// NO.22-059 | 11/2022

DISCUSSION PAPER

// JANNA AXENBECK, ANNE BERNER, AND THOMAS KNEIB

What Drives the Relationship Between Digitalization and Industrial Energy Demand? Exploring Firm-Level Heterogeneity

Leibniz Association

ZEW

What drives the relationship between digitalization and industrial energy demand? Exploring firm-level heterogeneity

Janna Axenbeck[a,*], Anne Berner[b], Thomas Kneib[b]

[a] *ZEW Mannheim and Justus-Liebig-University Giessen, Germany*

[b] *Chair of Statistics and Campus Institute Data Science, University of Göttingen, Germany*

Abstract

The ongoing digital transformation has raised hopes for ICT-based climate protection within manufacturing industries, such as dematerialized products and energy efficiency gains. However, ICT also consume energy as well as resources, and detrimental effects on the environment are increasingly gaining attention. Accordingly, it is unclear whether trade-offs or synergies between the use of digital technologies and energy savings exist. Our analysis sheds light on the most important drivers of the relationship between ICT and energy use in manufacturing. We apply flexible tree-based machine learning to a German administrative panel data set including more than 25,000 firms. The results indicate firm-level heterogeneity, but suggest that digital technologies relate more frequently to an increase in energy use. Multiple characteristics, such as energy prices and firms' energy mix, explain differences in the effect.

Keywords: digital technologies, energy use, manufacturing, machine learning.
JEL Codes: C14, D22, L60, O33, Q40.

*Corresponding author
Email addresses: Janna.axenbeck@zew.de (Janna Axenbeck), anne.berner@posteo.de (Anne Berner), tkneib@uni-goettingen.de (Thomas Kneib)

1. Introduction

The growing number of applications as well as the rapidly evolving performance of information and communication technologies (ICT) have raised hopes of increasing productivity while simultaneously reducing greenhouse gas emissions and energy use (Kander et al., 2015; IEA, 2019). New digital technologies, such as smart sensors and advanced data analytics, offer the opportunity to make energy use more efficient and thus save resources. However, as more and more digital devices are produced, used, and disposed, negative environmental impacts are increasingly being scrutinized (Williams, 2011; Andrae and Edler, 2015; Belkhir and Elmeligi, 2018; Lange et al., 2020). Thus, whether the ongoing digital transformation contains synergies or trade-offs between technological progress and environmental benefits is heavily debated. The use of digital technologies may also involve both, as effects are potentially context-dependent, i.e., heterogeneous. Despite this ambiguity and the need for more research, policies that propose pathways towards a sustainable economy consider digitalization as a key element in lowering environmental impacts. To enable a "green and digital transition", ICT shall support the decrease in energy use and the decarbonization of the energy mix (European Commission, 2021).

In general, selective targeting of digitalization that relates to lower levels of non-renewable energy demand may allow greater progress toward climate targets. Since the manufacturing sector is responsible for a large share of global CO_2 emissions, improving its environmental footprint is especially important. For example, manufacturing industries accounted for 26 % of global CO_2 emissions and for 38 % of global energy use in 2020 (IEA, 2021). Accordingly, associated industries receive special attention in policies that address the role of digital technologies in achieving climate targets at the European level (European Commission, 2019, 2020, 2021). In practice, however, a considerable share of European industrial digitalization policies also include subsidies and funding for small and medium-sized enterprises (SMEs) as well as for regions that are considered structurally weak. As digital technologies can also spark energy-relevant output growth (Lange et al., 2020), it is important to better understand the relationship between digital technologies and energy use and how it varies across firm and structural characteristics. For instance, the amount of energy consumed in a digitalized production process may depend on the industry association, such as the chemical or the automotive industry. Moreover, market characteristics such as market concentration or price levels can impact relationships. For example, less competition lowers firms' incentives to save costs, and this potentially influences a firm's willingness to accept an increase in energy costs due to ICT usage. Profound insights in this regard may enable policymakers to improve the synchronization of industrial and climate policy objectives.

Previous empirical studies on the relationship between ICT and environmental impacts mainly attempt to prove a linear and directional link between ICT and CO_2 emissions (Zhang and Liu, 2015; Chen et al., 2019; Kopp and Lange, 2019) or energy efficiency and energy use (Collard et al., 2005; Bernstein and Madlener, 2010; Schulte et al., 2016; Axenbeck and Niebel, 2021; Huang et al., 2022). However, as the relationship may be heterogeneous or non-linear (Ben Lahouel et al., 2021; Taneja and Mandys, 2022; Xu et al., 2022), standard regression models fall short of fully uncovering the complexity of the relationship. In view of effect heterogeneity, non-parametric econometric methods have the potential to provide more detailed insights.

In the article at hand, we analyze firm-level differences in energy demand with respect to the use of digital technologies in manufacturing. We aim to reveal effect heterogeneity by applying a non-parametric, flexible tree-based algorithm, which is called the Generalized Random Forest (GRF) algorithm (Athey et al., 2019). Related studies already demonstrate the usefulness of tree-based algorithms to analyze heterogeneous relationships (Davis and Heller, 2017; Johnson et al., 2020; Knaus et al., 2021), and apply them to evaluate environmental outcomes. For instance, previous literature identifies nonlinear relationships with respect to the introduction of new pricing schemes for households: Valente (2021) analyzes waste prices in Italian municipalities and O'Neill and Weeks (2019) as well as Prest (2020) focus on time-of-use electricity prices. Miller (2020) adapts the algorithm to panel data in order to analyze the temporal effect of exposure to environmental policies. Recent literature also uses the GRF to provide recommendations to improve environmental policies. For example, Knittel and Stolper (2021) show that reactions to behavioral nudges toward household energy conservation are heterogeneous and that selective targeting, i.e., treating households that are more likely to have desired treatment effects, can increase social benefits. Besides, Burlig et al. (2020) demonstrate the usefulness of machine learning techniques when forecasting counterfactual energy consumption.

We contribute to the GRF literature as well as to studies on the effects of digitalization on environmental outcomes by analyzing an extensive administrative panel data set on German manufacturing firms (AFiD data) for the years 2009 to 2017. We perform an analysis of subgroups using the GRF algorithm combined with R-learning (Nie and Wager, 2021) and a difference-in-difference approach to exploit the panel structure of our data. This procedure accounts for sources of self-selection and considerably reduces potential endogeneity issues. Allowing for heterogeneous effects of observables across decision-making units, this method enables the identification of specific firm-level and external characteristics that influence energy demand. For instance, we evaluate whether current industrial digitalization policies that especially involve subsidies and funding for small and medium-sized enterprises (SMEs) and for regions that are considered structurally weak reduce or foster energy demand. Besides, previous microeconometric studies on the relationship between digitalization and energy use solely focus on energy intensity or specific energy carriers. Hence, we also extend previous literature by firstly analyzing ICT-related changes in overall energy demand at the firm level.

Our results confirm a heterogeneous relationship, but generally indicate a trade-off between the use of digital technologies and absolute energy savings (for the majority of firms). On average, an increase in the firm-level degree of digitalization, which is approximated by a binary variable that is one if software capital rises and zero otherwise, relates to a simultaneous rise in energy use of 1.03 %. Analyzing electricity use and non-electric fossil fuel use separately reveals that the magnitude of the effect is even larger for electricity use (1.34 %), yet we do not find a significant effect for fossil fuel use, and the respective point estimate is close to zero. Thus, results suggest that the overall increase is driven by an intensified use of electricity, which is intuitive as digital technologies consume electric power. Linking this finding to policy objectives, contrary to expectations, digital technologies do not appear to largely decrease firm-level energy use, but are related to an increase in that energy source that is potentially renewable and thus allows for decarbonization.

Moreover, we identify multiple characteristics that explain heterogeneity. For instance, an increase in market concentration is associated with a higher rise in energy use. Also, we observe that sensitivity to the electricity price (and price policies) decrease for digital firms. A subgroup analysis additionally reveals that smaller firms in structurally weak regions show higher average growth in energy use than larger firms in regions that are considered economically strong. Therefore, the results also indicate a policy trade-off between lowering energy use and supporting technological progress in firms with a need for economic assistance.

The remainder of this paper is organized as follows: The next section deciphers the link between energy use and digitalization in the light of the current literature. Section 3 explains our empirical strategy with a focus on the Generalized Random Forest methodology to measure heterogeneous relationships. Our empirical analysis relies on an extensive administrative firm-level panel data set that will be described in Section 4. Section 5 presents and discusses the main results, while Section 6 discusses the robustness of our results and Section 7 concludes.

2. Digitalization and Energy Use in Manufacturing

In economic literature, the introduction of digital technologies is usually linked to productivity improvements, for example, due to increased process efficiency and the optimization of work practices (Brynjolfsson and Hitt, 2000; Brynjolfsson and McAfee, 2011; Cardona et al., 2013). Additionally, more and more studies focus on the environmental impacts connected to digitalized production processes, in particular, on the effect on energy consumption. In this context, the literature identifies four impact channels that drive or mitigate the overall effect on energy demand. At the economy-wide level, these transmission channels can be characterized by the following keywords: (1) direct effects, (2) economic growth, (3) energy efficiency, and (4) sectoral change (Lange et al., 2020).

Direct effects comprise the energy that is embodied in the production, usage, and disposal of ICT and lead to an increase in energy demand (Williams, 2011). The same holds for the second channel, which subsumes that digital technologies can act as a multiplier for *economic growth*. Subsequently, the resulting enhanced consumption of products and services can increase energy use indirectly (Belkhir and Elmeligi, 2018; Lange et al., 2020). The third channel implies that *energy efficiency* improvements may lower energy intensity. Especially, gray literature assigns high climate protection potentials to the application of ICT. For instance, GeSI & Accenture (2015) state that digital technologies could abate 2.7 Gt of CO_2 emissions by 2030 in manufacturing industries.[1] This is asserted because, for example, industrial control systems allow for an improved fault-detection, which potentially reduces per-unit energy and resource consumption as well as wastage (Berkhout and Hertin, 2001; Baer et al., 2002). Also, simulation methods and 3D printing can considerably decrease the environmental footprint during product design and engineering processes (OECD, 2017). More generally, Berkhout and Hertin (2001) identify five areas

[1] It should be noted here that the study is financially related to telecommunication companies.

in which ICT can lower relative energy use: a) Simulation of production processes, b) intelligent design and operation of products and services, c) intelligent distribution and logistics, e.g., supply chain efficiency or alternative distribution structures, d) changing seller-buyer relationships, e.g., mass customization, and e) work organization, e.g., teleworking. However, Lange et al. (2020) point out that the desired effects of energy efficiency improvements on energy demand can be mitigated by rebound effects. These describe the energy-increasing consequences that might be triggered by energy efficiency improvements and lead to a situation where potential savings will not be fully realized (c.f. Khazzoom, 1980; Gillingham et al., 2016). Last but not least, *sectoral change*, i.e., tertiarization, relates to a shift to a more service-oriented economy. For instance, software-based solutions do not need to be physically manufactured and thus potentially require less energy and capital.

In a nutshell, ICT directly consume energy and stimulate economic growth, which can increase energy use indirectly, but digital technologies can also foster energy-efficient manufacturing as well as the dematerialization of goods. Consequently, their usage may have simultaneous positive and negative impacts on energy use, and the respective net environmental impact is a priori ambiguous from a theoretical perspective.

The wide range of mechanisms may explain why it is still under debate whether digital technologies increase or decrease energy use. Studies that find synergies between energy savings and ICT highlight that the energy mix, sector association, production factors and regional characteristics may play a part: Analyzing ten OECD countries, Schulte et al. (2016) conduct a parametric econometric analysis at the sectoral level and confirm that reductions in relative energy demand can be linked to ICT usage. They highlight that relative demand decreases in particular for non-electric energy, while relative demand for electric energy is not significantly affected. Accordingly, the relationship may depend on the energy source. Bernstein and Madlener (2010) find mixed results with respect to the effect of computers and software on relative electricity demand for European manufacturing industries. They state that the sign of the effect depends heavily on the involved sector-specific production processes. Applying quantile regression, Taneja and Mandys (2022) find a reduction in relative energy demand, but the magnitude of the reduction varies depending on the level of energy intensity. Focusing on industrial robots, as well as considering 38 countries and 17 manufacturing industries, Wang et al. (2022) find energy intensity improvements due to robot usage.[2] A closer look at the mechanism reveals that the level of energy use is barely affected, while output increases in response to the intensified use of robots. Thus, the authors do not find absolute environmental improvements. In addition, their results indicate effect heterogeneity with respect to labor and capital intensity. Using a compound index to measure digitalization, Xu et al. (2022) find reductions in absolute energy use and improvements in the share of renewable energy in total energy at the country-level. They also show that effects are mediated by technological innovation and are more pronounced in low-income countries. Therefore, they postulate heterogeneous effects with respect to regional characteristics. Majeed (2018) confirms diverging effects of ICT on CO_2 emissions between developed and developing countries. Moreover, applying a nonlinear model Ben Lahouel et al. (2021) find that ICT have increased carbon efficiency in Tunisia within the last decades.

In contrast to these rather optimistic findings, other studies indicate a trade-off between environmental outcomes and technological progress. Ren et al. (2021) find that internet development can be linked to an increase in energy use per capita in China. Sadorsky (2012) measures that digital technologies are positively linked to an increase in electricity consumption in emerging economies. Covering 93 countries over the period 1995–2016, Alataş (2021) confirms that ICT increase CO_2 emissions at the country level.

Econometric evidence at the firm-level is scarce. To the best of our knowledge, no econometric study to date examines absolute energy use in the manufacturing sector, yet empirical evidence exists with respect to changes in energy intensity. Applying a large administrative panel data set on manufacturing firms, Axenbeck and Niebel (2021) observe only marginal average energy intensity improvements related to software usage. Besides, the authors find that, even though overall effects are small, the relationship is more pronounced in energy-intensive industries, which indicates effect heterogeneity with respect to different production processes. Applying propensity score matching and focusing on the effect of industrial robots on coal consumption, Huang et al. (2022) find improvements in coal intensity. However, as described in Wang et al. (2022) above, the origin of the improvements is mainly an increase in output. A study conducted by Wen et al. (2021) focuses on environmental pollution measured by chemical oxygen demand (COD) and sulfur dioxide (SO_2). The authors find that an increase in ICT investments and services at the provincial-city level relates to a significant firm-level reduction of pollutants. On the contrary, a

[2] Energy intensity denotes the ratio between energy demand and output.

study conducted by Brozzi et al. (2020) states that firms seldom consider digital improvements (summarized under the term ”Industry 4.0”) beneficial for environmental targets but pursue predominately economic opportunities in this regard. A questionnaire-based survey with 1,700 German firms indicates diverging effects. According to non-technical self-assessments, 65 % of all surveyed manufacturing firms said that their ICT-related energy use remained constant during the last three years, 22 % stated it decreased, and 13 % mentioned an increase (Bertschek et al., 2020).

To sum up, previous studies on the relationship between digital technologies and energy use show ambiguous results. One reason for different study outcomes could be that parallel mechanisms might lead to heterogeneous and diverging effects of ICT on the environment. In this vein, Berkhout and Hertin (2004, p. 903) argue for moving “beyond the dichotomy between pessimism and optimism” to recognize that the relationship between ICT and energy is “complex, interdependent, deeply uncertain and scale-dependent”. It is an urgent political task to create the conditions for placing digitalization at the service of sustainable development. To optimally use the potential of digital technologies for climate protection, Lange et al. (2020) argue that fields of application or mechanisms with a positive environmental impact should be promoted without favoring effects that have negative environmental impacts. Horner et al. (2016, p.16) also conclude from their review study that a “focus on identification of important parameters driving the energy use in ICT-infused systems” is important in future research studies. Our contribution is to identify certain characteristics that moderate effects of ICT on energy use in manufacturing.

3. Methodology

The literature review shows that identifying the role and importance of ICT for energy use is a complex endeavor. Accordingly, the identification of characteristics that moderate energy consumption in digitalized production processes by applying a linear OLS model would quickly result in estimating too many interaction coefficients. Interpreting all of them would get soon out of hand and hardly be useful from a scientific perspective (Prest, 2020; Gulen et al., 2021). As a consequence, we apply a flexible tree-based algorithm, which is suitable to measure complex nonlinear relationships. Our estimation approach builds on the Generalized Random Forest (GRF) algorithm (Athey et al., 2019), which is a non-parametric modeling approach that allows us revealing heterogeneity and uncovering subgroup differences by applying the potential outcome framework (Rubin, 1974). In particular, we are interested in the three following questions:

1. What role do digital technologies generally play for energy consumption in manufacturing firms?
2. Which firm-level and external characteristics relate to heterogeneity?
3. To what extent does current targeting of industrial digitalization policies influence energy use?

