positive predictions against the share of false positive predictions. In our case, the classifications are based on the predicted propensity scores for $Y_{h,0}$ from the logit model and a cutoff probability to predict $Y_{h,0}$. This cutoff probability again is found by minimizing the model's misclassification error rate on the test data. The ROC is plotted for each cutoff level. An AUC between 0.5 corresponds to random guessing, an AUC of 1 to perfect predictions. For a more detailed discussion of the AUC for model evaluation refer to Bradley (1997).

[9] In an approximately linear model all regressors enter linearly or as their transformations. Small approximation errors are permitted.

In the given regression formula, the left-hand side of the formula gives the logit for the number of ICEVs in a household, the regression coefficients are given by β and ε_h is household h's linear projection error. As we assume the model to be approximately sparse,[10] we can use the Least Absolute Shrinkage and Selection Operator (LASSO) introduced by Frank et al. (1993) and Tibshirani (1996) to select the relevant predictors X from the overall set of covariates χ. The LASSO regularization optimizes an objective function based on mean squared error plus a shrinkage term λ which penalizes the size of the model and thus leads to improved out of sample predictions in comparison to the unpenalized model.[11] Under the sparsity assumption and a set of fairly general regularity conditions, the LASSO identifies the set of variables with non-zero coefficients and shrinks all other coefficients to zero (cf. Chetverikov et al., 2016). Yet in our data the majority of observed covariates is scaled ordinally or nominally. As categorical variables are included as separate dummy variables, the LASSO is likely to exclude single categories of a variable without excluding the entire covariate (Meier et al., 2008). The modified group LASSO as first introduced by Bakin (1999) and later generalized by Yuan et al. (2006) summarizes these dummy variables into a group which can only be exempt as a whole. The group LASSO optimizes the following penalized regression formula:

$$\hat{\beta} = \arg\min \beta \frac{1}{2} \|y - \chi\beta\|_2^2 + \lambda \sum_{g=1}^{G} \sqrt{p_g} \left\| \beta^{(g)} \right\|_2$$

where $\lambda \geq 0$ controls the degree of penalization, $\left\|\beta^{(g)}\right\|_2$ is equal to $\sqrt{\sum_{j \in I_g} \beta_j^2}$,[12] G is the number of non-overlapping groups I_g of covariates in χ and p_g is the number of covariates in group g (Yang et al., 2015). By shrinking the coefficients of the irrelevant covariates in χ to zero, the group LASSO selects the relevant covariates X for our logistic regression model. We apply a group LASSO and group dummy regressors with respect to the underlying categorical variable as suggested by Yang et al. (2015).[13] To obtain the optimal rate of convergence, λ is chosen via K-fold cross-validation (cf. Chetverikov et al., 2016).[14] We use $K = 10$ and choose the value of λ corresponding to the highest amount of shrinkage which is still within a one standard error range of the lowest cross-validated average misclassification error rate, referred to as λ_{1se}, with the missclassification error

[10] Among all available confounders χ, only s variables X have non-zero coefficients, with $s << N$, N being the number of observations.

[11] Out of sample predicion performance becomes relevant when a prediction model is estimated on one data set, and then used to make predictions on another data set. Without the penalty term, there is a risk of over-fitting the model to the data used for model estimation.

[12] It is important to note that all covariates in χ are standardized before model estimation, such that the scale of a covariate does not matter for the degree od penalization.

[13] We use the R-package "gglasso" to implement the group LASSO.

[14] For K-fold cross-validation, the data set is partitioned into K subsets and a LASSO-regression fitted for each subset and each candidate value of the penalty term λ. For a more detailed discussion, see for instance Chetverikov et al. (2016). In each cross-validation step, the group LASSO formula is optimized over a grid of 100 different lambda values. For details on how the grid is chosen refer to Yang et al. (2015). The average cross-validated error for each λ and the corresponding standard deviation are retrieved.

rate being $Error = \frac{1}{n}\sum_i I(y_i \neq \hat{y}_i)$. λ_{1se} is a heuristic choice of λ in cross-validated LASSO models aiming to produce a less complex but statistically equally good model.

In the third step, the model is applied to the EV households to predict the counterfactual number of cars in these households in absence of treatment, i.e. $\hat{Y}_{h,0}$. Using the predictions $\hat{Y}_{h,0}$ $\forall h \in H_{EV}$ and estimates of the predictive ability of $\hat{Y}_{h,0}$ retrieved from the test sample, we recover a feasible test for Hypothesis 1. As we develop in Appendix C, under H_0 and given the channels *addition* and *substitution* for EV adoption, the model's specificity $\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2)$ is equal to the probability of an EV being classified as an ICEV substitute $\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2)$. This is because the only channel left for EV-adoption under H_0 is the substitution of a previously owned ICEV. Thus, the number of cars on a household pre- and post-treatment are equal under H_0, $Y_{h,0} = Y_{h,1}$.
We can estimate the share of households with $\hat{Y}_{h,0} = 2$ for $h \in H_{EV}$ as $\frac{1}{|H_{EV}|}\sum_{h=1}^{|H_{EV}|} I(\hat{Y}_{h,0} = 2)$. A sample analogue for $\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2))$ is captured by the model's specificity estimated on the ICEV test data. Assuming that $\hat{Y}_{h,0}$ are independent and identically distributed over households[15], we find

$$\frac{1}{|H_{EV}|}\sum_{h=1}^{|H_{EV}|} I(\hat{Y}_{h,0} = 2) \xrightarrow{d} \mathcal{N}(\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2), \sigma_1)$$
$$Specificity \xrightarrow{d} \mathcal{N}(\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2), \sigma_2)$$

Taken together, this allows us to test Hypothesis 1 by testing

$$H_0' : \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2) = \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2) \; against$$
$$H_1' : \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2) \neq \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2)$$

using Welch's two-sample t-test (unpaired, two-sided).[16] The validity of testing H_0' to potentially reject H_0 is derived in Appendix C.

3.2 Intensive margin

Building on the extensive margin analysis, we investigate the effect of EV adoption on household mileage for ICEV substituters in a second step. Again we base our analysis on the potential outcome framework (Rubin, 1974). We are interested in household mileage completed by EV households in the absence of treatment which is unobserved. The households considered would have been two-ICEV households in absence of treatment. Hence the universe of two-ICEV households represents the relevant set of potential control households. We use a genetic matching approch to identify the subset of households which is most similar to the treatment households in all relevant pre-treatment characteristics.

[15]Given the model estimated in step 1, drawing $\hat{Y}_{h,0}$ is equivalent to drawing X_h, which we assume to be an independent and identically distributed (iid) random vector. Thus, $\hat{Y}_{h,0}$ is iid as well.