3.1. Measuring Heterogeneous Relationships

In order to capture the effect that digital technologies may have on energy use, we compare a sample of $i = 1, \ldots, n$ firms F over a time period of $t = 1, \ldots, T$ years. For each firm, we define a binary variable $W_{i,t} = \mathbb{1}\{\Delta\, D_{i,t} > 0\}$ that indicates whether the firm i increases its use of digital technologies D in period t or not. As we follow a method that has its origin in the causal inference literature, we consider firms for which $W = 1$ as "treated" and firms for which $W = 0$ as "untreated" or "control group". Our variable of interest is energy consumption, $Y_{i,t}$, at time t. We denote the potential energy consumption of a firm that increases its use of digital technologies in period t as $Y_{i,t}(W_{i,t} = 1)$ and the corresponding energy consumption that we would have observed if the firm had not increased its use of digital technologies as $Y_{i,t}(W_{i,t} = 0)$. We define the expected difference between the two potential energy outcomes as the average treatment effect (ATE) τ. If we additionally condition on different covariates $X_{i,t} = x$, we receive the conditional average treatment effect (CATE), which is formally defined as (Athey and Wager, 2019):

$$\tau(x) = \mathbb{E}\left[Y_{i,t}(W_{i,t} = 1) - Y_{i,t}(W_{i,t} = 0) \mid X_{i,t} = x\right]. \tag{1}$$

3.2. Generalized Random Forests

A promising method to reveal these heterogeneous treatment effects from observational data is the causal forest algorithm (Wager and Athey, 2018; Knaus et al., 2021). While the name promises to automatically determine causal relationships, in fact it allows the measurement of high-dimensional interaction. The causal forest is a special case of the GRF approach introduced by Athey et al. (2019). This approach builds on the recursive partitioning, sampling, and split selection of the random forest algorithm (Breiman, 2001), an aggregation method applied to decision trees, i.e., classification and regression trees (CART). The goal is to predict an outcome

$\hat{y}$ using a non-parametric function of splitting variables, for instance, various covariates. Within one decision tree, the sample is recursively split into subgroups, optimizing the accuracy of the prediction. If a further split does not result in accuracy improvements, we call the subgroup at this node a final “leaf” of the tree.

Variation, and hence, decorrelation between decision trees is achieved, on the one hand, by basing each tree on a subsample $\mathcal{S}_b$ of the entire data set (bagging), and on the other hand, by choosing a random subset of all possible covariates to build each tree. This procedure also allows for out-of-bag predictions. Hence, we only consider trees where $i \notin \mathcal{S}_b$ to determine relationships and predict $\hat{y}^{-i}(X_{i,t})$ (Athey and Wager, 2019). This encounters problems, when working with panel data, as a firm constitutes a cluster of observations. This means that we have to exclude trees containing the same observation, i.e., firm i at period t, and trees including the same firm i at period $t+s$ to avoid information leakage.

To account for the clustered structure of our data when drawing subsamples for each decision tree, we manipulate the sampling of observations as follows (Athey et al., 2020): Instead of directly drawing $\mathcal{S}_b$, we first sample clusters J_b from $\{F_1, ..., F_n\}$. Based on each sampled J_b, we then draw k observations to build each tree.

The ensemble method applied to single trees can be described as a data-adaptive kernel method and formulated by the following when considering clusters:

$$\hat{y}(x) = \sum_{i=1}^{n} \sum_{t=1}^{T} \alpha_{i,t}(x) Y_{i,t}, \quad \alpha_{i,t}(x) = \frac{1}{B} \sum_{b=1}^{B} \frac{\mathbb{1}\left(\{X_{i,t} \in L_b(x), F_i \notin J_b\}\right)}{|\{i : X_{i,t} \in L_b(x), F_i \notin J_b\}|}, \tag{2}$$

where B indicates the number of “grown” trees, indexed by $b = 1, \ldots, B$, $L_b(x)$ is the leaf of the b-th tree containing test point x. Accordingly, $\alpha_{i,t}(x)$ indicates how often an observation falls in the identical leaf as x and it can be used to calculate a weighted average of $Y_{i,t}$ based on the forest-based adaptive neighborhood of x.

The weighting procedure is one of the main building blocks of the “Generalized Random Forest” framework (Athey et al., 2019). It is implemented in the `grf` package in `R`, on which we base our analysis.

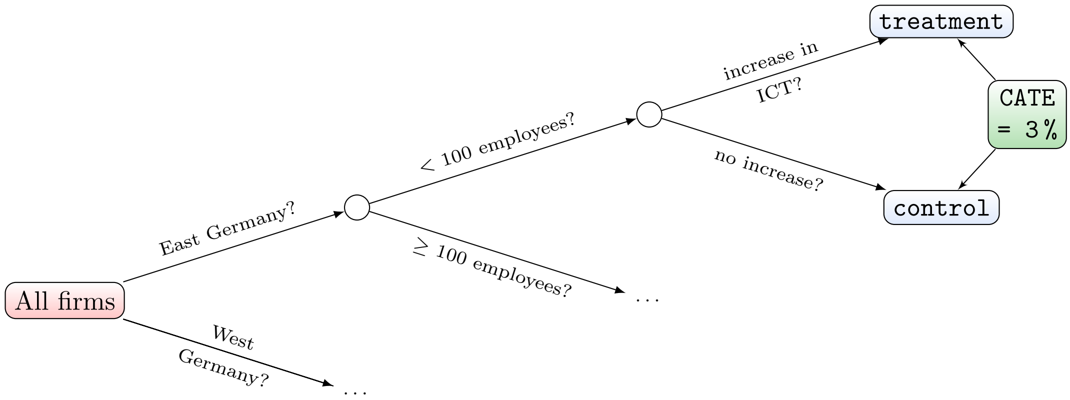

Figure 1: **Illustration of causal forest partitioning.** The conditional average treatment effect (CATE) is calculated by comparing the effect of an increase of digital technologies between firms within groups of similar firms.

The causal forest algorithm aims to predict treatment effects $\hat{\tau}$, which denote the difference between treated and untreated observations within leaves. Accordingly, splits are conducted by maximizing treatment effect heterogeneity. Nevertheless, the work horse of the algorithm remains a decision tree. See Figure 1 for a graphical illustration of a respective causal tree. The sample is split at each node recursively into two child nodes according to the covariates that maximize the discrepancy between the subgroup ATE. Unequal child node sizes are penalized. Final nodes report the estimated ATE conditional on the covariates that were responsible for the splitting, which is also known as CATE (Athey et al., 2019).

The size of our database allows us to follow an “honest” estimation procedure, which means that we split the firm panel in two groups: With the first half of the sample, we build the tree structure to calculate weights. Based on these weights, we use the second half of the training sample to estimate CATEs. Next to the decorrelation of single decision trees, this procedure prevents overstating the goodness of fit (Athey and Imbens, 2019).

For the analysis, we grow a forest of 10,000 trees. In addition to the size of the sample and the covariates used, the forest estimation is also influenced by the maximum split imbalance (between treatment and control in the child-node) and the minimum node size (minimum number of observations in a final leaf). We tune all parameters by using cross-validation.[3]

3.2.1. Identifying Assumptions

Since it is not possible to observe both, firm i increasing its use of digital technologies and not increasing its use in period t, we need the following additional assumptions to accurately estimate Equation (1). We refer here to Knaus et al. (2021):

A.1 Overlap: $0 < \mathbb{P}[W_{i,t} = 1 \mid X_{i,t} = x] < 1, \qquad \forall x \in [0,1]^d$.

A.2 Unconfoundedness: $\{Y_{i,t}(1), Y_{i,t}(0)\} \perp W_{i,t} \mid X_{i,t}$.

A.3 Exogeneity of covariates: $X^1_{i,t} = X^0_{i,t}$.

A.4 Stable Unit Treatment Value Assumption (SUTVA): $Y_{i,t} = W_{i,t}Y^1_{i,t} + (1 - W_{i,t})Y^0_{i,t}$.

The first assumption requires that no subgroup of firms defined by the covariates $X_i = x$ is located in either the treatment or the control group only, which implies that the (inverse) treatment probability must be bounded away from zero and one. The second assumption ensures that potential outcomes are independent of the treatment status, conditional on the covariates. The third assumption imposes that covariates are not affected by the treatment. The fourth assumption requires that there is no interference or no spillover effects between treated and untreated observations.

In our analysis, all assumptions might be challenged. For instance, selection effects are induced if investments in digital technologies correlate with specific firm characteristics (Athey and Wager, 2019; Gulen et al., 2021). To provide an example, firms that generate more output might consume more energy and have a higher probability to invest in digital technologies. This phenomenon may result in confounding effects and also increase the difficulty to identify counterfactual observations for these firms.

To ensure a more substantial degree of overlap, we trim our sample and only use observations which have propensity scores that match the counterfactual group (Dehejia and Wahba, 1999, 2002). For instance, we drop all observations in the group that does not increase its use of

[3]See Athey and Imbens (2019) for details.

digital technologies with an estimated propensity score lower (larger) than the smallest (highest) estimated propensity score in the group that increases its use of digital technologies and vice versa.

Second, we have to ensure unconfoundedness. In Section 3.2.2 and 3.2.3, we describe how we improve robustness to confounding by employing orthogonalization and exploiting the panel structure of our data.

Since the use of digital technologies can, in addition to energy use, influence other production function inputs, such as tangible capital, labor as well as output, the assumption of exogeneous covariates might also be violated. To solve this issue, we refrain from including critical variables measured concurrently in the same period as the treatment status. Instead, we incorporate them in lagged levels. This procedure allows for the consideration of these variables without risking that the assumption of exogeneity of covariates is violated.

We cannot assume with certainty that the fourth assumption of Stable Unit Treatment Values (SUTVA) is fulfilled a priori. Network effects are a potential argument why the SUTVA may not be valid. For instance, digital technologies can improve the efficiency of entire supply chains and alternate distribution structures. We expect network effects to be most strongly pronounced between subsidiaries within a firm. As we consider companies and not plants as one observational unit, we are able to integrate these type of network effects. However, also network effects between companies are possible, which is an issue that we will not solve but acknowledge here.

3.2.2. Orthogonalization

The assumption of independent assignment of treatment conditional on the features X is important for unbiased estimates (Assumption A.2). To fulfill this assumption, we orthogonalize treatment and outcome variables by regressing X on Y and W, and then subtracting predictions (Robinson, 1988; Nie and Wager, 2021). This procedure allows differencing out the variation in outcome and treatment due to covariates. To this end, we train separate random forests to compute estimates of propensity scores $e(x) = \mathbb{P}\left[W_{i,t} \mid X_{i,t} = x\right]$ and expected outcomes $m(x) = \mathbb{P}\left[Y_{i,t} \mid X_{i,t} = x\right]$. This approach is also known as R-learning or local centering.

The $(-i)$-superscript in this case stands for leave-one-out estimates, indicating that the i-th observation was not used to compute, e.g., $\hat{m}^{(-i)}\left(X_{i,t}\right)$. The resulting residualized outcome $(Y - m(x))$ and treatment $(W - e(x))$ variable, as well as the weights are combined in the estimation. Hence, treatment effects are estimated by solving the following equation:

$$\hat{\tau} = \frac{\sum_{i=1}^{n} \sum_{t=1}^{T} \alpha_{i,t}(x) \left(Y_{i,t} - \hat{m}^{(-i)}\left(X_{i,t}\right)\right) \left(W_{i,t} - \hat{e}^{(-i)}\left(X_{i,t}\right)\right)}{\sum_{i=1}^{n} \sum_{t=1}^{T} \alpha_{i,t}(x) \left(W_{i,t} - \hat{e}^{(-i)}\left(X_{i,t}\right)\right)^2}. \tag{3}$$

Table 1 summarizes the main steps of the causal forest algorithm including orthogonalization and honesty.

Table 1: **Summary of the steps of the causal forest algorithm with orthogonalization and honesty.**

1.	Regress $W_{i,t}$ on $X_{i,t}$ to obtain a prediction model for $\hat{e}^{(-i)}(X_{i,t})$.
2.	Regress $Y_{i,t}$ on $X_{i,t}$ to obtain a prediction model for $\hat{m}^{(-i)}(X_{i,t})$.
3.	With the first half of the sample generate causal trees but replace $W_{i,t}$ and $Y_{i,t}$ with $W_{i,t} - p^{(-i)}\left(X_{i,t}\right)$ and $Y_{i,t} - m^{(-i)}\left(X_{i,t}\right)$. Then calculate $\alpha_{i,t}$ as in Equation (2).
4.	Use the second half of the sample and weights obtained in step 3 to calculate $\hat{\tau}(x)$ by solving Equation (3).

3.2.3. Panel Structure

To reduce confounding due to unobservable characteristics, which can either be time-invariant or time-varying, we exploit the panel structure of our data. Firstly, similar to Athey et al. (2020) and Knittel and Stolper (2021), we take first differences from our outcome variable as well as from control variables to remove individual fixed effects.[4] This enables the elimination of a potential time-invariant omitted variable bias. Secondly, in the spirit of Prest (2020), Knittel and Stolper (2021), and Valente (2021), we additionally include a lagged outcome variable, to reduce possible time-varying confounding due to unobservables. The intuition why this lowers potential confounding is that the outcome from the previous period may be influenced by the same unobservables as current firm characteristics (Lechner, 2015). In other words, conditioning on pre-treatment outcomes allows controlling for previous behavior that might motivate investment in ICT.

4. Data

4.1. Microdata on the German Manufacturing Sector

Our analysis builds on firm-level data on the German manufacturing sector (AFiD[5]) collected by the Research Data Centres of the Statistical Offices of the Federation and the Federal States

[4]Note here that our treatment is also dichotomized based on the growth rate of ICT usage.
[5]"Amtliche Firmendaten für Deutschland"

(RDC) between 2009 and 2017 (Research Data Centre of the Statistical Offices Germany, 2022). We combine two different AFiD data sources: (1) The AFiD-Panel Industrial Units and (2) the AFiD-Module Use of Energy with additional information such as energy prices and deflators.[6]

Our final panel contains annual information on German manufacturing firms with at least 20 employees at the firm level (yielding around 90,000 observations in total). The longitudinal data set covers basic information about production value, employees, wages as well as details on production function inputs (e.g., machines and resources). Most importantly, it contains information about energy use, the related energy sources, and software investments.

Even though our data set covers an extensive set of firms, it is a rolling window survey, which means that not every firm is participating in the survey every year. This makes it difficult to assess whether firms are exiting or entering the market, and therefore aggregated effects at the sectoral level cannot be assessed properly. However, as our analysis concentrates on the firm level, this is only a minor limitation.

We base our variable selection on an energy consumption model that is introduced in the next subsection, before we describe the main variables of the analysis in Section 4.3.

4.2. Modeling Energy Demand

To select relevant variables for our estimation, we model observed energy use Y^* as follows:

$$Y^* = Y\left(Q, p_E, p_M, K, L, D, \vartheta(X,t)\right). \tag{4}$$

Energy demand Y is a function of a given level of output Q. In addition to energy, Q is generated by the inputs of tangible capital, labor, and materials. Energy and materials are treated as flexible inputs. Therefore, they are included via the price for energy p_E and materials p_M. In contrast, we consider labor L and tangible capital K as quasi-fixed inputs. Thus, they are directly integrated in the energy demand function. Following Stiroh (2002) and Schulte et al. (2016), we additionally augment the energy demand function by the firm-level degree of digitalization D. In addition, energy demand depends on the level of energy efficiency, which is approximated by ϑ. Firm characteristics X influence energy efficiency, such as the use of different energy sources or investments in R&D as well as market attributes, such as the competitive situation or existing regulations (cf. Porter and Van der Linde, 1995). Also, energy efficiency improvements certainly

[6]The dataset is also used and described in detail in Axenbeck and Niebel (2021)

depend on disembodied technological change, which can be captured by time t. Thus, energy efficiency is formulated as a function of X and t.

4.3. Variable Description

In this section, we briefly characterize the variables included in the analysis. Unless stated explicitly, first differences are taken. Please find an overview of all employed variables in Table 2 and a detailed description of the variables in Appendix A. We provide detailed descriptive statistics in Appendix B.

We look at three different outcomes of interest (denoted by Y): energy use, electricity use, and non-electric energetic fossil fuel use (hereafter abbreviated by fossil fuel use). Energy use represents the sum of consumed energy sources (renewable and fossil, e.g., natural gas or biomass) plus electricity consumption. All variables are measured in kWh and are log-transformed.[7]

The degree of firm-level digitalization D is approximated via a software capital stock. We consider software capital to be a suitable indicator for firm-level ICT usage, as it is a precursor to almost all digital hardware. In particular in manufacturing, technologies that optimize production processes usually require additional software. The monetary measurement of the software capital stock makes it easy to compare the proxy across different sectors and provides a certain generality in contrast to investments in single technologies, such as Cloud Computing or robotics.[8] Not without reason, it is a commonly used indicator at the firm level (cf. Almeida et al., 2020; Bessen and Righi, 2020; Axenbeck and Niebel, 2021; Barth et al., 2022). We also integrate a tangible capital stock K. We calculate both capital stocks by applying the perpetual inventory method (PIM), which allows generating a productivity-relevant capital stock (c.f. Griliches, 1980; Lutz et al., 2017). For this purpose, we use deflated investments. Moreover, we base the calculation of software capital on information on software investments, while tangible capital is approximated using information on investments in property, plants, and equipment.

Furthermore, we take first differences of the software capital stock. Based on this transformation, we define a binary treatment indicator W that approximates an increase in the use of digital technologies. Accordingly, the indicator is one if firm i shows an increase in software capital in the year t and zero otherwise. We include a dichotomized treatment indicator, as a continuous

[7]Note here that electricity consumption and fossil fuel use do not sum up to energy use, since non-electric non-fossil energy, such as biomass, cannot be accounted to either of the two.

[8]For a detailed description of the capital stock approximation as well as for a descriptive analysis of the software capital indicator, we refer to Axenbeck and Niebel (2021).

Table 2: **Variable overview.**

variable	description	variation	transformation
Outcome			
Y	energy use	firm, year	Δln
	electricity use	firm, year	Δln
	(non-electric) fossil fuel use	firm, year	Δln
Treatment			
W	binary indicator for an increase in digitalization	firm, year	$\mathbb{1}\{\Delta D > 0\}$
Covariates			
Q	output	firm, year	$t-1$
K	tangible capital	firm, year	$t-1$
L	number of employees	firm, year	$t-1$
p_M	producer price index	year, sector	$\Delta \ln$
p_E	energy price	year, sector/district	$\Delta \ln$
	prices for electricity and gas	year, consumption level	$\Delta \ln$
	prices for other energy sources	year	$\Delta \ln$
X	lagged outcome (Y_{t-1})	firm, year	$t-1$
	share of energy source (e.g., natural gas/energy use)	firm, year	$t-1$
	R&D intensity (R&D divided by Q)	firm, year	Δ
	tax intensity	firm, year	Δ
	subsidy intensity	firm, year	Δ
	trading intensity	firm, year	Δ
	HHI	year, sector	Δ
	relative use of self-produced fossil-based energy	firm, year	Δ
	relative use of self-produced renewable energy	firm, year	Δ
	proxy for renewable levy (EEG) exemption	firm, year	one-hot
	multi/single unit	firm, year	one-hot
	main industrial grouping	firm, year	one-hot
	structurally weak region	district	one-hot
	sector association	sector	LASSO vector
	location	federal state	LASSO vector
t	time or disembodied technological change	year	LASSO vector

treatment would result in the estimation of linear treatment effects, which we believe is a rather unrealistic assumption. In summary, we observe for approximately 30 percent of firms an increase in software capital. Thus, we consider them as quasi-treated. We include tangible capital in logarithmized lagged levels in the estimation.

For each firm, we additionally observe numerous covariates that we group in five categories: Production function in- and outputs, external factors, firm structure, policy situation, and energy mix. We provide a brief description of the variables here, but refer to the overview in Appendix A (Table A.4) for a detailed description and the data sources.

As motivated in the previous section, in addition to the tangible capital stock, *production function in- and outputs* relate to labor use L, which is approximated by the number of employees, the price for materials, and the energy price as well as the firm-level production value Q. The number of employees and the firm-level production value are integrated in lagged levels to ensure that the exogeneity of the covariates is fulfilled (Assumption A.3 in Section 3.2.1). We approximate the price for materials p_M by the producer price index. For the energy price p_E, we use the location-specific industry average of firm-level expenditure for one kilowatt-hour of energy. We additionally add prices for different energy carriers from external data sources: We merge electricity, natural gas, coal, heating oil, district heat, biomass, and liquid gas prices. Also, we log-transform all price variables.

The variables described in the following pre-dominantly relate to firm and sector characteristics X that potentially influence energy efficiency ϑ. Data on *external factors* cover variables, such as location (federal state), year of observation, which approximates disembodied technological change t, and industry association. In a standard OLS regression, all three characteristics would typically be included as one-hot-encoded fixed effects. However, trees-based algorithms have difficulties with large one-hot-encoded matrices. Therefore, we follow Jens et al. (2021) and modify them in a two-step procedure. First, we estimate the effect of each variable, coded as fixed effects dummies in a LASSO regression, on Y. For instance, we estimate the effect of each manufacturing industry, such as the automotive industry, on energy use. Second, we create a vector of the respective estimation coefficients for each variable and include this vector as a feature in the GRF estimation instead of a one-hot-encoded matrix. Jens et al. (2021) show the effectiveness of this approach in Monte Carlo simulations. Further external factors that are integrated in the estimation are the competitive situation in each industry approximated by the Herfindahl–Hirschman Index (HHI) and a dummy indicating whether the firm is situated in a region considered "structurally weak" due to its limited economic productivity.

Additionally, we include information on the *firm structure*, such as information on the number of plants, industrial grouping (intermediate goods, capital goods, durable consumer goods, non-durable consumer goods, and energy producer), and the volume of traded commodities as well as investment in research and development (R&D) relative to output. Except for the last two variables, which are continuous, we integrate all information in levels and one-hot-encoded. Here, we refrain from applying LASSO-based fixed effects vectors, as the number of categories is small.

The *policy situation* of the firm is characterized by paid taxes and received subsidies. The information is considered proportional to output. Information whether the firm is potentially fully or partly exempt from the EEG levy[9], which is the case for various energy-intensive firms, is also included as a categorical variable.

Last but not least, we include covariates that describe the *energy mix* of the firm, starting with the share of different energy sources used in the production process. All shares are integrated as lagged levels, as current changes may be strongly correlated with the outcome variable. Moreover, we add the share of self-generated energy (fossil and renewable energies).

After this preprocessing our dataset contains $p = 78$ covariates and our sample includes 92,315 observations based on 28,734 firms.

5. Results

We structure our results according to the three research questions posed in Section 3. Hence, we first discuss the general role of digital technologies for energy consumption in the manufacturing sector. Then we turn to heterogeneity-driving characteristics before presenting results on selective targeting of current industrial digitalization policies and their influence on energy use.

5.1. Conditional Average Treatment Effects

We start by estimating the conditional average treatment effects of an increase in the use of digital technologies, approximated by a binary indicator, on total energy use, electricity, and fossil fuels. We use each outcome in a separate analysis, i.e., we estimate three separate causal forest models.

[9]A levy paid in Germany for electricity consumption to promote renewable energies.

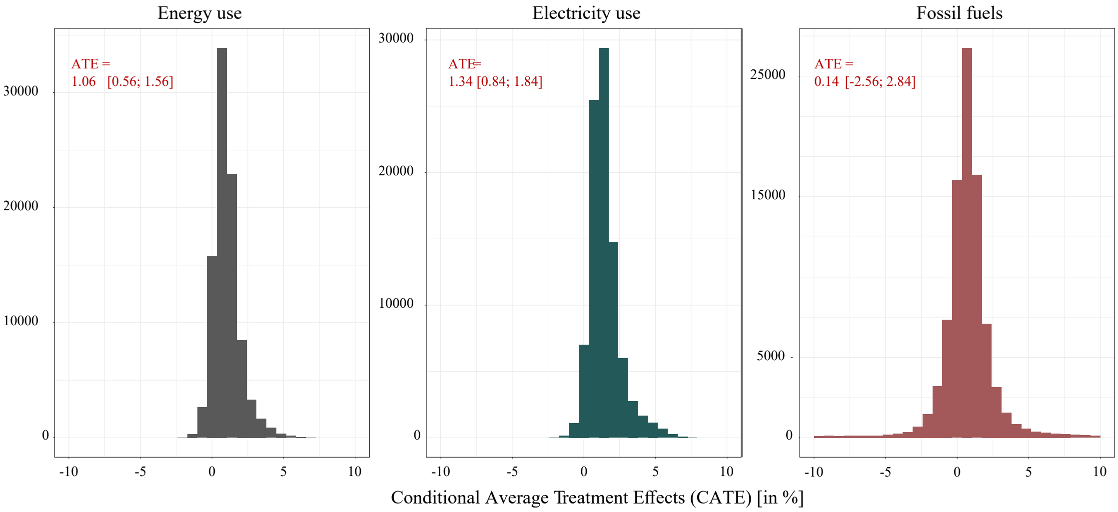

Figure 2: **Distribution of the Conditional Average Treatment Effect (CATE)** for the three different outcome variables: energy use, electricity use, and fossil fuels use.

Figure 2 depicts the distributions of the treatment effects predicted by the causal forest for the three different outcomes. All panels show out-of-bag (OOB) predictions, which are average predictions for each observation, using only trees that do not include the respective observation (James et al., 2021).[10] We find for total energy use that the ICT-related increase in energy consumption ranges roughly from $-3\,\%$ to $6\,\%$ and has its mean at $1.03\,\%$. When electricity use is our dependent variable, the ICT-related increase in electricity consumption is slightly higher and at $1.34\,\%$. The contrary holds for fossil fuels, where the ATE decreases to $0.23\,\%$ and becomes insignificant. It has to be acknowledged here that the range of the CATE distribution is now much broader and spans roughly from $-30\,\%$ to $30\,\%$. Overall, the average treatment effect indicates that an increase in the firm level degree of digitalization is significantly related to higher levels of energy use. However, there is a small share of firms for which the potential outcome declines. Thus, we conclude that for some firms we can observe both, energy savings and an increase in digital technologies, i.e., potential synergies. Nonetheless, firms for which an increase in the software capital stock relates to growing energy use are far more frequent. The positive relationship seems to be particularly pronounced for the electricity use of a firm, while we cannot determine an unambiguous direction of ICT-related changes in energy consumption

[10]We excluded a small test sample of 2% of our observations from the training procedure of the Causal Forest model to analyze the external validity. Figure E.13 in the Appendix shows the CATE predictions for this test sample. The similarity between the distributions indicates that the model is well calibrated.

for fossil fuel use. Accordingly, results suggest that the change in overall energy use is driven by an increase in electricity use. This is in line with the reasoning that ICT consume mainly electric energy.

At first sight, this finding contradicts previous results from Schulte et al. (2016), who observe that ICT relate to a reduction in non-electric energy, but do not significantly affect the demand for electric energy. However, comparing both studies reveals that Schulte et al. (2016) use different outcome variables. For instance, instead of considering absolute electricity use, they use the share of electricity costs in variable costs as a dependent variable. This divergence may explain the differences between the two studies.

We evaluate the Causal Forest fit by applying the Best Linear Prediction Test (Chernozhukov et al., 2018). The test uses the OOB predictions of ICT-related changes in energy consumption to predict actual changes and thereby evaluates the quality of estimates with the following linear model:[11]

$$\begin{aligned}(Y_{i,t} - Y_{i,t-1}) - \hat{m}^{(-i)}(X_{i,t}) = {} & \beta_{\text{ATE}}\bar{\tau}\left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t})\right) \\ & + \beta_{\text{CATE}}\left(\hat{\tau}^{(-i)}(X_{i,t}) - \bar{\tau}\right)\left(W_{i,t} - \hat{e}^{(-i)}(X_{i,t})\right) + \epsilon_{i,t}.\end{aligned} \tag{5}$$

The results for the two β-coefficients are reported in Table 3 with respect to overall energy use. Since β_{ATE} is close to 1, the model captures the average ICT-related changes in energy consumption well. We also find evidence that the covariates adequately capture the underlying heterogeneity, as the second coefficient (β_{CATE}) is also close to 1 and significant. The results of the other two outcomes are reported in Table E.7. Although the results for the electricity model are comparable to those of the overall energy model, the fossil fuel model does not appear to adequately predict ICT-related changes in fossil fuel consumption. Thus, for fossil fuels, we cannot reject the null that no heterogeneity exists.

Table 3: **Best Linear Predictor Test** for the forest with total energy use as outcome.

	Estimate	SE	t-stat	p-value
β_{ATE}	0.998	0.235	4.245	$1.09e-05^{***}$
β_{CATE}	1.261	0.366	3.448	0.0003^{***}

Notes: Results of the best linear predictor test for model calibration and heterogeneity that seeks to fit the estimated CATE as a linear function of the out-of-bag predictions (see Equation (5)).

[11] Accordingly, the model is calibrated well if β_{ATE} and β_{CATE} are close to one.

As our results confirm an increase in energy use at the firm-level. It is intuitive to ask how this result affects the overall energy consumption of the manufacturing sector. However, we have to face a limitation in this regard, as even though our data set covers an extensive set of firms, it is a rolling window survey, which means that not every firm of the manufacturing sector has to answer the survey every year. This makes it difficult to assess whether firms exit or entry the market and therefore aggregated effects at the sectoral level cannot be assessed properly. Thus, we refrain from conclusions with respect to changes in aggregated energy consumption.

5.2. Analyzing Effect Heterogeneity

While the CATE distributions indicate that the relationship is heterogeneous for total energy and electricity use, it does not clarify how the observed covariates are associated with ICT-related changes in energy consumption.

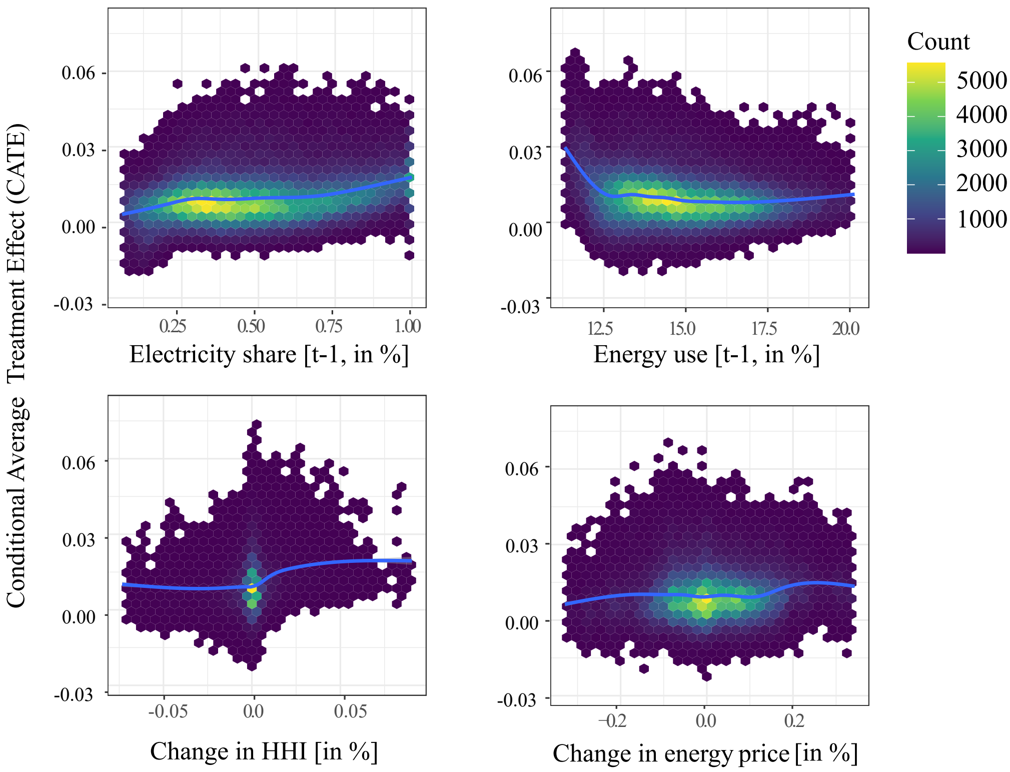

Figure 3: **Bivariate distributions and smoothed regression lines for ICT-related changes in energy consumption and selected variables (total energy use).** The color of the hexagons symbolizes the density of the observations and each hexagon comprises at least 5 individual observations. Individual observations cannot be presented due to anonymity constraints.