[16]Notice that $\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2))$ and $\mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2)$ are expectations over different sub-populations. Thus, the validity of testing H_0' instead of H_0 is not trivial.

Thus, the set of relevant pre-treatment characteristics needs to be established in a first step. We refer to this step as *variable selection*. In this paper, we contrast an ad hoc variable selection based on existing theoretical and empirical evidence with a data-driven variable selection procedure (double LASSO). We compare the sets of variables selected with respect to congruence. The methods are described in detail in the following. The data-driven variable selection is based on a group LASSO as described in section 3.1. We again assume an approximately linear, sparse model. In addition to sociodemographic factors, attitudes and preferences may play a role for the adoption decision. As the MiD does not elicit attitudes directly, we additionally assume the following for the conditional independence assumption (CIA) to hold:

Assumption 3: All relevant attitudinal determinants of EV adoption and VKT can be proxied by observables.

Our equation of interest reads as follows. For household h,

$$VKT_h = \alpha EV_h + X_h\theta + r_h + \zeta_h \tag{1}$$

with outcome VKT_h, treatment EV_h, the set of relevant control variables X_h, approximation error r_h and error term ζ_h with $E[\zeta_h \mid EV_h, x_h, r_h] = 0$. α is the treatment effect of interest. However, running group LASSO on this single equation is problematic as the LASSO may shrink α to zero if the treatment is correlated with the confounders. Moreover, regularization tends to underestimate non-zero coefficents and may erronously exclude variables with moderate non-zero coefficients. The CIA posits that treatment assignment becomes ignorable only after controlling for all relevant confounders which are related to both treatment status and outcome. Hence, if the mentioned variables are predictors of treatment, omitted variable biases may arise. In order to safeguard against such biases, it is necessary to split the analysis into two separate LASSO estimations as suggested by Belloni et al. (2014). The so-called double LASSO procedures subjects both treatment and the outcome variable of interest to a separate LASSO.[17] We first fit a group LASSO to predict the outcome variable household mileage with the observables (henceforth VKT LASSO).[18]

$$VKT_h = x_h'\theta_{VKT} + r_{VKT,h} + \nu_h \tag{2}$$

We retrieve the set of θ_{VKT} predictors. Secondly, we fit a group LASSO to predict EV ownership with the available observables in order to capture the relationship between treatment and controls (henceforth EV LASSO). As the outcome variable is binary, we

[17] An easy-to-read guideline can be found in Urminsky et al. (2016).

[18] For the VKT LASSO, we run OLS. We use the R package *gglasso* and choose the L2 norm as loss function for the cross-validation. Note that we cannot directly infer effect magnitudes as standard errors are incorrect (Belloni et al., 2013).

use logistic regression for the classification problem. In both cases, we again split our data into 80% used to train the model and 20 % for model evaluation. In the model selection step, we use 10-fold cross-validation on the train data to choose the optimal penalty parameter λ. Choosing the tuning parameter is less straightforward for variable selection (Kirkland et al., 2015). We again choose λ_{1se} as tuning parameter. Yet as EV ownership is a rare event, defined as the minority class comprising less than 5% of the sample, the data is imbalanced with respect to class. Resampling can be applied to reduce class imbalance (see e.g. Kotsiantis et al., 2006, for a review). We apply under-sampling, i.e. we randomly draw as many households from the set of potential control group households as we have treated households and discard all other control households so that both treatment and control group have the same class size. We then run the group LASSO on the re-sampled train set of treated and control households. Instead of a least squares error, we use the misclassification error rate as the loss function to be optimized in cross-validation for the EV LASSO. We repeat this procedure 100 times. In each of the 100 iterations, for household h we have

$$EV_h = x_h' \theta_{EV,i} + r_{EV,h,i} + \nu_{h,i} \tag{3}$$

In each iteration i, we retrieve the set of $\theta_{EV,i}$ predictors. We count in how many iterations a variable is selected and choose those selected in at least 75 iterations, $\theta_{EV} = \cup\, \theta_{EV,\, c \geq 75}$ for count c. This threshold is chosen ad hoc. As robustness checks, we choose 60 and 90 as alternative cutoffs. After fitting the Logit-estimator with the selected set of covariates to the entire train set, we evaluate the model's predictive ability on the test data. The matching is then conducted based on the union of variables selected from the double LASSO, $\theta_{DL} = \theta_{VKT} \cup\ \theta_{EV}$. θ_{DL} represents the set of matching variables. For the alternative ad hoc specification, we follow the same logic and select variables related to EV adoption and household mileage based on a review of the available empirical literature.[19] Assumption 2 guarantees that we do not face endogenous control problems. Nonetheless, an additional assumption is required for the matching procedure to be interpretable. As the survey questions do not provide information on when a vehicle was bought, we cannot differentiate EV households based on the duration of EV ownership. Instead, we observe the average effect in a population with unknown distribution in duration of EV ownership. The effect observed in the sample reflects the actual treatment effect only in the special case of constant treatment effect over time. For ease of interpretation, we posit that the treatment effect unfolds upon adoption and does not amplify or perish over time.

Assumption 4: The treatment effect is constant over time.

[19] Yet notice that the LASSO detects the relationship between the outcome variable and the characteristics for one- and two-vehicle households only while the literature considers EV adoption across all types of households and does not allow for further differentiation based on car portfolio. If the chosen subsample differs significantly from the overall population, the ad hoc variable selection may not precisely reflect the subsample of interest.

Given this assumption, we use *EV adoption* and *EV ownership* interchangeably. Based on the respective set of matching variables, we apply variable ratio genetic matching with replacement as introduced by Diamond et al. (2013).[20] The genetic matching algorithm is a multivariate matching method based on a search algorithm to optimize post-matching covariate balance. The algorithm matches treated and control households based on a generalized Mahalanobis distance

$$GMD(Z_h, Z_h|W) = \sqrt{(Z_h - Z_h)^T (S^{-\frac{1}{2}})^T W (S^{-\frac{1}{2}})(Z_h - Z_h)}$$

with $Z_h = \begin{pmatrix} X_h \\ PS_h \end{pmatrix}$, the set of covariates X and propensity score PS. In extension of a standard Mahalanobis distance, the weights W of the individual variables are optimized by minimizing the largest individual discrepancy based on p-values from a Kolmogorov-Smirnov test and paired t-tests for all variables chosen for matching. This procedure is particularly flexible, as it allows for dimensionality reduction by assigning potentially all weight to the propensity score but also for enhancing post-matching balance by including further variables to match on (Diamond et al., 2013). Finally, we investigate the average treatment effect on the treated (ATT) in the matched sample. We now estimate the equation of interest (1) outlined above on the matched subsets, with θ being the set of selected control variables θ_{DL}, respectively. Weighted OLS is used to reflect the fact that different controls assigned to the same treated households may have different values for some covariates.