Figure 3 shows bivariate distributions and smoothed regression lines for predicted ICT-related changes in overall energy consumption with respect to the following variables: energy use and relative electricity consumption in the previous period, changes in market concentration, and changes in the overall energy price. The four variables were chosen according to the variable importance in the splitting algorithm for overall energy consumption (see Figure D.11).[12]

Previous level of energy use and share of electricity. The upper right panel of Figure 3 indicates that firms that used relatively little energy in the previous period are associated with a greater increase in ICT-related energy use. This may imply that smaller firms increase their energy use to a greater extent when investing in ICT, which can be explained by the phenomenon that digital technologies spark economic growth. In addition, the joint distribution of ICT-related changes in overall energy consumption and the electricity share in the previous period indicates a positive relationship (upper left panel). This result potentially confirms that digitalization more strongly affects electricity-using firms.

HHI. The HHI is positively correlated with predicted CATEs (lower left). This might imply that digital firms that face less competition use relatively more energy than digital firms in less concentrated markets. Accordingly, fierce competition may provide larger incentives to save costs and mitigate the additional energy consumed by digital technologies.

Energy prices. The lower right panel of Figure 3 relates to the overall energy price. It suggests that the association between ICT-related changes in energy consumption and changes in the energy price is positive. Assuming a negative "baseline" effect for energy prices (Labandeira et al., 2017), i.e., a negative own price elasticity, this result indicates that the sensitivity to the energy price decreases for firms that increase their use of digital technologies, since the slope of their energy demand curve potentially becomes less steep compared to firms that do not increase their use of digital technologies.

The energy price only reflects the average price of the energy sources consumed, weighted by their usage. However, in fact, the effects for different energy outcomes may diverge with respect to prices for different energy sources. We assume this because different energy sources can be used as substitutes, and digital technologies may influence own and cross-price elasticities differently. As digital technologies consume electricity, we conjecture that firms that increase their use of digital technologies become more dependent on electricity. Thus, on the one hand, their sensitivity to an increase in the electricity price may decline. On the other hand, if firms increase their use of digital technologies, they may also respond differently to changes in fossil fuel prices. We

[12]The importance of prices is considered jointly.

assume this because they can potentially substitute fossil fuels more easily by electricity and, therefore, may become more sensitive to fossil fuel prices. In summary, we hypothesize that own and cross-price sensitivity for different energy sources is affected if a firm increases its use of digital technologies.

To analyze this claim, we compare the difference between the prices of different energy sources between the 20 % of firms (Q5) with the highest predicted increase in ICT-related energy consumption and the 20 % of firms (Q1) with the lowest predicted increase. The first panel of Figure 4 depicts results for the overall energy use. Each bar represents the price difference of an energy source. We see that the electricity price per kWh is higher in Q5 than in Q1. Hence, the firms for which the ICT-related difference in energy consumption is the largest face higher electricity prices. For natural gas, district heat, and coal, we do not observe any notable price differences between Q1 and Q5. For heating oil and liquid petroleum gas (LPG), we find a negative divergence. Hence, we observe lower respective prices where the difference in energy consumption between ICT-increasing and not ICT-increasing firms is the largest.

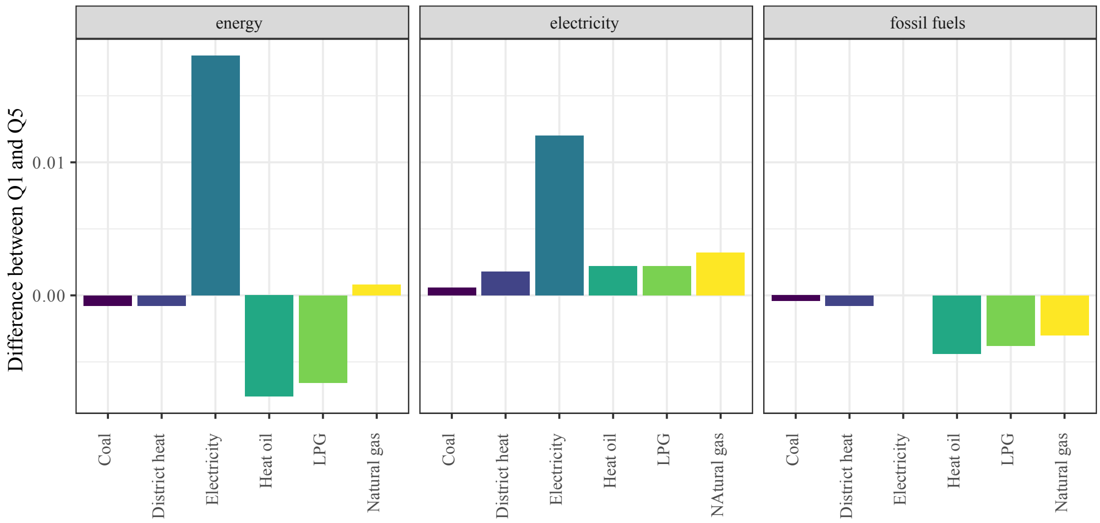

Figure 4: **Difference between energy prices with respect to the 20% percent of firms with the highest ICT-related change in electricity, fossil fuel, or energy consumption and the 20% with the lowest ICT-related change.** We calculate Q5 -Q1.

The second panel of Figure 4 shows price differences for changes in electricity consumption. It is straightforward to see that electricity and fossil fuel prices are higher where the difference in electricity consumption between ICT-increasing and not ICT-increasing firms is the largest. As explained above, two different mechanisms that work in parallel may explain this difference. On the one hand, sensitivity to electricity prices declines for digital firms. On the other hand, digital firms can more easily switch to electricity if prices of fossil fuels increase; hence, the sensitivity to other prices may increase.

The third panel of Figure 4, shows for fossil fuel use that a higher respective positive divergence between ICT-increasing and not ICT-increasing firms can be associated with lower fossil fuel prices. Furthermore, there is no difference between both quintiles with respect to the electricity price. This result is in line with our assumption that price sensitivity increases for fossil fuel prices. However, since the Best Linear Prediction Test does not confirm heterogeneity for fossil fuels, results for fossil fuel use should be interpreted with caution.

In summary, we find that when the electricity difference between ICT-increasing and not ICT-increasing firms is larger than electricity prices are also higher. Furthermore, we find that a smaller increase in energy consumption is more frequently linked to higher fossil fuel prices. Policymakers should be aware that this result suggests that digital firms may be less responsive to an electricity price policy, such as a levy to promote renewable energies, but may be more responsive to a fossil fuel price policy (targeting non-electric energy consumption).

5.3. Group Differences in the Light of Current Policies

In the following, we look at the differences between subgroups with respect to current digitalization policies. So far, German and also European digitalization policies,[13] involve subsidies and funding for small and medium-sized enterprises (SMEs) and for regions that are considered structurally weak. To analyze the interplay of this strategy with climate targets, we conduct a subgroup analysis investigating whether and how the estimated ICT-related increase in energy consumption varies along firm size and regional structure.

We use Group Average Treatment Effects (GATEs) for the analysis. GATEs refer to the average of individual treatment effects over pre-defined, low-dimensional characteristics (Knaus et al., 2021). Therefore, they are more granular than the overall ATE but are easier to interpret than the previously described firm-level effects. In the spirit of Athey et al. (2020), we split the sample into quintiles for the exercise.[14]

[13] For instance "go digital" `https://www.innovation-beratung-foerderung.de/INNO/Navigation/DE/go-digital/Foerdermodell/foerdermodell.html` [online; accessed 17. Mar. 2022] and "digital jetzt", see `https://www.foerderdatenbank.de/FDB/Content/DE/Foerderprogramm/Bund/BMWi/digital-jetzt-investitionsfoerderung-kmu.html` [online; accessed 17. Mar. 2022]

[14] Note that we do not estimate GATEs doubly robust, as AIPW-scores tend to not perform well on smaller samples and the overlap assumption may not be fulfilled anymore (Glynn and Quinn, 2010).

We estimate GATEs for firms located in regions that are considered either structurally weak or strong along three continuous variables that indicate firm size: the number of employees, the tangible capital stock and output. We consider these variables all from the previous period, as decision makers usually observe firm characteristics, and funding decisions are subsequently made. Figure 5 shows the GATEs for each of the three "size" variables separately. The horizontal axes depict quintiles for "size" variables. The vertical axes show the estimated ICT-related increase in energy consumption. Note that we calculate the quintiles before grouping the data by region. Green lines relate to firms in structurally weak regions and purple lines to firms in structurally strong regions.

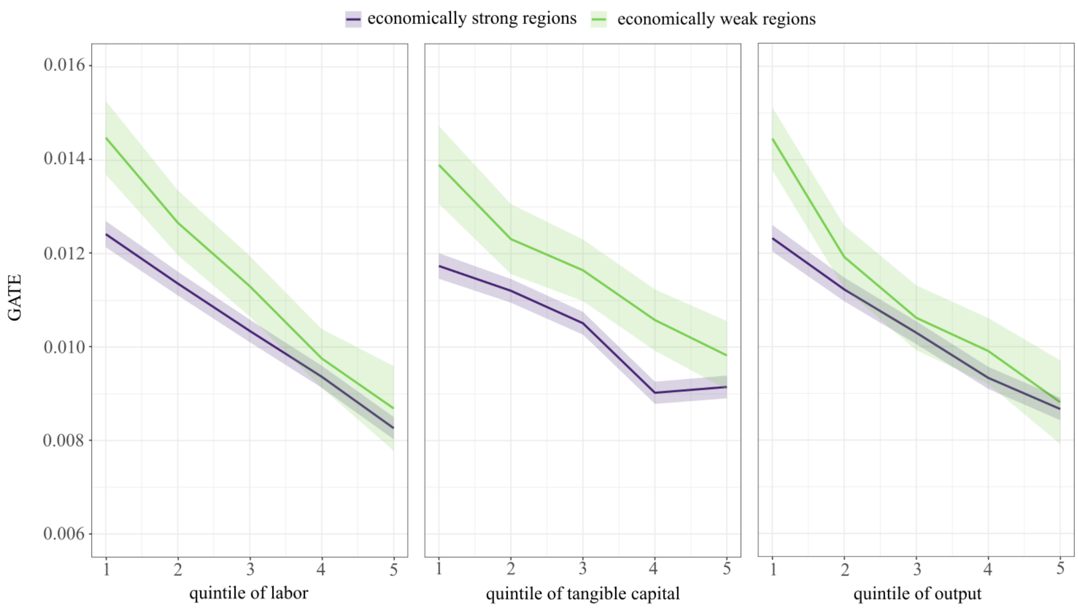

Figure 5: **Group Average Treatment Effects (GATE) grouped by the economic strength of the corresponding region** for different quintiles of labor (number of employees; L), tangible capital (K_N) and output (Q). The two lines relate to the economic strengths of the region, shaded areas denote 10 % confidence intervals

All three panels indicate that the ICT-related increase in energy consumption declines with firm size in both structurally weak and strong regions. The effects in structurally weak regions vary between 1.45 % for firms in the lowest quintiles of labor and output and 0.9 % for firms in respective highest quintiles. Furthermore, the effect size is smaller for structurally strong regions, for which the effect range is between 1.25 % – 0.85 % for quintiles of labor and output. Effect differences for quintiles of tangible capital are slightly less pronounced. Besides, the difference between the energy use of firms with increasing software capital and those without is in particular strong for small firms in structurally weak regions, while the difference between regions is partly insignificant for higher quintiles for each "size" variable and never significant for the highest quintile.

One potential explanation for the higher increase in energy consumption in small firms in structurally weak regions is the fact that digital technologies are a catalyst for economic growth by improving productivity, especially for laggard firms (Borowiecki et al., 2021). Related efficiency improvements exist for economic reasons, such as the generation of scale and scope economies and the reduction of transaction costs (Brynjolfsson and Hitt, 2000). Since larger firms in industrialized regions potentially have advantages in economies of scale and scope and fewer transaction costs, digital technologies may spark here efficiency improvements and economic growth to a lower magnitude. This phenomenon may explain why we observe a larger increase in energy consumption for smaller firms in structurally weak regions. We conclude that a policy trade-off between the goal of saving energy and economic assistance by increasing the use of digital technologies may be especially pronounced for those firms.

In the next step, we analyze group differences with respect to energy-intensive and other industries. We already put forward the hypothesis that relationships may diverge between industries as production processes vary and, hence, can be differently affected by digitalization. Considering that a large share of manufacturing's total energy consumption is driven by a few industries, differences between industries are policy relevant and may be crucial for achieving climate targets.

Similarly to differences between structurally strong and weak regions, we calculate GATEs for energy-intensive and other industries. We consider the following industries as energy-intensive, as they jointly account for more than 80 % of the total energy consumption in manufacturing: "Food, beverages, tobacco products (Division 10–12, 5.8 %)", "Paper & paper products" (Division 17, 5.7 %), "Coke, refined petroleum products" (Division 19, 14.4 %), "Chemicals & chemical products" (Division 20, 32.9 %) "Non-metallic products" (Division 23, 7.4 %), "Basic metals" (Division 24, 16.9 %).[15] We also calculate sector GATEs with respect to quintiles of different "size" variables, as the previous analysis shows large differences in this regard.

Figure 6 shows that the increase in energy consumption is less for firms in energy-intensive industries. For the lowest quintile of labor, for example, the increase in energy consumption for energy-intensive industries is only 1.1 %, whereas it is 1.35 % for other industries. However, the

[15] See German Environmental Agency; www.umweltbundesamt.de/daten/umwelt-wirtschaft/industrie/branchenabhaengiger-energieverbrauch-des#primarenergienutzung-des-verarbeitenden-gewerbes [online; accessed 2. August. 2022].

differences decrease for the higher quintiles of "size" variables and are only significant for the highest quintile of tangible capital.

As assumed, the results suggest that industry differences affect ICT-related changes in energy consumption.[16] Energy-intensive industries are part of the European Union Emissions Trading System (EU ETS). This system generates an additional incentive to save carbon emissions. Hence, an increase in energy consumption may be attenuated for energy-intensive firms by an increasing pressure to save energy-related carbon emissions.[17]

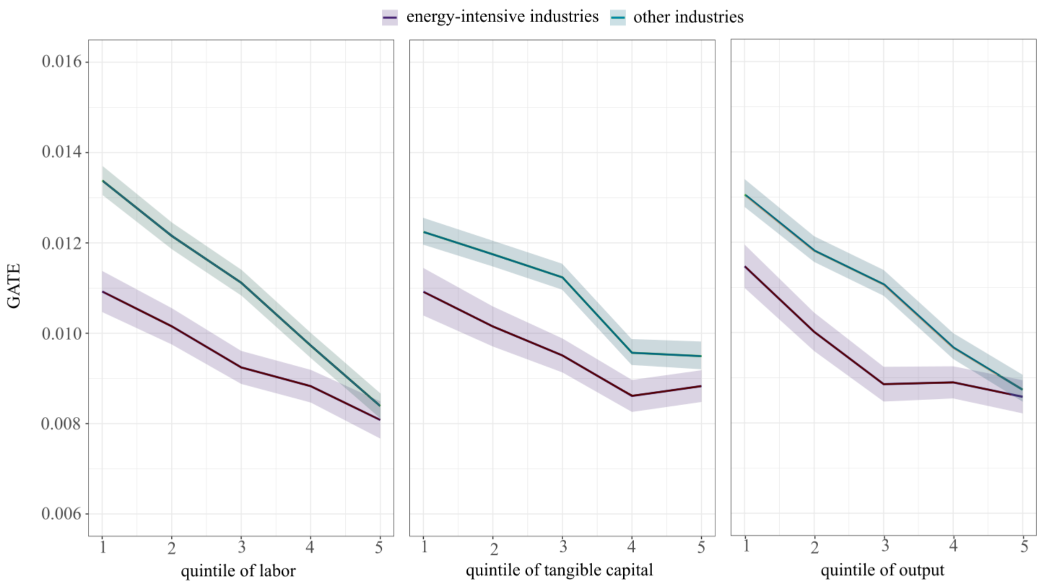

Figure 6: **Group Average Treatment Effects (GATE) grouped by energy-intensive (Divisions: 10-12, 17,19, 20, 23, 24) and remaining industries** for different quintiles of labor (number of employees; L), tangible capital (K_N) and output (Q). The two lines relate to the economic strengths of the region, shaded areas denote 10 % confidence intervals

6. Robustness

The fact that digital technologies can also spark economic growth and via this channel increase energy use points us to an issue. The role of output, labor use, and tangible capital for the relationship between ICT and energy use is ambiguous in our analysis. On the one hand, these variables might be influenced by digital technologies and are part of the impact mechanism.