4 Results

4.1 Extensive Margin Analysis

The model developed in section 3.1 predicts household car ownership to consist of one ICEV before treatment for 23.5% of the observed EV households.[21] As all EV households studied own two vehicles post-EV adoption, the EV is predicted to represent an additional car for these households. Conversely, the model predicts 76.5% of EV households to have substituted an ICEV. The model's specificity, i.e. the rate of true substitution predictions estimated on test data, is moderate. Only 75.3% of all test households owning two ICEVs were also predicted to own ICEVs. The remaining 24.7% were falsely predicted as one-ICEV households. Table 3 displays the results of testing Hypothesis 1 using H_0'. According to the developed test, the share of EVs bought as ICEV substitutes is not significantly different from the model's specificity ($t(190) = 0.39, p = 0.70$). Hence, we cannot reject the feasible hypothesis H_0' which implies we cannot reject H_0 either. Based on this test we

[20]For implementation, we use the MatchIt R package (Ho et al., 2011).
[21]The variables selected for the model can be found in Table 9 in Appendix A.

conclude that the share of EVs bought as additional cars is not significantly different from zero. Yet notice that the test reveals reduced power of testing hypothesis H_0' compared to testing H_0 which is infeasible in the potential outcome framework given, which may also explain the insignificance. Further model evaluation measures as reported in Table 3 suggest that the prediction is relatively precise. The sensitivity of 83% estimated on the test data shows that, given the model applied for the prediction of the counterfactual number of ICEVs in the treated households, we can expect that 83% of the EVs bought as an additional car are also predicted to be additional. The confusion matrix in Table 4 reveals that the share of ICEV households predicted to own one car is very close to the actual share of two-ICEV households in the test data. Given the AUC of 0.88, the model's prediction performance can be considered medium to high. Taken together, the evidence on whether an extensive margin rebound effect exists for the subsample of EV owners considered is inconclusive. While the test developed does not indicate the share of additional EVs to be significant, other performance measures suggest that a substantial share of additional EVs exists in the sample.

Measures of Prediction Performance	
Sensitivity	0.8296
Specifity	0.7527
AUC	0.8801
Test Results	
Share of substitution EV households	0.7650
Specificity	0.7527
Difference	0.0124
Confidence Interval	[-0.0503, 0.0750]
t-Statistic	0.3894
p-Value	0.6974
Degrees of Freedom	190.0993

Table 3: Measures of Prediction Performance and Test for Extensive Margin Rebound Effects

	$Y_{h,0} = 1$	$Y_{h,0} = 2$	Share
$\hat{Y}_{h,0} = 1$	11103	2118	60.24
$\hat{Y}_{h,0} = 2$	2280	6445	39.76
Share	60.98	39.02	100.00

Table 4: Confusion Matrix

4.2 Intensive Margin Analysis

We first elaborate on how the ad hoc variable selection is built and then compare it to the data-driven variable selection results. In a second step, we run the respective matching algorithms and report the results. Table 5 provides an overview of the sociodemographic

Study	Country	Sample size	Socio-demographics
EV adoption			
Anable et al. (2011)	UK	N=2,729	gender, income
Hidrue et al. (2011)	U.S.	N=3029	age, education
Plötz et al. (2014)	Germany	N=969/N=210	full-time employment, regional type, household size
Vassileva et al. (2017)	Sweden	N=247	gender, education, income
Sovacool et al. (2018)	Nordic countries	N=5,067	gender, education, full-time employment, age
VKT			
Büchs et al. (2013)	UK	N=24,446	income, household size, regional type, education, age structure
Munyon et al. (2018)	U.S.	N=82,485	income, regional type, home ownership, household size, number of drivers, number of workers
Sovacool et al. (2018)	Nordic countries	N=5,067	household size, age structure, employment

Table 5: Literature review on sociodemographic factors related to EV adoption

factors related to EV adoption or VKT identified in the existing literature. As there are important caveats in the literature, little causal knowledge is yet established. Importantly, the evidence on EV adoption is almost exclusively survey-based. Except Vassileva et al. (2017), none of the studies is conducted among actual EV owners. Instead, respondents face a hypothetical choice setting or report their intention to adopt. However, an intention-behavior-gap is well established in the psychological literature (e.g. Sheeran, 2002, for an overview). The surveys may thus not necessarily reveal the characteristics of actual early adopters. Moreover, the studies were conducted in different countries with heterogeneous institutional arrangements for EV adoption. Country-specific results may not be generalizable. In addition, most studies date back several years while conditions for EV adoption and usage have improved in many countries in recent years. As results are inconsistent among the different studies, we mainly build on Plötz et al. (2014) to capture the German context as closely as possible. As Plötz et al. (2014) do not elicit income-related information, we add economic status as socioeconomic drivers which early adoption theory has pointed out as relevant (Rogers, 2003) and which is univocally confirmed in other empirical studies (Anable et al., 2011; Hidrue et al., 2011). Moreover, more recent studies reveal that home ownership is important for EV adoption (Vassileva et al., 2017; Davis, 2019). Given the German legal context, the lack of home ownership represents a plausible impediment to EV adoption, particularly for tenants in apartment buildings, and we hence include home ownership status as a matching variable.[22] With re-

[22]According to German law (Wohnungseigentumsgesetz, §22, Abs. 1), all owners of a building have to agree to the installation of a charging facility for electric vehicles. Renters living in an appartment thus need to seek approval from all other apartment owners and convince the landlord to cover the installation

spect to attitudinal determinants, Plötz et al. (2014) find that pro-environmental attitudes and affinity to technology play an important role for intended EV adoption.[23] As the MiD does not elicit attitudinal variables, we use the share of retirees as proxy for openness to technology and the level of household education as proxy for pro-environmental orientation as e.g. found in Meyer (2015). With respect to household mileage, we follow Munyon et al. (2018) which represents the most comprehensive analysis but add the share of the household's retirees as well as the share of adults as indicators of age structure identified to be relevant in both Büchs et al. (2013) and Sovacool et al. (2018). We again use education as a proxy for environmental awareness as the relevance of environmentally-friendly attitudes for car mobility is demonstrated in e.g. Kahn (2007). Table 9 in Appendix A presents the variables selected ad hoc alongside the LASSO results. With respect to EV adoption, the LASSO confirms only the subset of ad hoc variables reflecting economic status and age structure. As can be seen in Table 9, the variable selection is relatively robust across the three cut-offs. Therefore, we use the intermediate cut-off of 75 for further analysis.[24] With respect to household VKT, the LASSO confirms all variables identified ad hoc except household size but also adds an extensive set of additional variables. Among others, the additional variables again reflect economic status (highest car segment, share unemployed) and age structure. Moreover, access to and habit to use different (environmentally-friendly) means of transport as reflected in mobility portfolio and the share of respondents who use low- or zero-carbon transport on a regular basis is plausibly indicative of household VKT. Taken together, home ownership and household composition in terms of age and gender reveal extraordinary importance across outcome variables and thus for mobility behavior in general.

To sum up, the data-driven LASSO is only partially congruent with the ad hoc selection. While too many variables have been identified for EV adoption ad hoc, a substantial set of factors relevant for VKT goes unnoticed. It is yet unclear whether this indicates a weakness of the literature or whether this results from the different subsets of households considered.

In the extensive margin analysis, we identified 140 EV households for which the EV is predicted to be an ICEV substitute. For these households, we match on the union of variables identified, respectively. Based on the ad hoc variable selection, the matching algorithm finds 9838 suitable control households among all two-ICEV households. As a result of the more comprehensive list of matching variables, the LASSO-based matching algorithm finds 355 control households only. While substantial imbalances exist prior to matching, the genetic matching strongly improves balance such that the standardized difference in means is close to zero for all variables in both specifications, as Table 6

costs.

[23]This is in line with previous literature which has established a strong link between pro-environmental attitudes and intentions towards pro-environmental behavior (e.g. Bamberg et al., 2007). However, the relationship between attitudes and actual behavior is less clear due a widely observed attitude-behavior-gap (e.g. Heslop et al., 1981; Gatersleben et al., 2002).

[24]For a model using all variables selected at $c = 75$, the AUC is 0.83.

demonstrates. Figure 1 plots the distribution of household mileage by treatment status

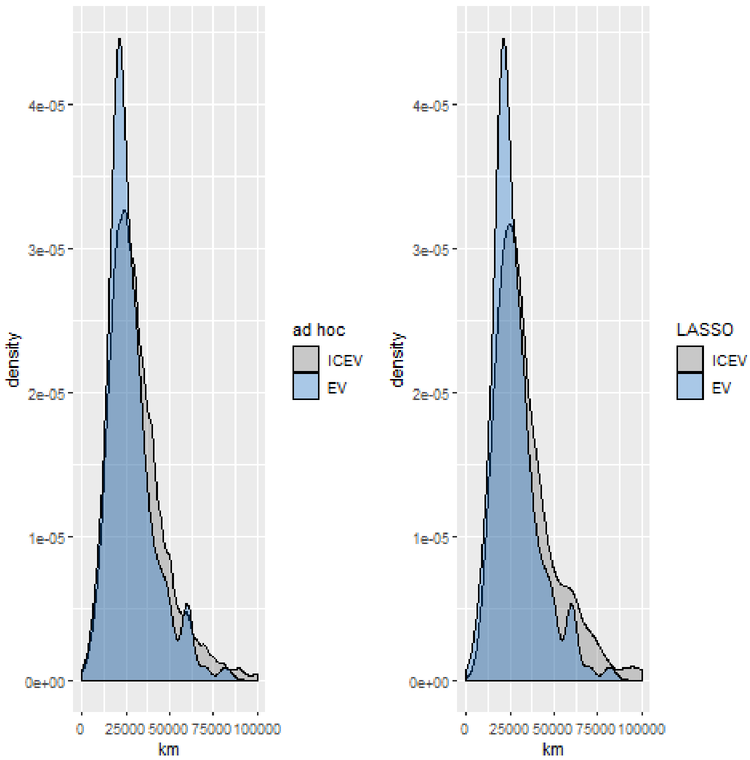

Figure 1: Household mileage by treatment status

post-matching. In both specifications household mileage is more concentrated and reveals a lower mean among EV households than among ICEV households. The regression results in Table 7 confirm this finding.[25] When controlling for all relevant confounders, household mileage is *significantly lower* among EV households than among ICEV households in both specifications. The difference amounts to -16.8% of the sample mean VKT in the ad hoc specification. In the LASSO-based specification, EV ownership is associated with a reduction in annual VKT of -23.3% compared to the sample VKT average which is even more pronounced than in the ad hoc case. Overall, we do not find statistical evidence for an intensive margin rebound effect for the subgroup of ICEV substituters irrespective of the specification. Instead, EV adoption is consistently related to significantly lower household VKT in two-vehicle households. The negative relationship may be explained by the strive for behavior consistent with the environmentally-friendly car choice. Adopting an environmentally-friendly car may induce households to reconfigure their mobility patterns, particularly to reduce the mileage traveled by car. In any case, lower household mileage post EV adoption is associated with lower GHG emissions from driving the car, irrespective of the split in mileage between the household's EV and ICEV.

However, we treat our results with caution due to the strong assumptions required for tractability. Also note that the adjusted R^2 of the LASSO-based model is overall rather low. It is even slightly lower in the LASSO-based regression than in the ad hoc model

[25]We only report the coefficient of interest due to the extensive list of control variables.

	prior	post	
		ad hoc	LASSO
household size	0.27	0.00	
no. employed	0.28	0.00	
no. adults with driving license	0.03	0.00	0.02
regional type	0.19	0.00	-0.02
home ownership	0.69	0.00	0.00
economic status	0.62	0.01	0.00
share retirees	-0.52	0.00	0.00
share adult females	-0.06	0.00	0.00
highest educational level	0.47	0.00	0.00
share adults	-0.32	0.00	0.00
income (eq)	0.37		0.01
mobility portfolio	0.25		0.04
highest car segment	-0.09		0.00
log income (eq)	0.49		0.01
share ENVI users	0.01		0.00
share unemployed	0.02		0.00

Table 6: Standardized Mean Differences prior and post-matching

	household VKT	
	ad hoc	LASSO
EV household	-5425.8^{***}	-7744.0^{***}
	(3266.2)	(2072.0)
.	.	.
R^2	0.14	0.21
Adj. R^2	0.13	0.12
Num. obs.	9978	495
RMSE	17877.60	20173.19

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 7: Regression results

which is likely an outcome of the higher number of regressors but lower sample size. Although sign and significance level are identical across specifications, the magnitude of the effect differs. The ad hoc specifications underestimates the effect by 30% compared to the LASSO-based results. Methodologically, the analysis shows that data-driven variable selection is a useful complement for literature-based variable selection: As results differ among ad hoc and LASSO selection, data-driven approaches allow the verification of confounder selection and estimation precision established ad hoc. They can be considered particularly (but not exclusively) useful when little theoretical or empirical knowledge is yet established and if the specific subsample of interest may differ from the overall population.