[16]This result does not contradict findings of Axenbeck and Niebel (2021), who find that energy intensity improvements are rather statistically significant in energy-intensive industries. If output increases parallel to ICT investments than energy intensity can decrease. It will probably improve to a greater extent where the rise in absolute energy consumption is smaller.

[17]We have to acknowledge that prices of emission allowances were rather low between 2011 and 2017.

Therefore, they are a potential source of biased results, as this would violate Assumption A.3, and we cannot consider contemporaneous changes in variables. However, on the other hand, production function in- and outputs may also be potential confounders. For instance, an increase in tangible capital may correlate with the use of digital technologies. This might lead to the rise in tangible capital being the reason for a higher energy consumption, while digital technologies had actually no impact on energy use. Consequently, by integrating only lagged levels, we cannot fully control for confounding due to simultaneous changes in production function in- and outputs. Therefore, we reestimate our model, and replace lagged output, tangible capital, and labor use by logarithmized growth rates (see Appendix E). This enables controlling for respective simultaneous changes. Since the results are comparable to those of our main model, we conclude that the results are robust and contemporaneous changes only play a minor role.

Moreover, we conduct a second robustness check in which we constrain our definition of an increase in digitalization and only consider firms as digital for which the software capital stock per employee increases additionally. In this specification, the ATE is now 0.006 percent. With a p-value of 0.12, the null hypothesis of no increase in energy use cannot be rejected (see Appendix E). The Best Linear Prediction Test shows significant results at the 95 %-level. Hence, considering software capital per employee, our results confirm firm-level heterogeneity. However, it should be acknowledged that this result is significant at a much lower level. An explanation for an attenuated statistical power may be, on the one hand, that we now control more strictly for firm growth (also for the ICT-induced one). On the other hand, we observe firms now in the control group which have been previously considered as digital, i.e., firms with an increasing software capital stock but with a decreasing software capital stock per employee. Observing these firms in the control group may decrease measured ICT-related changes in energy consumption. A further study should shed more light on whether electricity consumption rises because digital firms grow faster or whether ICT capital and electricity are complements and, thus, digital technologies spark electricity use independent of economic growth.

7. Summary and Conclusion

On the one hand, the ongoing digital transformation has raised hopes of climate protection potentials in the energy-intensive manufacturing sector. On the other hand, digital technologies may actually contribute further to environmental damage because they themselves consume energy and resources. However, there is little evidence in the literature that identifies key parameters that determine this relationship.

The main contribution of the article is to disentangle the heterogeneity at the firm level regarding the relationship between ICT and energy use in manufacturing. For this purpose, we apply the Generalized Random Forest algorithm proposed by Athey et al. (2019) to a large administrative panel data set. We harness the panel structure of the data to reduce confounding and mitigate endogeneity issues.

We find that for most firms with an increase in ICT capital, energy use increases relative to firms that do not or barely invest in ICT. Comparing electricity and non-electric fossil fuel use, we additionally show that the relationship differs with respect to different energy sources. We find no significant changes in the use of non-electric fossil fuels, but an average increase in electricity use of 1.34 %. Contrary to political hopes, digital technologies seem to increase energy use at the firm level. However, the increase is particularly related to electricity consumption, for which decarbonization can be realized by renewable energy sources. Furthermore, there is a small share of firms for which energy use declines. Looking closer at the external and firm-level characteristics that may explain heterogeneity, our analysis confirms anticipated rationales. Most interestingly, we observe a growing ICT-related increase in energy consumption with respect to the electricity price, which indicates that the sensitivity to the electricity price declines for digital firms.

Analyzing current policy rationales to target SMEs and firms in regions that are considered structurally weak, the analysis reveals that digitalization policies might not mitigate energy use, while simultaneously fostering technological progress. However, since our study is the first to shed light on characteristics that determine a change in firms' energy consumption as a response to the ongoing digital transformation, there is a strong need for further research. As digital technologies become even more important in the next few years, so will the question of how to actively shape this process into a direction that supports sustainability goals. To be able to systematically align both policies that support technological progress and instruments that reduce energy use, a better understanding of drivers and moderators, i.e., of firm-level heterogeneity, is essential.

Acknowledgements

We thank Susan Athey, Peter Winker, Anthony Strittmatter, Vanessa Behrens, Thomas Niebel, and Petrik Runst for valuable feedback as well as Kerstin Stockmayer for supporting access to the German administrative data. We are grateful for the seminar invitation and feedback received from participants of the seminars at ENSAE ParisTech, ifh Göttingen and the ZEW QUEST seminar. This work also benefited from presentations at the 27th Annual Conference of the European Association of Environmental and Resource Economists in Rimini, the 9th Atlantic Workshop on Energy and Environmental Economics in A Toxa, and the 14th International Conference of the European Society for Ecological Economics in Pisa. This paper has been written with support of the research project "CliDiTrans" Climate Protection Potential of Digital Transformation. The project received funding from the German Federal Ministry of Education and Research (funding ID: 01LA1818B), which played no role in the research.

References

Alataş, Sedat. 2021. The role of information and communication technologies for environmental sustainability: Evidence from a large panel data analysis. *Journal of Environmental Management* 293: 112889.

Almeida, Rita K., Ana M. Fernandes, and Mariana Viollaz. 2020. Software Adoption, Employment Composition, and the Skill Content of Occupations in Chilean Firms. *Journal of Development Studies* 56 (1): 169–185.

Andrae, Anders SG, and Tomas Edler. 2015. On global electricity usage of communication technology: Trends to 2030. *Challenges* 6 (1): 117–157.

Athey, Susan, Rina Friedberg, Nicolaj Søndergaard Mühlbach, Henrike Steimer, and Stefan Wager. 2020. *Essays in Applied Econometrics and Causal Machine Learning*, chapter Between Work, Public Programs, and Retirement: Heterogeneous Responses to a Retirement Reform. Dissertation. `https://pure.au.dk/portal/files/198663023/Nicolaj_Mu_hlbach_PhD_dissertation.pdf#page=63` [Online; accessed 9. Mar. 2022].

Athey, Susan, and Guido W. Imbens. 2019. Machine learning methods that economists should know about. *Annual Review of Economics* 11 (1): 685–725.

Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. Generalized random forests. *Annals of Statistics* 47 (2): 1148–1178.

Athey, Susan, and Stefan Wager. 2019. Estimating treatment effects with causal forests: An application. *Observational Studies* 5 (2): 37–51.

Axenbeck, Janna, and Thomas Niebel. 2021. Climate protection potentials of digitalized production processes: Microeconometric evidence ZEW Discussion Paper No. 21-105, Mannheim.

Baer, Walter S., Scott Hassell, and Ben A. Vollaard. 2002. *Electricity Requirements for a Digital Society.* Santa Monica, CA: RAND Corporation.

Barth, Erling, James C. Davis, Richard B. Freeman, and Kristina McElheran. 2022. Twisting the demand curve: Digitalization and the older workforce. *Journal of Econometrics* `https://doi.org/10.1016/j.jeconom.2021.12.003` (Forthcoming).

Belkhir, Lotfi, and Ahmed Elmeligi. 2018. Assessing ICT global emissions footprint: Trends to 2040 & recommendations. *Journal of Cleaner Production* 177: 448–463.

Ben Lahouel, Béchir, Lotfi Taleb, Younes Ben Zaied, and Shunsuke Managi. 2021. Does ICT change the relationship between total factor productivity and CO2 emissions? Evidence based on a nonlinear model. *Energy Economics* 101: 105406.

Berkhout, Frans, and Julia Hertin. 2001. Impacts of information and communication technologies on environmental sustainability: Speculations and evidence. Report to the OECD, Brighton.

———. 2004. De-materialising and re-materialising: Digital technologies and the environment. *Futures* 36 (8): 903–920.

Bernstein, Ronald, and Reinhard Madlener. 2010. Impact of disaggregated ICT capital on electricity intensity in European manufacturing. *Applied Economics Letters* 17 (17): 1691–1695.

Bertschek, Irene, Daniel Erdsiek, Thomas Niebel, Bettina Schuck, Mareike Seifried, Johannes Ewald, Thorsten Lang, et al. 2020. Schwerpunktstudie Digitalisierung und Energieeffizienz - Erkenntnisse aus Forschung und Praxis. Bundesministerium für Wirtschaft und Energie (BMWi), Berlin.

Bessen, James E., and Cesare Righi. 2020. Information technology and firm employment. Law and Economics Research Paper No. 19-6 (2019), Boston Univ. School of Law. https://ssrn.com/abstract=3371016 [Online; accessed 20 Sep 2022].

Borowiecki, Martin, Jon Pareliussen, Daniela Glocker, Eun Jung Kim, Michael Polder, and Iryna Rud. 2021. The impact of digitalisation on productivity: Firm-level evidence from the Netherlands. *OECD Economics Department Working Papers* (1680).

Breiman, Leo. 2001. Random forests. *Machine Learning* 45 (1): 5–32.

Brozzi, Riccardo, David Forti, Erwin Rauch, and Dominik T Matt. 2020. The advantages of industry 4.0 applications for sustainability: Results from a sample of manufacturing companies. *Sustainability* 12 (9): 3647.

Brynjolfsson, Erik, and Lorin M. Hitt. 2000. Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives* 14 (4): 23–48.

Brynjolfsson, Erik, and Andrew McAfee. 2011. *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy.* Lexington, MA: Digital Frontier Press.

Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram. 2020. Machine learning from schools about energy efficiency. *Journal of the Association of Environmental and Resource Economists* 7 (6): 1181–1217.

Cardona, M., T. Kretschmer, and T. Strobel. 2013. ICT and productivity: Conclusions from the empirical literature. *Information Economics and Policy* 25 (3): 109–125.

Chen, Xiaohong, Xiaomei Gong, Dayuan Li, and Jiaping Zhang. 2019. Can information and communication technology reduce CO2 emission? A quantile regression analysis. *Environmental Science and Pollution Research* 26 (32): 32977–32992.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. 2018. Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in india, National Bureau of Economic Research.

Collard, Fabrice, Patrick Fève, and Franck Portier. 2005. Electricity consumption and ICT in the French service sector. *Energy Economics* 27 (3): 541 – 550.

Davis, Jonathan M.V., and Sara B. Heller. 2017. Using causal forests to predict treatment heterogeneity: An application to summer jobs. *American Economic Review* 107 (5): 546–50.

Dehejia, Rajeev H, and Sadek Wahba. 1999. Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs. *Journal of the American Statistical Association* 94 (448): 1053–1062.

Dehejia, Rajeev H., and Sadek Wahba. 2002. Propensity score-matching methods for nonexperimental causal studies. *The Review of Economics and Statistics* 84 (1): 151–161.

European Commission. 2019. COM(2019) 640 final: The European green deal. `https://eur-lex.europa.eu/resource.html?uri=cellar:b828d165-1c22-11ea-8c1f-01aa75ed71a1.0002.02/DOC_1&format=PDF` [Online; accessed 9. Mar. 2022].

———. 2020. COM(2020) 562 final: Stepping up Europe's 2030 climate ambition – Investing in a climate-neutral future for the benefit of our people. `https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020DC0562&from=EN` [Online; accessed 9. Mar. 2022].

———. 2021. Com(2021) 350 final: Updating the 2020 new industrial strategy: Building a stronger single market for Europe's recovery.

https://ec.europa.eu/info/sites/default/files/communication-industrial-strategy-update-2020_en.pdf [Online; accessed 9. Mar. 2022].

GeSI & Accenture. 2015. Smarter 2030 - ICT solutions for 21st century challenges http://smarter2030.gesi.org/downloads/Full_report.pdf [Online; accessed 9. Mar. 2022].

Gillingham, Kenneth, David Rapson, and Gernot Wagner. 2016. The Rebound Effect and Energy Efficiency Policy. *Review of Environmental Economics and Policy* 10 (1): 68–88.

Glynn, Adam N., and Kevin M. Quinn. 2010. An introduction to the augmented inverse propensity weighted estimator. *Political Analysis* 18 (1): 36–56.

Griliches, Zvi. 1980. R&D and the productivity slowdown. *American Economic Review* 70 (2): 343–348.

Gulen, Huseyin, Candace E Jens, and T Beau Page. 2021. The heterogeneous effects of default on investment: An application of causal forest in corporate finance. https://mays.tamu.edu/department-of-finance/wp-content/uploads/sites/2/2021/03/jens.pd [Online; accessed 26. Oct. 2022].

Horner, Nathaniel C., Arman Shehabi, and Ines L. Azevedo. 2016. Known unknowns: Indirect energy effects of information and communication technology. *Environmental Research Letters* 11 (10): 103001.

Huang, Geng, Ling-Yun He, and Xi Lin. 2022. Robot adoption and energy performance: Evidence from Chinese industrial firms. *Energy Economics* 107: 105837.

IEA. 2019. World Energy Outlook 2019, IEA, Paris, France. https://www.iea.org/reports/world-energy-outlook-2019 [Online; accessed 9. Mar. 2022].

———. 2021. Tracking industry 2021, IEA, Paris, France. https://www.iea.org/reports/tracking-industry-2021 [Online; accessed 14. Feb. 2022].

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2021. *An Introduction to Statistical Learning.* New York, NJ: Springer.

Jens, Candace, T Beau Page, and James Reeder III. 2021. Controlling for group-level heterogeneity in causal forest. https://ssrn.com/abstract=3907601 [Online; accessed 26. Oct. 2022].

Johnson, Matthew S, David I Levine, and Michael W Toffel. 2020. Improving regulatory effectiveness through better targeting: Evidence from OSHA. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (20-019).

Kander, Astrid, Paolo Malanima, and Paul Warde. 2015. *Power to the people: Energy in Europe over the last five centuries.* Boston, NJ: Princeton University Press.

Khazzoom, J Daniel. 1980. Economic implications of mandated efficiency in standards for household appliances. *The Energy Journal* 1 (4): 21–40.

Knaus, Michael C, Michael Lechner, and Anthony Strittmatter. 2021. Machine learning estimation of heterogeneous causal effects: Empirical Monte Carlo evidence. *The Econometrics Journal* 24 (1): 134–161.

Knittel, Christopher R., and Samuel Stolper. 2021. Machine learning about treatment effect heterogeneity: The case of household energy use. *AEA Papers and Proceedings* 111: 440–444.

Kopp, Thomas, and Steffen Lange. 2019. The climate effect of digitalization in production and consumption in OECD countries. In *Proceedings of the 6th International Conference on ICT for Sustainability.*

Labandeira, Xavier, José M. Labeaga, and Xiral López-Otero. 2017. A meta-analysis on the price elasticity of energy demand. *Energy Policy* 102: 549–568.

Lange, Steffen, Johanna Pohl, and Tilman Santarius. 2020. Digitalization and energy consumption. Does ICT reduce energy demand? *Ecological Economics* 176 (June): 106760.

Lechner, Michael. 2015. *The Oxford Handbook of Panel Data*, chapter Treatment Effects and Panel Data. Oxford, UK: Oxford University Press.

Lutz, Benjamin Johannes, Philipp Massier, Katrin Sommerfeld, and Andreas Löschel. 2017. Drivers of energy efficiency in German manufacturing: A firm-level stochastic frontier analysis ZEW Discussion Paper No. 17-68, Mannheim.

Majeed, Muhammad Tariq. 2018. Information and communication technology (ICT) and environmental sustainability in developed and developing countries. *Pakistan Journal of Commerce and Social Sciences* 12 (3): 758–783.

Miller, Steve. 2020. Causal forest estimation of heterogeneous and time-varying environmental policy effects. *Journal of Environmental Economics and Management* 103: 102337.

Nie, Xinkun, and Stefan Wager. 2021. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108 (2): 299–319.

OECD. 2017. *OECD Digital Economy Outlook 2017*. Paris, France: OECD Publishing. `https://www.oecd-ilibrary.org/content/publication/9789264276284-en` [Online; accessed 9. Mar. 2022].

O'Neill, Eoghan, and Melvyn Weeks. 2019. Causal tree estimation of heterogeneous household response to time-of-use electricity pricing schemes. `https://arxiv.org/abs/1810.09179` [Online; accessed 26. Oct. 2022].

Porter, Michael E, and Claas Van der Linde. 1995. Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives* 9 (4): 97–118.

Prest, Brian C. 2020. Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing. *Journal of the Association of Environmental and Resource Economists* 7 (1): 103–143.