5 Discussion

This paper uses recent household-level micro data from the national travel survey *Mobilität in Deutschland 2017* to investigate the potential rebound effects EV adoption may induce both at the extensive and the intensive margin. In contrast to most previous studies, we were able to analyze the mobility behavior of actual EV owners based on data from a large-scale national travel survey. For a well-defined extensive margin analysis, we restrict the treatment group to the subset of households owning one EV and one ICEV which is the largest subgroup of EV households in the data set. We observe 183 treated households. At the extensive margin, the EV is predicted to represent an additional car for 24% of the EV households by means of supervised learning. We cannot conclusively assert the extensive margin rebound effect to be significant effect among the subgroup of EV owners considered. More generally, note that we only investigate a specific subgroup of EV owners. Evidently, the results may vary for other subgroups and are thus not conclusive for policy advice. Particularly, households without a car but high environmental awareness may be particularly prone to buying an EV as additional car.

For intensive margin analysis, we focus on the subset of treated households who adopted an EV as an ICEV substitute which reduces the size of the treatment group to 140 households. Technically, treatment of interest becomes the substitution of one of two household ICEVs with an EV. In a genetic matching approach, we recruit a control group for the treated households from a large set of 42,898 two-ICEV households. For matching variable selection, we contrast ad hoc variable selection based on the limited empirical and theoretical literature with a data-driven variable selection algorithm (double LASSO). Overall, we do not find a significant rebound effect of EV adoption on household mileage for the subset of EV households considered. Instead, household mileage significantly lower among EV households than in the control sample which is indicative of a strive for behavior consistent with the environmentally-friendly car choice rather than a rebound effect. Lower annual mileage is related to a decrease in GHG emissions, irrespective of how mileage is divided among the household cars. The data-driven variable selection identifies a refined set of variables and suggests a 30% higher effect size than the ad hoc selection. It can thus be considered a useful complement to literature-based selection.

However, there are several limitations to our analysis. Most importantly, we cannot rule out reverse causality. Households may select into EV ownership based on the mileage required to satisfy their mobility demand. This selection may on the one hand be driven by the lower variable cost per km. Vassileva et al. (2017) find this consideration to be a relevant aspect for EV adoption among EV owners. On the other hand, the limited driving range of EVs could play a role for car purchase decisions and household VKT. For the household VKT, we argue that the limited driving range is likely to play a minor role in the multi-vehicle households considered. Firstly, even moderate driving ranges of EVs available today are sufficient to capture the daily mobility demand of most households

(Pearre et al., 2011; Habla et al., 2020). Additionally, as we restrict the analysis to multi-engine households, between-vehicle substitution is possible if the required distance cannot be completed with the EV. Thus, EV range restrictions are not likely to affect annual household mileage substantially in the EV-owning households considered. Nonetheless, an EV's limited driving range might psychologically deter households with high mobility demand from adoption (Franke et al., 2012). If causality works in this direction, the lower household mileage found among EV households may rather reflect a selection effect instead of a treatment effect. Moreover, although the survey is large-scale overall, the number of EV households both in the sample as well as in the overall population is still low. Consequently, one can assume that EV households still constitute early adopters. As early adopters are a highly selected subsample of the overall population (Rogers, 2003), the effects established may differ from the overall population once market penetration increases and are thus not generalizable.

Several challenges for identification are posed by the temporal structure of the data and the scope of the variables contained in the survey. We rely on four crucial assumptions for which validity cannot be assessed directly. In the following, we discuss their plausibility. Due to the absence of a panel data structure, we are required to make three assumptions to recover sufficient temporal information for estimation. Assumption 1 posits that the EV is the last car the household has added to its car portfolio. This is a largely plausible assumption as EV market penetration has only taken off in recent years in Germany. Therefore, most EV households are likely to have bought the EV only recently at the time of survey in 2017 . However, it may not hold for all cases and evidently represents a generalization. As the analysis requires observables to reflect pre-treatment information, assumption 2 states that the covariates controlled for have not changed since EV adoption. If the EV was adopted long ago, it is conveivable that household characteristics may have changed since adoption. In this case, we observe household characteristics with measurement error. Given that EVs are overall a rather new phenomenon and substantial adjustments to living circumstances are likely to occur in the long-run, if at all, we argue that assumption 2 can still be defended as we mostly rely on socio-economic information about the household. However, the treatment may also directly influence the control variables. In this case, we face a problem of endogenous controls in the intensive margin analysis. Most plausibly, EV adoption may induce changes in the use of other means of transportation that are considered environmentally-friendly like cycling, public transport and walking, which would invalidate the LASSO-based matching where the share of regular users of environmentally-friendly means of transportation is selected as a matching variable. Even more critically, the lack of information on when a particular car was adopted also requires us to presume in assumption 4 that the treatment effect is constant over time. We therefore implicitly exclude learning over time or lagged effects. If the rebound effect in mileage was predominantly driven by an income effect, it seems possible that the effect does not occur instantly as households may take some time to realize the cost differences in their expenses. Lagged effects would imply that the measured effect represents a lower bound if a substantial share of treated households adopted the EV only recently. By contrast, if the effect is primarily driven by moral licensing, it does not seem implausible to assume immediate effect occurence. However, also in this case adaptation and learning over time may amplify or weaken effect size over time. All in all, the assumption of homogenous effect size over time is a strong assumption without empirical or analytical backing. The second threat to identification arises from the limited scope of variables elicited. Most importantly, only sociodemographic variables are elicited while attitudes are unobserved. In order to plausibilize the conditional independence assumption required for the matching to be valid, we assume that attitudinal variables relevant for EV adoption are captured to a sufficient extent by the given observable proxy variables (assumption 3). While the positive correlation between education and pro-environmental attitudes has been demonstrated repeatedly, the age proxy seems rather weak as it is unable to capture heterogeneity in openness towards technology within age groups. On the other hand, the list of observeable matching variables is substantive in scope so that similarity in unobservables becomes likely (cf. Altonji et al., 2005). Moreover, the evidence with respect to the role of attitudes for EV adoption primarily bases on the relationship between attitudes and intention to adopt with Vassileva et al. (2017) as an exception. If the link between attitudes and actual adoption is weak as the presence of a well-documented intention-behavior-gap may suggest, the attitudinal variables we need to proxy for may not even be relevant drivers of EV adoption. As a consequence, the quality of the proxy variables may be of limited concern. However, if we have insufficient control over relevant attitudinal drivers of both EV adoption and VKT, the results may rather evidence a selection instead of a treatment effect. Especially pro-environmental attitudes could lead to both lower mileage and higher propensity to adopt an EV (cf. Kahn, 2007). Given these caveats, we caution against strong causal interpretation.