Ren, Siyu, Yu Hao, Lu Xu, Haitao Wu, and Ning Ba. 2021. Digitalization and energy: How does internet development affect China's energy consumption? *Energy Economics* 98: 105220.

Research Data Centre of the Statistical Offices Germany, . 2022. AFiD-Panel Industrial Units (doi: 10.21242/42111.2018.00.01.1.1.0), AFiD-Module Use of Energy (doi: 10.21242/43531.2017.00.03.1.1.0) and the panel of the cost structure survey (doi: 10.21242/42251.2017.00.01.1.1.0), for the years 2009-2017, Further information: http://www.forschungsdatenzentrum.de/en [10.01.2022]. [Data Set].

Rhoades, Stephen A. 1993. The Herfindahl-Hirschman Index. *Federal Reserve Bulletin* Mar: 188–189.

Robinson, P. M. 1988. Root-n-consistent semiparametric regression. *Econometrica* 56 (4): 931–954.

Rubin, Donald B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66 (5): 688.

Sadorsky, Perry. 2012. Information communication technology and electricity consumption in emerging economies. *Energy Policy* 48: 130–136.

Schulte, Patrick, Heinz Welsch, and Sascha Rexhäuser. 2016. ICT and the demand for energy: Evidence from oecd countries. *Environmental and Resource Economics* 63 (1): 119–146.

Stiroh, Kevin. 2002. Are ICT spillovers driving the new economy? *Review of Income and Wealth* 48 (1): 33–57.

Taneja, Shivani, and Filip Mandys. 2022. The effect of disaggregated information and communication technologies on industrial energy demand. *Renewable Sustainable Energy Reviews* 164: 112518.

Valente, Marica. 2021. Policy evaluation of waste pricing programs using heterogeneous causal effect estimation DIW Berlin Discussion Paper No. 1980, 2021.

Wager, Stefan, and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113 (523): 1228–1242.

Wang, En-Ze, Chien-Chiang Lee, and Yaya Li. 2022. Assessing the impact of industrial robots on manufacturing energy intensity in 38 countries. *Energy Economics* 105: 105748.

Wen, Huwei, Chien-Chiang Lee, and Ziyu Song. 2021. Digitalization and environment: How does ICT affect enterprise environmental performance? *Environmental Science and Pollution Research* 28 (39): 54826–54841.

Williams, Eric. 2011. Environmental effects of information and communications technologies. *Nature* 479 (7373): 354–358.

Xu, Qiong, Meirui Zhong, and Xin Li. 2022. How does digitalization affect energy? International evidence. *Energy Economics* 107: 105879.

Zhang, Chuanguo, and Cong Liu. 2015. The impact of ICT industry on CO2 emissions: A regional analysis in China. *Renewable and Sustainable Energy Reviews* 44: 12–19.

Appendix A. Description of Variables

Table A.4: Description of Variables

Main variables	
total energy use	Overall firm-level energy use, i.e. the sum of energetic use of different energy carriers plus electricity use (in kWh) observed in the AFiD-Module Use of Energy.
electricity use	Total electricity consumption (in kWh) observed in the AFiD-Module Use of Energy.
fossil fuel use	The sum of firm-level use of natural gas, coal, heating oil, district heat and liquid petroleum gas (in kWh) observed in the AFiD-Module Use of Energy.
treatment	Software capital approximates the degree of firm-level digitalization (in €). We calculate firm-level software capital stocks as in Axenbeck and Niebel (2021) and base them on software investments reported in the AFiD-Panel Industrial Units. Firstly, we generate real software investments using software deflators from Eurostat. Secondly, we apply the perpetual inventory method (PIM) to estimate capital stocks (Griliches, 1980; Lutz et al., 2017). We consider a depreciation rate of 31.5%. The value is retrieved from the EU KLEMS database.[18] Based on these software capital stocks, we calculate software capital growth rates and dichotomize them to generate treatment W. Accordingly, W is one if the software capital stock of firm i increases in period t and zero otherwise. It has to be acknowledged that we only account for purchased software capital and firms may also use open source software. We refer to Axenbeck and Niebel (2021) for a detailed description of the calculation of software capital stocks and a discussion of their representativeness for the firm-level degree of digitalization. Moreover, Axenbeck and Niebel (2021) conduct robustness checks with different depreciation rates and find that the link between software capital intensity and energy intensity is robust to different depreciation values (25%, 33%, and 50%).

Covariates	
lagged outcome	We include the lagged outcome in log-levels in the estimation. If we integrated change rates from the previous period, we would need to consider $t-2$ as well. This would imply that we loose a large share of observations, as our panel is imbalanced.
output	We take the gross production value and subtract turnover from trade and other activities to calculate output (in €). All variables are observed in the AFiD-Panel Industrial Units.
tangible capital	Tangible capital is calculated using real investments in property, plant and equipment (AFiD-Panel Industrial Units, in €) and applying the PIM. Deflators and depreciation rates are taken from the EU KLEMS data. We refer also to Axenbeck and Niebel (2021) for a detailed description of the calculation of tangible capital stocks.

[18] See EU KLEMS database - 2019 release, Germany capital input data, see Stehrer, R., A. Bykova, K. Jäger, O. Reiter and M. Schwarzhappel (2019): Industry level growth and productivity data with special focus on intangible assets, wiiw Statistical Report No. 8. link to data (Retrieved on: 18.04.2020).

labor use	Labor use is measured by the number of employees observed in the AFiD-Panel Industrial Units. We convert part-time employees to full-time employees and adjust the number of employees in this regard.
producer price index	Average material prices are approximated by the index of producer prices of industrial products (domestic sales) retrieved from Destatis. Link to data (retrieved on: 12.11.2020).
price energy	We calculate overall firm-level energy prices in a two-step procedure. Firstly, we divide firm-level energy costs by firm-level energy use to approximate the firm-level energy price (in €/kWh). However, this approach may be endogenous and prone to issues resulting from misreporting. Consequently, we calculate in a second step, based on the firm-level energy price, the average energy price within an industry (4-digit NACE level) in one region (5-digit AGS level [Kreisbene]). We then approximate the firm-level energy price by the regional industry average. This allows considering a more robust energy price. Moreover, if we observe less than five firms in a region within one industry, we approximate the firm-level energy price by the federal-state average at the 2-digit NACE level.
price electricity	Electricity prices are retrieved from Eurostat (status: 08.04.2019, in €/kWh). We consider prices for non-household consumers, which is bi-annual data and we, therefore, take the yearly average. Moreover, prices dependent on the consumption level and we exclude VAT and other recoverable taxes and levies. As firms switch their consumption level over time, we consider the consumption level of the first period we observe a firm to match prices. This allows for not considering price variations due to changes in the consumption level. Link to data (retrieved on: 15.07.2020).
price natural gas	Natural gas prices are also retrieved from Eurostat (status: 10.02.2020, in €/ GJ). We consider bi-annual natural gas prices (average price per year is calculated) for non-household consumers. Natural gas prices dependent on the consumption level. Accordingly, we consider the consumption level of the first period we observe a firm to match prices. Prices are retrieved excluding VAT and other recoverable taxes and levies. Natural gas prices are converted from GJ to kWh. Link to data (retrieved on: 15.07.2020).
prices of other energy carriers	Other energy prices are retrieved from IEA (liquid petroleum gas, retrieved on: 04.09.2019), Destatis & DEPI (biomass, retrieved on: 16.07.2020 [Destatis], and retrieved on: 13.09.2019 [DEPI]), AGFW (district heat, retrieved on: 14.08.2019), BMWK former BMWi (heating oil, retrieved on: 01.04.2020). For a more detailed description on sources for energy prices see Axenbeck and Niebel (2021).
share of energy sources	To consider the energy mix, we divide the use of electricity, natural gas, coal, heating oil, district heat, liquid petroleum gas and biomass by overall energy consumption and consider each share as a variable in the causal forest model. All variables are observed in the AFiD-Module Use of Energy. We include each energy share in lagged levels in our estimation.
R&D intensity	We divide the total expenditure on research & development observed in the AFiD-Panel Industrial Units by output.
tax intensity	The amount of taxes (e.g. property tax, motor vehicle tax, excise duties; excluding income and corporation tax, equalization levies on burdens and VAT) observed in the AFiD-Panel Industrial Units is divided by output.

subsidy intensity	The amount of subsidies received for current production in the business year observed in the AFiD-Panel Industrial Units is divided by output.
trading intensity	The total turnover of trading goods during the business year observed in the AFiD-Panel Industrial Units is divided by output. Trading goods are considered to be goods of foreign origin that are generally resold unprocessed and without a production-related connection to own products.
Herfindahl–Hirschman Index	The HHI captures the competitive situation that a firm has to face. It is calculated using yearly revenue-based market shares at the 4-digit NACE level observed in the AFiD-Panel Industrial Units. For a detailed description of the HHI see Rhoades (1993). We exclude industries for which we observe less than five firms per year.
share renewable production	Own electricity generation from renewable power observed in the AFiD-Module Use of Energy is divided by overall energy use.
share fossil production	Own electricity generation from fossil sources observed in the AFiD-Module Use of Energy is divided by overall energy use.
weak region	We include a dummy indicating whether a firm is situated in a region (5-digit AGS level) that is considered as "structurally weak" (0) due to its limited economic productivity or "structurally strong" (1). An overview map of structurally weak regions can be found at: https://www.bmwi.de/Redaktion/DE/Dossier/Digital-Jetzt/digital-jetzt-infografik-strukturschwache-regionen.html
EEG exemption	A one-hot encoded variable is generated that indicates whether a firm is partly (1) or fully (2) exempted from charges under the law on renewable energies (EEG). This is calculated by means of the approximated ratio between electricity costs and value added as well as electricity use. For this purpose, we combine information from the AFiD-Panel Industrial Units and the AFiD-Module Use of Energy.
energy intensive industry	We define an energy-intensive industry as an industry or a group of industries at the 2-digit NACE level that accounts for more than 5 % of total energy consumption of the manufacturing sector (Divisions: 10-12, 17,19, 20, 23, 24). The information is retrieved from the German Environmental Agency.
main industrial grouping	We add a one-hot encoded variable that indicates the industrial main group of the firm (intermediate goods producer (1), capital goods producer (2), durable goods producer (3), consumer goods producer (4), and energy producer(5))
single- / multi-unit firm	A one-hot encoded variable is included that indicates whether a firm is a single-unit firm (1), a multi-unit firm in one federal-state (2), or a multi-unit firm in several federal states (3).
industry association	A LASSO-based fixed effects vector controlling for the industry assignment is calculated based on 2-digit NACE codes. For a detailed description see Jens et al. (2021).
year	A LASSO-based fixed effects vector controlling for the observation period (year) is generated based on Jens et al. (2021).
federal states	A LASSO-based fixed effects vector controlling for the federal state of the company's registered office is calculated based on Jens et al. (2021).

Appendix B. Descriptive Statistics

Table B.5: Averages and (standard errors) of firm characteristics for treated and untreated firms (2010 to 2017).

variable	mean control	s.d. control	mean treated	s.d. treated
total energy use (in GWh)	29.64	414.65	45.14	396.57
total energy use $ln\Delta$	0.02	0.27	0.03	0.27
total energy use log $t-1$	14.72	1.9	15.2	1.92
electricity use (in GWh)	9.61	74.67	16.18	132.21
electricity use $ln\Delta$	0.02	0.27	0.03	0.28
electricity use log $t-1$	13.88	1.9	14.42	1.9
fossil fuel use (in GWh)	17.71	281.98	25.89	262.55
fossil fuel use $ln\Delta$	0.05	1.38	0.06	1.34
fossil fuel use log $t-1$	13.23	3.55	13.84	3.29
treatment (W)	0.0	0.0	1.0	0.0
output (in million €)	62.15	572.48	122.86	1250.08
output $ln\Delta$	0.04	0.2	0.06	0.21
output log $t-1$	16.53	1.44	17.11	1.44
tangible capital (in million €)	16.55	152.71	33.75	326.58
tangible capital $ln\Delta$	0.02	0.24	0.05	0.22
tangible capital log $t-1$	14.78	1.93	15.51	1.76
number of employees	241.07	1699.02	431.28	2836.52
number of employees $ln\Delta$	0.01	0.13	0.03	0.13
number of employees log $t-1$	4.59	1.09	5.04	1.16
producer-price index	99.5	3.63	99.57	3.74
producer-price index $ln\Delta$	0.01	0.03	0.01	0.03
price energy	0.13	0.03	0.13	0.03
price energy $ln\Delta$	0.01	0.12	0.01	0.13
price biomass	0.04	0.0	0.04	0.0
price biomass $ln\Delta$	0.02	0.08	0.02	0.07
N	64933		27382	

variable	mean control	s.d. control	mean treated	s.d. treated
price coal	0.01	0.0	0.01	0.0
price coal $ln\Delta$	0.02	0.19	-0.0	0.21
price district heat	0.07	0.0	0.07	0.0
price district heat $ln\Delta$	0.01	0.05	0.01	0.05
price electricity	0.14	0.03	0.14	0.02
price electricity $ln\Delta$	0.04	0.06	0.05	0.06
price heating oil	0.06	0.01	0.06	0.01
price heating oil $ln\Delta$	0.0	0.22	-0.01	0.24
price liquefied petroleum gas	0.06	0.01	0.06	0.01
price liquefied petroleum gas $ln\Delta$	-0.02	0.16	-0.03	0.18
price natural gas	0.04	0.01	0.04	0.01
price natural gas $ln\Delta$	-0.01	0.09	-0.02	0.1
biomass share [in %]	0.02	0.1	0.02	0.09
biomass share [in %] $t-1$	0.02	0.1	0.01	0.09
coal share [in %]	0.01	0.06	0.01	0.06
coal share [in %] $t-1$	0.01	0.05	0.01	0.06
district heat share [in %]	0.03	0.12	0.04	0.13
district heat share [in %] $t-1$	0.03	0.12	0.03	0.13
electricity share [in %]	0.49	0.25	0.51	0.24
electricity share [in %] $t-1$	0.49	0.25	0.51	0.24
natural gas share [in %]	0.32	0.29	0.33	0.28
natural gas share [in %] $t-1$	0.31	0.29	0.32	0.28
heating oil share [in %]	0.11	0.22	0.09	0.19
heating oil share [in %] $t-1$	0.12	0.22	0.1	0.2
liquefied petroleum gas share [in %]	0.01	0.06	0.01	0.05
liquefied petroleum gas share [in %] $t-1$	0.01	0.06	0.01	0.05
R&D intensity [in %]	0.01	0.03	0.02	0.04
R&D intensity [in %] Δ	-0.0	0.02	0.0	0.02
tax intensity [in %]	0.01	0.02	0.01	0.03
tax intensity [in %] Δ	-0.0	0.01	-0.0	0.01
subsidy intensity [in %]	0.0	0.01	0.0	0.01
N	64933		27382	

variable	mean control	s.d. control	mean treated	s.d. treated
subsidy intensity [in %] Δ	-0.0	0.01	0.0	0.01
trading intensity [in %]	0.11	1.06	0.12	0.5
trading intensity [in %] Δ	0.01	1.0	-0.0	1.54
share self-produced fossil-based energy [in %]	0.0	0.02	0.01	0.03
share self-produced fossil-based energy [in %] Δ	0.0	0.01	0.0	0.02
share self-produced renewable energy [in %]	0.01	0.04	0.01	0.04
share self-produced renewable energy [in %] Δ	0.0	0.03	0.0	0.03
HHI	0.07	0.09	0.07	0.09
HHI Δ	0.0	0.03	0.0	0.03
no EEG exemption [in %]	0.92	0.28	0.92	0.28
partial EEG exemption [in %]	0.08	0.27	0.07	0.26
full EEG exemption [in %]	0.01	0.09	0.01	0.1
single-unit firm	0.78	0.41	0.75	0.43
multi-unit firm in one federal state [in %]	0.08	0.27	0.07	0.26
multi-unit firm in several federal states [in %]	0.14	0.35	0.18	0.38
intermediate goods producer [in %]	0.44	0.5	0.44	0.5
capital goods producer [in %]	0.29	0.45	0.34	0.47
durable goods producer [in %]	0.04	0.19	0.04	0.2
consumer goods producer [in %]	0.23	0.42	0.18	0.39
energy producer in [in %]	0.0	0.02	0.0	0.03
year 2010 [in %]	0.14	0.35	0.13	0.34
year 2011 [in %]	0.14	0.34	0.15	0.36
year 2012 [in %]	0.06	0.23	0.07	0.25
year 2013 [in %]	0.14	0.35	0.13	0.34
year 2014 [in %]	0.15	0.36	0.13	0.34
year 2015 [in %]	0.15	0.35	0.13	0.34
year 2016 [in %]	0.08	0.28	0.13	0.33
year 2017 [in %]	0.14	0.35	0.12	0.33
Industry: Food products [in %]	0.14	0.35	0.09	0.29
Industry: Beverages [in %]	0.02	0.13	0.02	0.12
Industry: Tobacco products [in %]	0.0	0.0	0.0	0.0
N	64933		27382	

variable	mean control	s.d. control	mean treated	s.d. treated
Industry: Textiles [in %]	0.02	0.15	0.02	0.14
Industry: Wearing apparel [in %]	0.01	0.11	0.01	0.1
Industry: Leather & related products [in %]	0.01	0.08	0.0	0.07
Industry: Wood, wood & cork products [in %]	0.02	0.15	0.02	0.13
Industry: Paper & paper products [in %]	0.03	0.17	0.03	0.16
Industry: Printing, recorded media [in %]	0.02	0.14	0.02	0.13
Industry: Coke, refined petroleum products [in %]	0.0	0.02	0.0	0.03
Industry: Chemicals & chemical products [in %]	0.06	0.23	0.06	0.24
Industry: Basic pharmaceutical products [in %]	0.01	0.1	0.02	0.12
Industry: Rubber & plastic products [in %]	0.06	0.23	0.06	0.24
Industry: Other non-metallic mineral prod. [in %]	0.05	0.22	0.05	0.21
Industry: Basic metals [in %]	0.05	0.21	0.05	0.21
Industry: Fabricated metal products [in %]	0.13	0.34	0.12	0.33
Industry: Computer, electro, optical prod. [in %]	0.04	0.2	0.05	0.23
Industry: Electrical equipment [in %]	0.06	0.24	0.07	0.25
Industry: Machinery and equipment n.e.c. [in %]	0.13	0.34	0.18	0.38
Industry: Motor vehicles, (semi-)trailers [in %]	0.04	0.19	0.05	0.21
Industry: Other transport equipment [in %]	0.01	0.12	0.02	0.12
Industry: Furniture [in %]	0.02	0.14	0.02	0.14
Industry: Other manufacturing [in %]	0.04	0.19	0.04	0.19
Industry: Repair and installation [in %]	0.03	0.17	0.02	0.15
N	64933		27382	

Appendix C. Evaluating Assumptions and the Fit of the Causal Forest

To assess the overlap assumption, we plot the propensity scores that indicate the probability of treatment for each observation in Figure C.7. Note that we lose four observation due to trimming. The histograms for the (trimmed) treated and not treated firms overlap in a way that makes it impossible to deterministically decide on the treatment status of a firm, as the scores are bounded away from zero and one. Hence, the overlap assumption is fulfilled. Furthermore, the

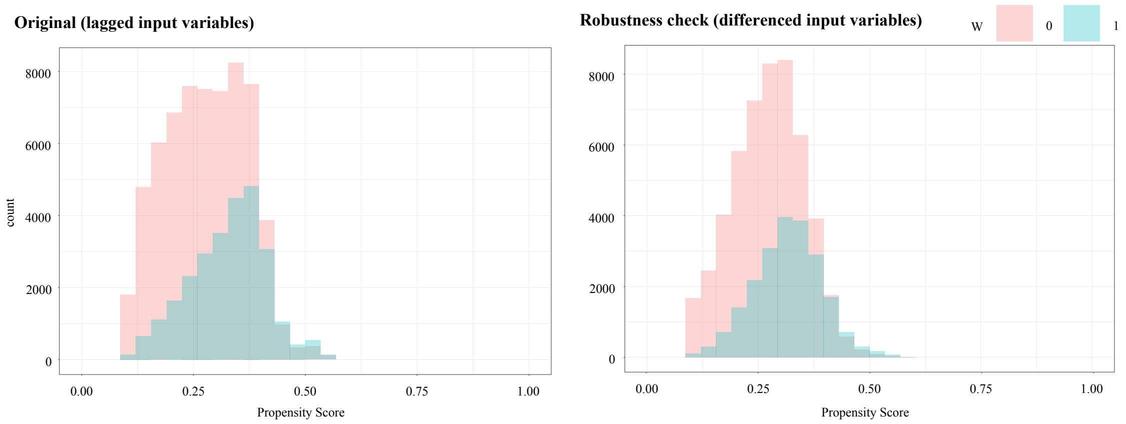

Figure C.7: **Distribution of propensity scores** between treatment and control group.

internal validity of the causal forest approach is based on the idea, that the sample is reweighted in a way that makes the treatment random. Therefore, we assess the balance of the covariates between firms that increase software capital and firms that do not increase software capital. Figure C.10 depicts these differences in the distributions of the treated and untreated samples after reweighting the covariates with the inverse propensity score. Except for some very rare outliers, the distributions do not show any notable differences between both groups. Thus, our model is able to appropriately balance covariates.

Additionally, we test the model calibration by comparing OOB predictions to actual changes in energy consumption using the training sample (see Equation 5). The results for all three model outcomes are presented in Table E.7. According to the test results, the model for electricity use seems to be calibrated well and the performance is comparable to the model with energy use as an outcome variable. In contrast, the model using fossil fuels as an outcome variable fails in predicting an average treatment effect that is different from zero and does not seem to capture the underlying heterogeneity adequately. Model results of this outcome should therefore be interpreted cautiously.

Table C.6: **Best Linear Predictor Test** for the forest with all energy use outcomes.

Outcome variable	Coefficient	Estimate	SE	t-stat	p-value
Energy use	β_{ATE}	0.998	0.235	4.245	1.09e05***
	β_{CATE}	1.261	0.366	3.448	0.0003***
Electricity use	β_{ATE}	0.980	0.172	5.695	1.96e$-$09***
	β_{CATE}	0.914	0.316	2.897	0.002***
Fossil fuel use	β_{ATE}	1.442	3.870	0.373	0.355
	β_{CATE}	-0.833	0.784	-1.061	0.856

Notes: Results of the best linear predictor test for model calibration and heterogeneity that seeks to fit the estimated CATE as a linear function of the out-of-bag predictions (see 5)

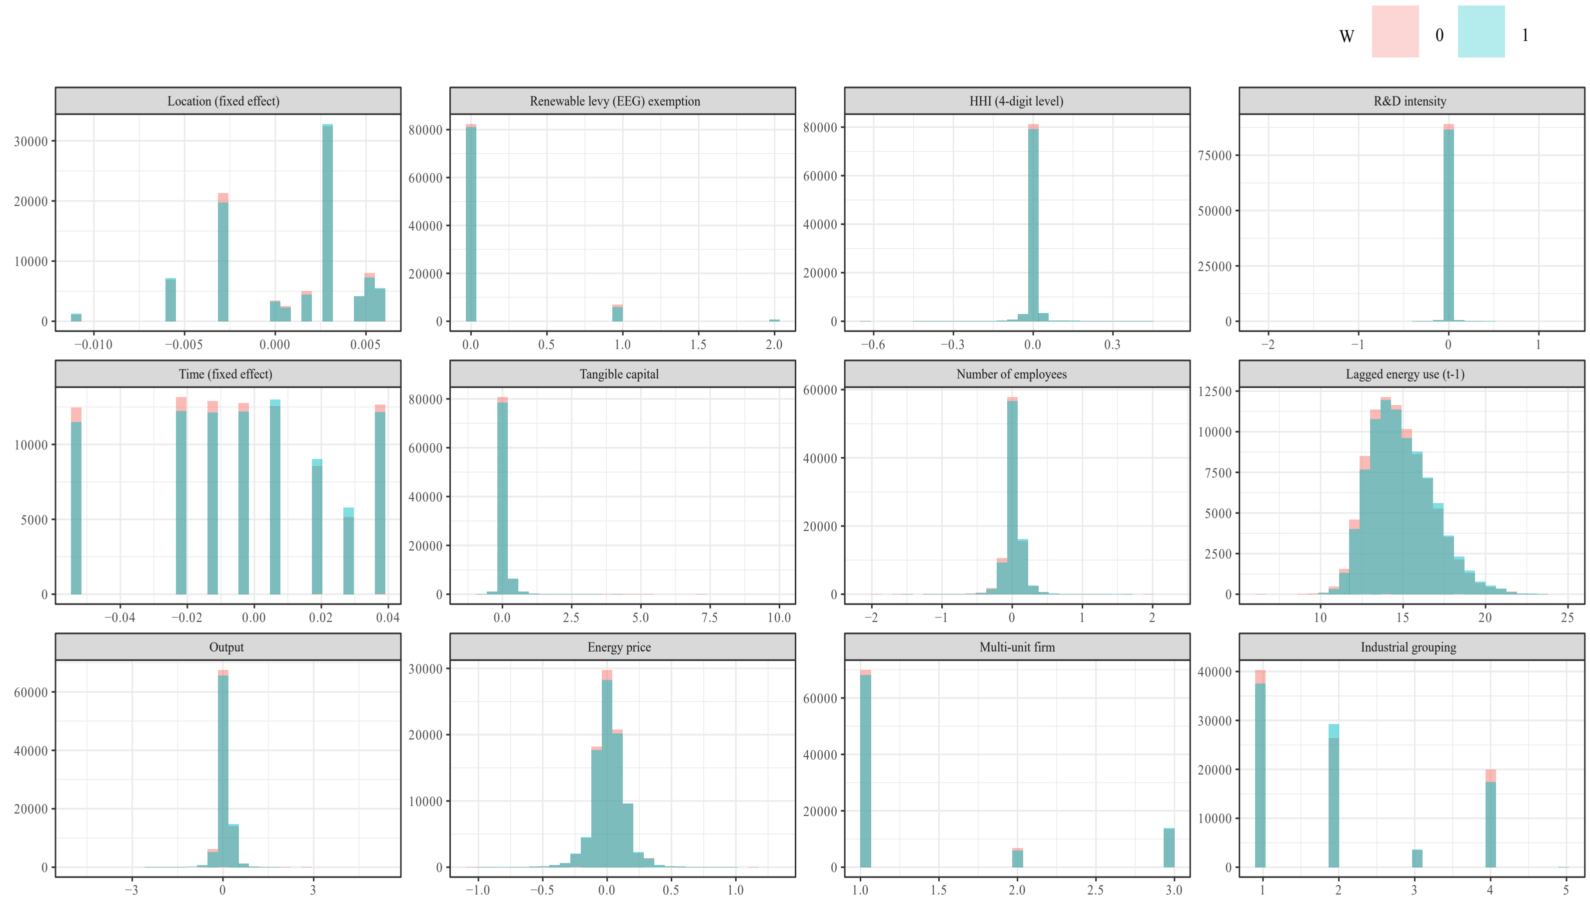

Figure C.8: **Inverse-propensity weighted histograms** for treated and untreated observations (Part I).

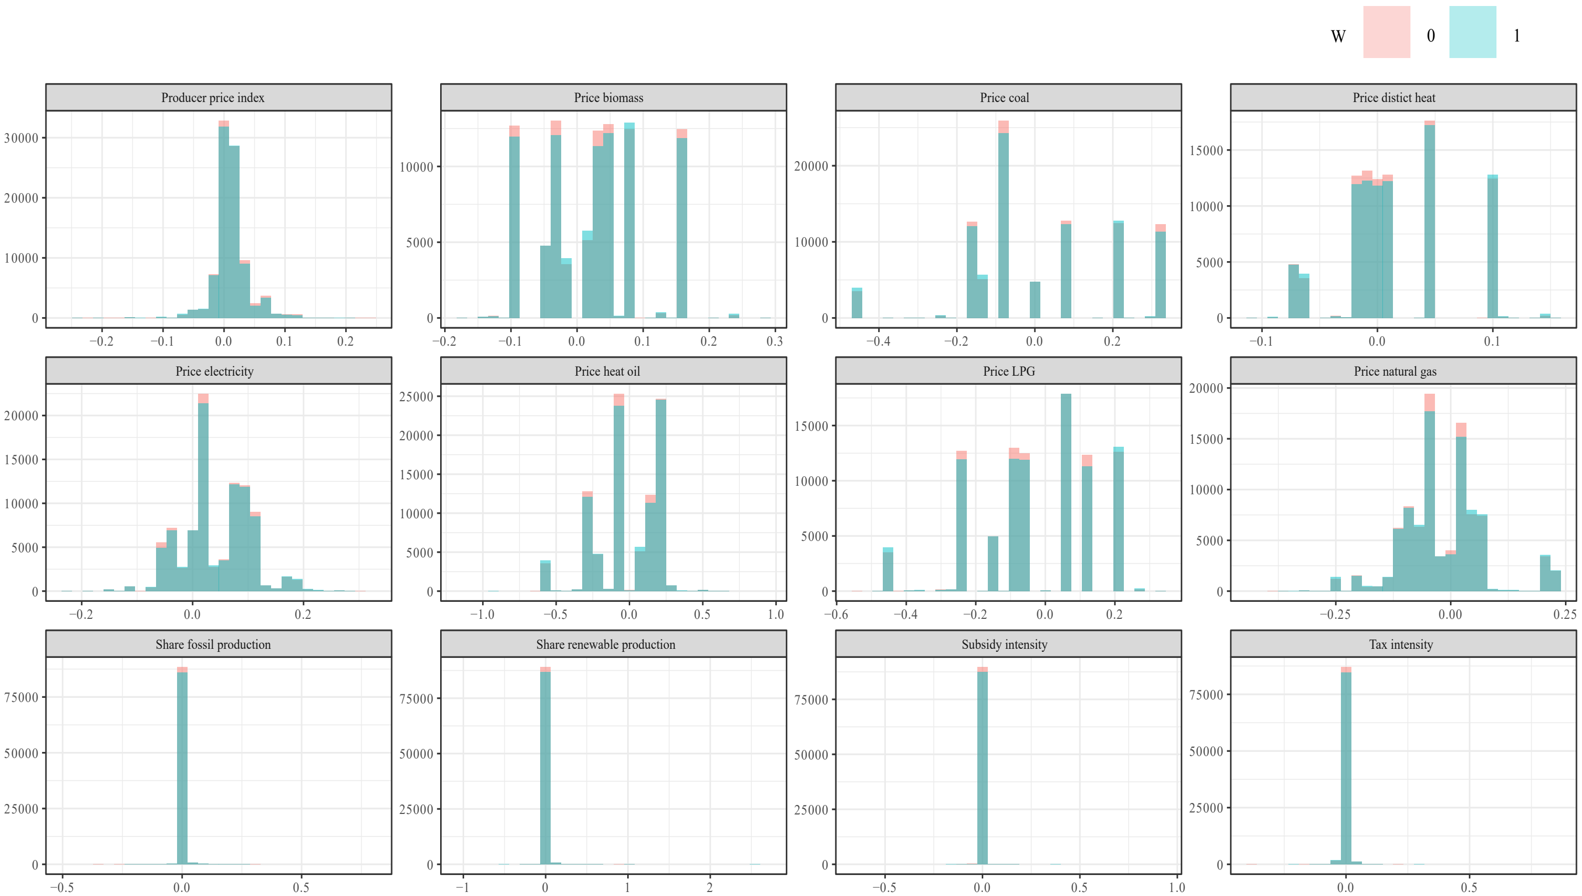

Figure C.9: **Inverse-propensity weighted histograms** for treated and untreated observations (Part II).