Future research is needed to provide clear causal evidence with respect to the effects of EV adoption on car ownership, household mileage as well as on GHG emissions taking both emissions from vehicle production and operation into account. In order to address the limitations we face, micro-level data with an even more comprehensive scope and temporal dimension are required. As decarbonizing the transport sector becomes an increasingly important policy goal, highly granular data is required to allow researchers and policy advisors to develop scientific guidelines for effective policies and regulations.

Acknowledgements

We thank Beate Thies, Ulrich Wagner, Robert Germeshausen and participants at the ZEW internal seminar for valuable feedback.

References

Agora Energiewende (2018). "The European Power Sector in 2018". In: URL: https://www.agora-energiewende.de/fileadmin2/Projekte/2018/EU-Jahresauswertung_2019/Agora-Energiewende_European-Power-Sector-2018_WEB.pdf.

Altonji, Joseph G, Todd E Elder, and Christopher R Taber (2005). "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools". In: *Journal of political economy* 113.1, pp. 151–184.

Anable, J., S. Skippon, G. Schuitema, and N. Kinnear (2011). "Who will adopt electric vehicles?: A segmentation approach of UK consumers". In: *European Council for an Energy Efficient Economy*.

Bakin, Sergey (1999). "Adaptive regression and model selection in data mining problems". PhD thesis. School of Mathematical Sciences, Australian National University.

Bamberg, S. and G. Möser (2007). "Twenty years after Hines, Hungerford, and Tomera: A new meta-analysis of psycho-social determinants of pro-environmental behaviour". In: *Journal of environmental psychology* 27.1, pp. 14–25.

Belloni, A. and V. Chernozhukov (2013). "Least squares after model selection in high-dimensional sparse models". In: *Bernoulli* 19.2, pp. 521–547.

Belloni, A., V. Chernozhukov, and C. Hansen (2014). "Inference on treatment effects after selection among high-dimensional controls". In: *The Review of Economic Studies* 81.2, pp. 608–650.

Bradley, Andrew P. (1997). "The use of the area under the ROC curve in the evaluation of machine learning algorithms". In: *Pattern Recognition* 30.7, pp. 1145–1159.

Büchs, M. and S. V. Schnepf (2013). "Who emits most? Associations between socio-economic factors and UK households' home energy, transport, indirect and total CO2 emissions". In: *Ecological Economics* 90, pp. 114–123.

Chetverikov, D., Z. Liao, and V. Chernozhukov (2016). "On cross-validated Lasso". In: *arXiv preprint arXiv:1605.02214*.

Davis, L. W. (2019). "Evidence of a homeowner-renter gap for electric vehicles". In: *Applied Economics Letters* 26.11, pp. 927–932.

Davis, L. W. and J. M. Sallee (2019). "Should Electric Vehicle Drivers Pay a Mileage Tax?" In: *Environmental and Energy Policy and the Economy*. University of Chicago Press.

Diamond, A. and J. S. Sekhon (2013). "Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies". In: *Review of Economics and Statistics* 95.3, pp. 932–945.

European Environmental Agency (EEA) (2019). *Greenhouse gas emissions from Transport in Europe*. URL: https://www.eea.europa.eu/data-and-maps/indicators/transport-emissions-of-greenhouse-gases/transport-emissions-of-greenhouse-gases-12.

Frank, L. E. and J. H. Friedman (1993). "A statistical view of some chemometrics regression tools". In: *Technometrics* 35.2, pp. 109–135.

Franke, Thomas, Isabel Neumann, Franziska Bühler, Peter Cocron, and Josef F. Krems (2012). "Experiencing Range in an Electric Vehicle: Understanding Psychological Barriers". In: *Applied Psychology* 61.3, pp. 368–391.

Gatersleben, B., L. Steg, and C. Vlek (2002). "Measurement and determinants of environmentally significant consumer behavior". In: *Environment and behavior* 34.3, pp. 335–362.

Habla, Wolfgang, Vera Huwe, and Martin Kesternich (2020). *Beyond monetary barriers to electric vehicle adoption: Evidence from observed usage of private and shared cars.* Tech. rep. ZEW-Centre for European Economic Research Discussion Paper 20-026.

Held, Michael and Michael Baumann (2011). "Assessment of the environmental impacts of electric vehicle concepts". In: *Towards life cycle sustainability management.* Springer, pp. 535–546.

Heslop, L. A., L. Moran, and A. Cousineau (1981). ""Consciousness" in energy conservation behavior: an exploratory study". In: *Journal of Consumer Research* 8.3, pp. 299–305.

Hidrue, M. K., G. R. Parsons, W. Kempton, and M. P. Gardner (2011). "Willingness to pay for electric vehicles and their attributes". In: *Resource and energy economics* 33.3, pp. 686–705.

Ho, D. E., K. Imai, G. King, and E.A. Stuart (2011). "MatchIt: Nonparametric Preprocessing for Parametric Causal Inference". In: *Journal of Statistical Software* 42.8, pp. 1–28. URL: http://www.jstatsoft.org/v42/i08/.

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016). "Are there environmental benefits from driving electric vehicles? The importance of local factors". In: *American Economic Review* 106.12, pp. 3700–3729.

Kahn, Matthew E. (2007). "Do greens drive Hummers or hybrids? Environmental ideology as a determinant of consumer choice". In: *Journal of Environmental Economics and Management* 54.2, pp. 129–145. ISSN: 00950696.

Kirkland, Lisa-Ann, Frans Kanfer, and Sollie Millard (2015). "LASSO tuning parameter selection". In: *Annual Proceedings of the South African Statistical Association Conference.* Vol. 2015. Congress 1. South African Statistical Association (SASA), pp. 49–56.

Klöckner, C. A., A. Nayum, and M. Mehmetoglu (2013). "Positive and negative spillover effects from electric car purchase to car use". In: *Transportation Research Part D: Transport and Environment* 21, pp. 32–38.

Kotsiantis, S., D. Kanellopoulos, P. Pintelas, et al. (2006). "Handling imbalanced datasets: A review". In: *GESTS International Transactions on Computer Science and Engineering* 30.1, pp. 25–36.

Kraftfahrtbundesamt (2017). *Bestandsbarometer: Personenkraftwagen am 1. Januar 2017 nach ausgewählten Merkmalen.* URL: https://www.kba.de/DE/Statistik/Fahrzeuge/Bestand/Jahresbilanz/2017/2017_b_barometer.html?nn=1873496.

Lévay, P. Z., Y. Drossinos, and C. Thiel (2017). "The effect of fiscal incentives on market penetration of electric vehicles: A pairwise comparison of total cost of ownership". In: *Energy Policy* 105, pp. 524–533.

Meier, Lukas, Sara van de Geer, and Peter Bühlmann (2008). "The group lasso for logistic regression". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70.1, pp. 53–71.