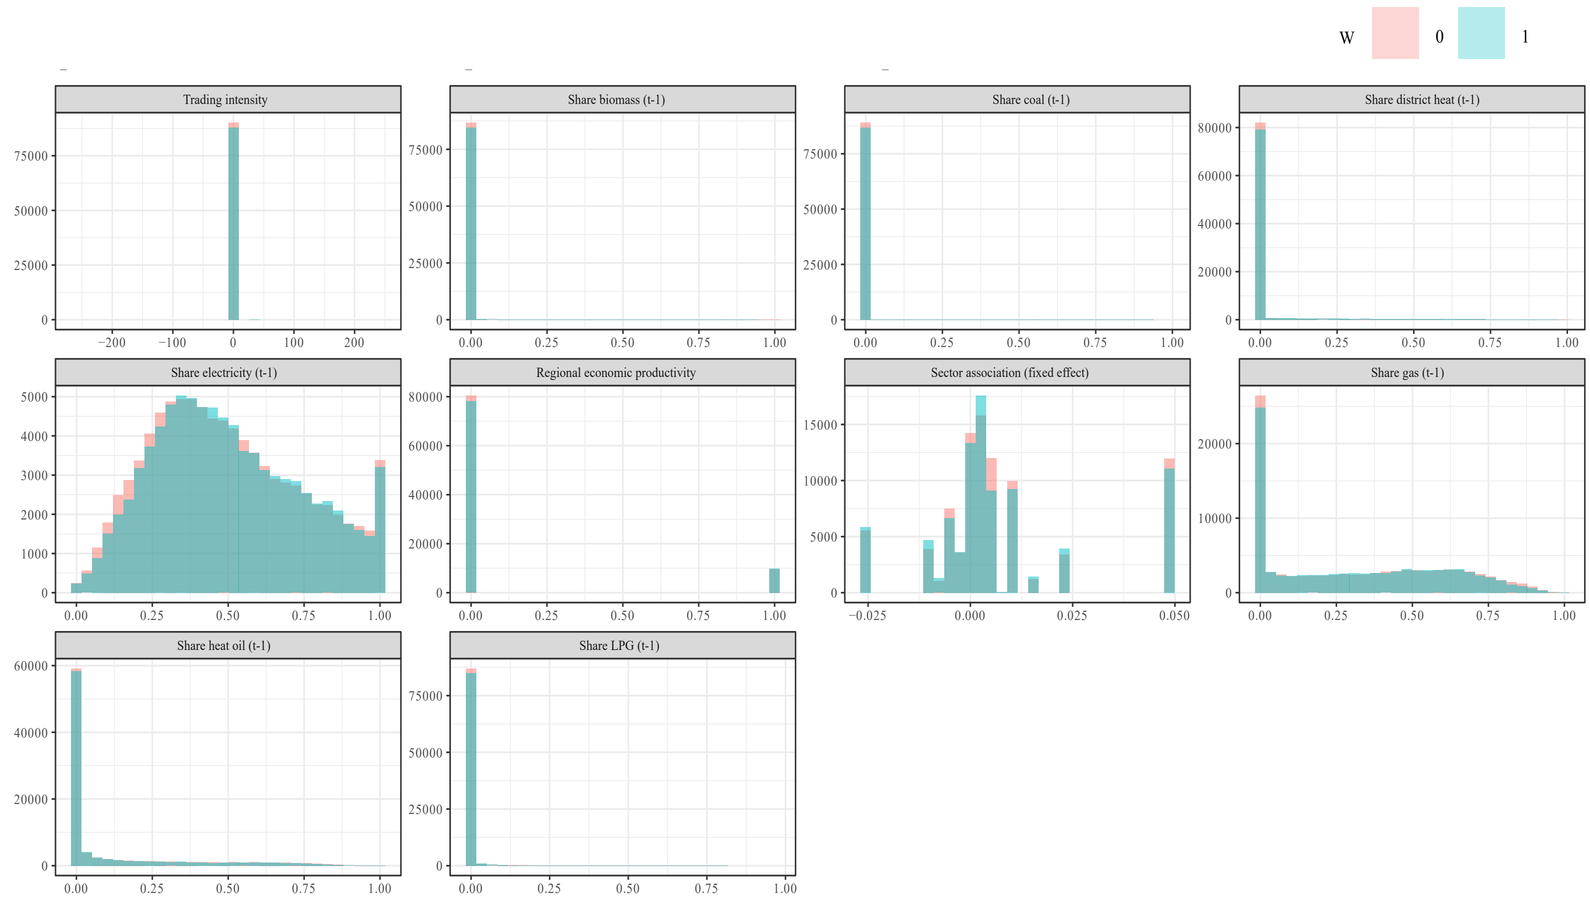

Figure C.10: **Inverse-propensity weighted histograms** for treated and untreated observations (Part III).

Appendix D. Variable Importance

To understand the main drivers of treatment heterogeneity, we analyze the firm characteristics that were used as splitting variables in the forest. Variable importance (VI) measures how many times a covariate was used for splitting at level l across all trees t, where relative split frequency (RSF) denotes the split frequency (SF) of variable m divided by all splits at level l. Additionally, weights ($w_l = l^{-2}$) are used that exponentially favor higher tree levels.

$$VI_m = \frac{\sum_{l=1}^{L} RSF_{ml} * w_l}{\sum_{l=1}^{L} w_l} \tag{D.1}$$

$$RSF_{ml} = \frac{SF_{ml}}{\sum_{m=1}^{M} SF_{ml}} \tag{D.2}$$

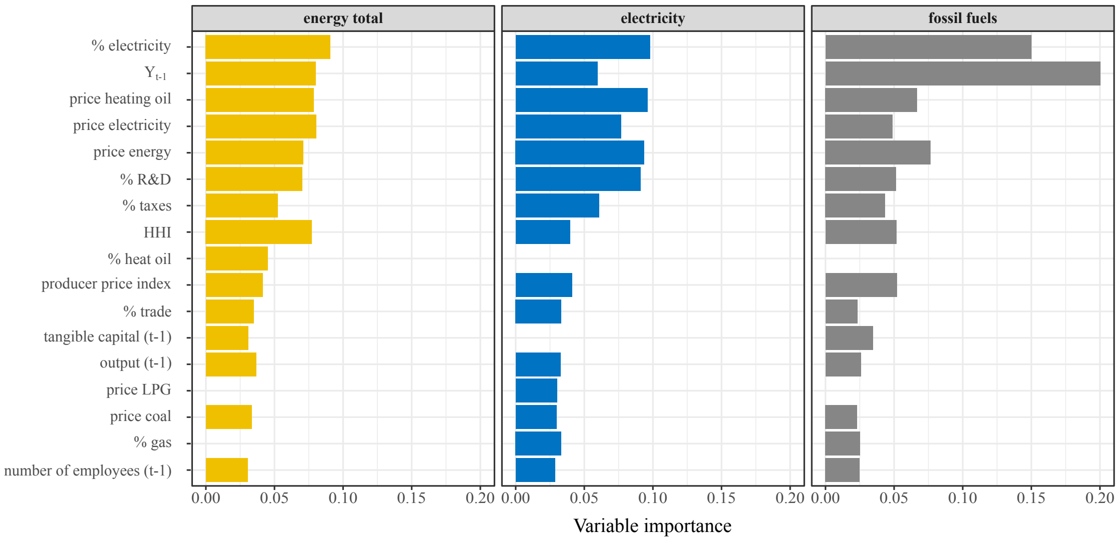

Figure D.11: **Variable Importance** for the three causal forests with the outcome variables total energy use, electricity use, and fossil fuels.

Figure D.11 lists the 15 most important variables for splitting the sample into groups. Combined, energy prices are by far the most important variable, if we sum up the importance values of the energy prices (total energy price, price of heat oil, price of LPG, price of electricity and price of coal). Furthermore, the share of electricity use (of the previous period) is important for the splitting procedure of all three outcomes.

Appendix E. Robustness Analysis

We conduct the following robustness analysis:

Growth rates: We repeat our analysis and replace output, tangible capital, and labor use in lagged levels by logarithmized growth rates as one alternative specification.

D/L: In a further robustness test, we modify our treatment variable and consider only firms as treated if their software capital per employee increases. We repeat the analysis and replace the treatment dummy D with a relative dummy that represents the change in the capital stock (ΔK_{ICT}) relative to the number of employees (L).

Table shows estimated ATEs as well as the performance of the Best Linear Prediction Test for both robustness checks. In the first specification (growths rates), the ATE is now at 0.01 and significant. The mean and differential forest predication indicate that treatment effects are well calibrated. Hence, different formulations of production function in- and outputs does not alter main results. However, it is noteworthy that the variable importance of these critical variables increases when they are included in logarithmized growth rates (not displayed). In the second robustness test (D/L), the ATE is now at 0.006, which is slightly smaller than in our main specification and the p-value is at 0.12. The Best Linear Prediction Test confirms that the model is well calibrated. Hence, we can also confirm heterogeneity by our modified digitalization indicator.

Table E.7: **Robustness tests.**

Robustness type	Outcome variable	Variable	Estimate	SE	t-stat	p-value
Growth rates	Energy use	ATE	0.010	0.003	3.330	0.0004^{***}
		β_{ATE}	1.052	0.308	3.413	0.0003^{***}
		β_{CATE}	1.117	0.420	2.660	0.0004^{**}
D/L	Energy use	ATE	0.006	0.005	1.180	0.119
		β_{ATE}	1.011	0.459	2.202	0.012^{*}
		β_{CATE}	0.815	0.418	1.950	0.026^{*}

Notes: Results for the different robustness models.

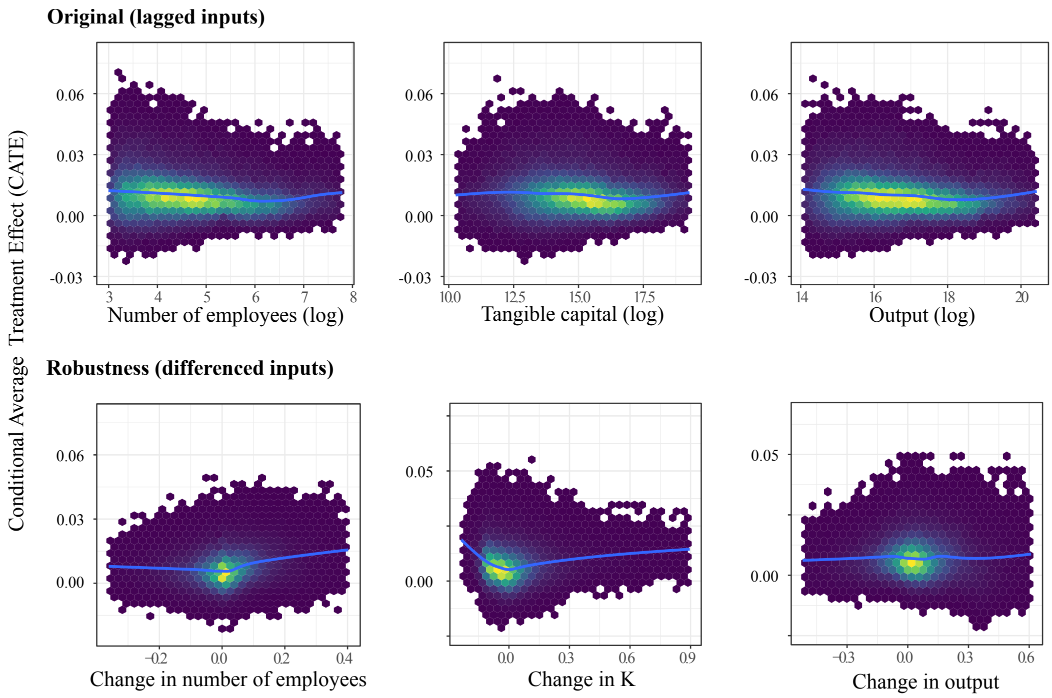

Figure E.12: **Bivariate distributions of the predicted treatment effect and production factors.** Comparison of the original model (lagged input variables) and the robustness check (differenced inputs).

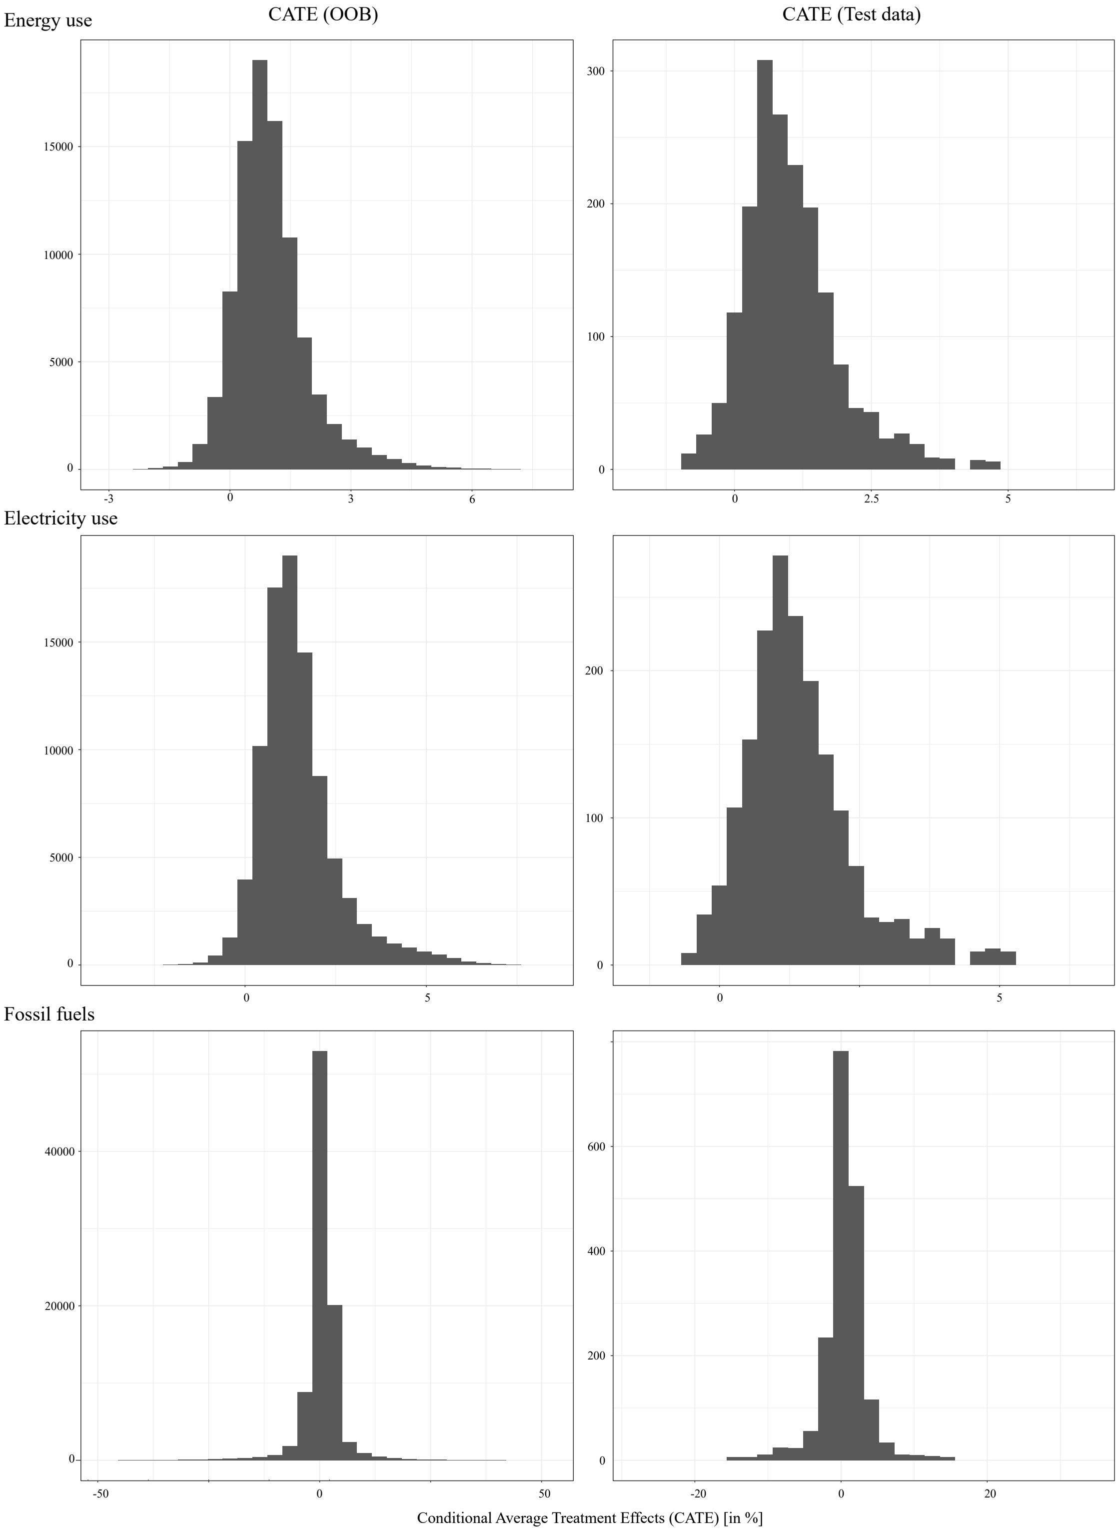

Figure E.13: Comparison between OOB predictions and predictions of the test sample.

Download ZEW Discussion Papers:

https://www.zew.de/en/publications/zew-discussion-papers

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//

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