Meyer, A. (2015). "Does education increase pro-environmental behavior? Evidence from Europe". In: *Ecological economics* 116, pp. 108–121.

Miller, Dale T and Daniel A Effron (2010). "Psychological license: When it is needed and how it functions". In: *Advances in experimental social psychology*. Vol. 43. Elsevier, pp. 115–155.

Munyon, V. V., W. M. Bowen, and J. Holcombe (2018). "Vehicle fuel economy and vehicle miles traveled: An empirical investigation of Jevon's Paradox". In: *Energy Research & Social Science* 38, pp. 19–27.

Nobis, C. and K. Köhler (2018). *Mobilität in Deutschland - MiD Nutzerhandbuch. Studie von infas, DLR, IVT and Infas 360 on behalf of Bundesministerium für Verkehr und digitale Infrastruktur.*

Nordic Council of Ministers (2014). *10 Insights into the Nordic energy systems.* URL: https://www.nordicenergy.org/wp-content/uploads/2018/06/10-Insights-A4.pdf.

Pearre, Nathaniel S., Willett Kempton, Randall L. Guensler, and Vetri V. Elango (2011). "Electric vehicles: How much range is required for a day's driving?" In: *Transportation Research Part C: Emerging Technologies* 19.6, pp. 1171–1184. ISSN: 0968-090X.

Plötz, P., U. Schneider, J. Globisch, and E. Dütschke (2014). "Who will buy electric vehicles? Identifying early adopters in Germany". In: *Transportation Research Part A: Policy and Practice* 67, pp. 96–109.

Rogers, E. M. (2003). *Diffusion of Innovations: 5th ed.* New York.

Rubin, D. B. (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of educational Psychology* 66.5, p. 688.

Sheeran, P. (2002). "Intention—behavior relations: a conceptual and empirical review". In: *European review of social psychology* 12.1, pp. 1–36.

Sovacool, B. K., J. Kester, L. Noel, and G. Z. de Rubens (2018). "The demographics of decarbonizing transport: the influence of gender, education, occupation, age, and household size on electric mobility preferences in the Nordic region". In: *Global environmental change* 52, pp. 86–100.

Tibshirani, R. (1996). "Regression shrinkage and selection via the lasso". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 58.1, pp. 267–288.

Urminsky, O., C. Hansen, and V. Chernozhukov (2016). *Using double-lasso regression for principled variable selection.* Tech. rep. Available at SSRN 2733374.

Vassileva, I. and J. Campillo (2017). "Adoption barriers for electric vehicles: Experiences from early adopters in Sweden". In: *Energy* 120, pp. 632–641.

Xing, Jianwei, Benjamin Leard, and Shanjun Li (2019). *What does an electric vehicle replace?* Tech. rep. National Bureau of Economic Research.

Yang, Yi and Hui Zou (2015). "A fast unified algorithm for solving group-lasso penalize learning problems". In: *Statistics and Computing* 25.6, pp. 1129–1141.

Yuan, Ming and Yi Lin (2006). “Model selection and estimation in regression with grouped variables”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68.1, pp. 49–67.

A Variable List

Variable	Characteristic	Description
household size	categorical $\in [0, 8]$	
home ownership	1	rent
	2	property
car sharing membership	0	no
	1	yes (with at least one provider)
type of household	1	young household (under 35 years)
	2	family
	3	adult-only household
	4	household with persons over 65 years
mobility portfolio	1	car
	2	car and bicycle
	3	car and carsharing
	4	car, bicycle and carsharing
	5	bicycle
	6	bicycle and carsharing
	7	carsharing
	8	without car, bicycle and carsharing
highest car segment	1	small
	2	compact
	3	mid-range
	4	large
regional type	51	metropolis
	52	regiopolis/ large city
	53	medium-sized town
	54	urban area
	55	small town, rural area
number of motorbikes/mopeds in household, grouped	0	no motorbike/moped
	1	1 motorbike/moped
	2	2 motorbikes/mopeds
	3	3 motorbikes/mopeds
	4	4+ motorbikes/mopeds
number of electric bicycles/pedelecs in household, grouped	0	no pedelec/electric bicycle
	1	1 pedelec/electric bicycle
	2	2 pedelecs/electric bicycles
	3	3pedelecs/electric bicycles
	4	4+ pedelecs/electric bicycles
number of bicyles in household, grouped	1	1 bicyle

	2	2 bicyles
	3	3 bicycles
	4	4+ bicycles
secondary residence	0	no
	1	yes
kids (< 18)	0	no
	1	yes
highest reported educational level	0	no degree (yet)
	1	lowest educational degree (Volks- or Hauptschulabschluss)
	2	medium educational degree (Realschulabschluss)
	3	A-levels
	4	university degree
Eastern Germany indicator	0	Western Germany
	1	Eastern Germany
number of employed adults	categorical $\in [0, 8]$	
number of adults with driving license	categorical $\in [0, 8]$	
share of unemployed household members	continous $\in [0, 1]$	
share of regular ENVI (bikes and public transit) users	continous $\in [0, 1]$	
share of adult household members	continous $\in [0, 1]$	
share of adult females	continous $\in [0, 1]$	
share of retired household members	continous $\in [0, 1]$	
equivalized income (OECD-scale)	continuous $\in [0.8, 9,000]$	
logarithm of eq. income		
eq. income squared		
economic status categorized	1	very low
	2	low
	3	medium
	4	high
	5	very high

B Additional Figures

Variable Selection

Variable	extensive margin	intensive margin					
		EV ownership				VKT	
		ad hoc	LASSO			ad hoc	LASSO
			60	75	90		
economic status		x				x	
economic status categorized	x		x				x
eq. income	x		x	x			
eq. income squared							
log eq. income	x						x
mobility portfolio	x						x
highest car segment	x						x
regional type	x	x				x	x
home ownership	x	x	x	x	x	x	x
share of retirees	x	x	x	x	x	x	x
share of adults	x	x	x	x	x	x	x
share of adult females	x	x	x	x			x
level of education	x	x				x	x
Eastern Germany	x						
household size	x	x				x	
no. employed adults	x	x					
no. adults with driving license	x	x				x	x
share of regular ENVI users	x		x				x
share unemployed	x						x
secondary residence	x						
type of household	x						
car sharing membership	x						
kids	x						
number of motorbikes	x						
number of bicyles	x						
number of electric bicycles	x						
no. of motorbikes							

Table 9: Variable selection

C Derivation of the Extensive Margin Hypothesis Test

We show that by testing $H_0' : \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2) = \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2)$ against $H_1' : \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,1} = 2) \neq \mathbb{E}(I(\hat{Y}_{h,0} = 2)|Y_{h,0} = 2)$ using Welch's two-sample t-test, we recover a means to test H_0*:* $P(Y_{h,1} = 2 \wedge Y_{h,0} = 1) = 0$ against H_1*:* $P(Y_{h,1} = 2 \wedge Y_{h,0} = 1) \neq 0$ in the potential outcome setup where $\forall h \in H_{EV}$ $Y_{h,0}$ and $\forall h' \in H_{ICEV}$ $Y_{h',1}$ is unknown. However, the test developed here is not equivalent to testing H_0 in the complete information setup, as there are possible scenarios in which H_0 would be rejected but H_0' would not.

On the EV sample, we can estimate $\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2)$. Under H_0:

$$\begin{aligned}
\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2) &= P(\hat{Y}_{h,0}=2|Y_{h,1}=2) \\
&= P(\hat{Y}_{h,0}=2|Y_{h,1}=2, Y_{h,0}=2)P(Y_{h,0}=2|Y_{h,1}=2) \\
&+ P(\hat{Y}_{h,0}=2|Y_{h,1}=2, Y_{h,0}=1)P(Y_{h,0}=1|Y_{h,1}=2) \\
&= P(\hat{Y}_{h,0}=2|Y_{h,0}=2, Y_{h,1}=2)
\end{aligned} \tag{4}$$

as

$$\begin{aligned}
P(Y_{h,1}=2) &= P(Y_{h,1}=2 \wedge Y_{h,0}=2) + P(Y_{h,1}=2 \wedge Y_{h,0}=1) \\
&= P(Y_{h,1}=2 \wedge Y_{h,0}=2),
\end{aligned}$$

it follows that

$$\begin{aligned}
P(Y_{h,0}=1|Y_{h,1}=2) &= \frac{P(Y_{h,0}=1 \wedge Y_{h,1}=2)}{P(Y_{h,1}=2)} = 0 \\
P(Y_{h,0}=2|Y_{h,1}=2) &= \frac{P(Y_{h,0}=2 \wedge Y_{h,1}=2)}{P(Y_{h,1}=2)} = 1.
\end{aligned}$$

The second component $\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2)$ can be recovered from the model evaluation measures on the ICEV test sample. Defining the set of households for which $Y_{h,0}=2$, the "non-event" as defined in section 3, as $H_2 = \{h \in H_{ICEV}|_{testdata} : Y_{h,0}=2\}$, this expected value can be estimated by the model's specificity defined as $Specificity = \frac{1}{|H_2|}\sum_{h\in H_2} I(\hat{Y}_{h,0}=2)$ We now derive an expression for $\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2)$ under H_0 as

$$\begin{aligned}
\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2) &= P(\hat{Y}_{h,0}=2|Y_{h,0}=2) \\
&= P(\hat{Y}_{h,0}=2|Y_{h,0}=2 \wedge Y_{h,1}=2)P(Y_{h,1}=2|Y_{h,0}=2) \\
&+ P(\hat{Y}_{h,0}=2|Y_{h,0}=2 \wedge Y_{h,1}=3)P(Y_{h,1}=3|Y_{h,0}=2) \\
&= P(\hat{Y}_{h,0}=2|Y_{h,0}=2 \wedge Y_{h,1}=2).
\end{aligned} \tag{5}$$

As under H_0,

$$\begin{aligned}
P(Y_{h,1}=3|Y_{h,0}=2) &= \frac{P(Y_{h,1}=3 \wedge Y_{h,0}=2)}{P(Y_{h,0}=2)} = 0 \\
P(Y_{h,1}=2|Y_{h,0}=2) &= \frac{P(Y_{h,1}=2 \wedge Y_{h,0}=2)}{P(Y_{h,0}=2)} = 1
\end{aligned}$$

as

$$
\begin{aligned}
P(Y_{h,0}=2) &= P(Y_{h,0}=2 \wedge Y_{h,1}=2) + P(Y_{h,0}=2 \wedge Y_{h,1}=3) \\
&= P(Y_{h,0}=2 \wedge Y_{h,1}=2)
\end{aligned}
$$

if we maintain the restriction on possible adoption channels and assume that under H_0, the addition channel occurs with zero probability also for higher pre-treatment car endowment $Y_{h,0}=2$.

Therefore, under H_0, it must hold that $\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2) = P(\hat{Y}_{h,0}=2|Y_{h,0}=2, Y_{h,1}=2) = \mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2)$.
To derive the asymptotic distribution of the two estimators $\frac{1}{|H_{EV}|}\sum_{h\in H_{EV}} I(\hat{Y}_{h,0}=2|Y_{h,1}=2)$ and $Specificity = \frac{1}{|H_2|}\sum_{h\in H_2} I(\hat{Y}_{h,0}=2)$, it is important to notice that the model for $\hat{Y}_{h,0}$ is estimated on the ICEV train sample and is thus invariant under increases of the size of the EV sample and the ICEV test sample. We can thus consider $\hat{Y}_{h,0}$ as a given random variable as the sample sizes of the EV and the ICEV test sample approach infinity. Assuming that the vector of covariates X_h is an independent and identically distributed random vector (iid), it follows that $\hat{Y}_{h,0}$ is iid as well and therefore

$$
Specificity \xrightarrow{d} \mathcal{N}(\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2), \sigma_2)
$$

and

$$
\frac{1}{|H_{EV}|}\sum_{h=1}^{H_{EV}} I(\hat{Y}_{h,0}=2|Y_{h,1}=2) \xrightarrow{d} \mathcal{N}(\mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2), \sigma_1)
$$

Thus, if the hypothesis $H_0' : \mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2) = \mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2)$ is rejected by Welch's two-sample unpaired t-test, we know that also H_0 cannot hold. Testing H_0' however is not equivalent to testing H_0, as there might be a scenario in which H_1 $P(Y_{h,1}=2 \wedge Y_{h,0}=1) \neq 0$ holds true, but at the same time $H_0' : \mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,1}=2) = \mathbb{E}(I(\hat{Y}_{h,0}=2)|Y_{h,0}=2)$ holds true. A possible such scenario would be given by

$$
\begin{aligned}
&P(Y_{h,1}=2|Y_{h,0}=2) = P(Y_{h,0}=2|Y_{h,1}=2) \\
&P(Y_{h,1}=3|Y_{h,0}=2) = P(Y_{h,0}=1|Y_{h,1}=2) \\
&P(\hat{Y}_{h,0}=2|Y_{h,1}=2, Y_{h,0}=1) = P(\hat{Y}_{h,0}=2|Y_{h,0}=2 \wedge Y_{h,1}=3).
\end{aligned}
$$

This is an illustrative examples to show that the power of testing Hypothesis 1 using H_1' is clearly lower than the test in the full information setup. Overall, the test we develop is not entirely conclusive about the effect's significance level.

// IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de