

CAMP Working Paper Series
No 5/2019

On the use of machine learning for causal inference in climate economics

Isabel Hovdahl



© Authors 2019

This paper can be downloaded without charge from the CAMP website <http://www.bi.no/camp>

On the use of machine learning for causal inference in climate economics

Isabel Hovdahl ^{*†}

June 2019

Abstract

One of the most important research questions in climate economics is the relationship between temperatures and human mortality. This paper develops a procedure that enables the use of machine learning for estimating the causal temperature-mortality relationship. The machine-learning model is compared to a traditional OLS model, and although both models are capturing the causal temperature-mortality relationship, they deliver very different predictions of the effect of climate change on mortality. These differences are mainly caused by different abilities regarding extrapolation and estimation of marginal effects. The procedure developed in this paper can find applications in other fields far beyond climate economics.

Keywords: Climate change, machine learning, mortality.

JEL Classification: C14, C51, I10, Q54

^{*}Department of Economics, Norwegian University of Science and Technology, 7491 Trondheim, Norway.
Phone: +47 90280774. E-mail address: isabel.hovdahl@ntnu.no

[†]Affiliations: BI Norwegian Business School Centre for Applied Macroeconomics and Commodity Prices (BI CAMP) and Norwegian Centre for Energy Transition Strategies (NTRANS).

1 Introduction

In light of climate change, environmental economists have devoted increasing attention to the investigation and quantification of the effect of temperatures and other climatic variables on human health and mortality. It is conjectured that the largest share of the global cost of climate change will be the adverse effects on human mortality (Hsiang et al., 2017), both through increased mortality and through increased expenditures on technologies to reduce exposure to high temperatures (Deschênes & Greenstone, 2007). For example, Carleton et al. (2018) estimates that by the end of the century, the impact of climate change on mortality will equal 3.7% of global GDP.

This has spurred environmental economists to estimate the causal effect of temperatures on mortality. The goal of this research is often to use the estimated temperature-mortality relationship to predict the effect of future climate change on mortality (e.g. Carleton et al., 2018). Although recent innovations in climate economics has led to improvements in the estimation of the temperature-mortality relationship (Dell, Jones, & Olken, 2014), there still remains several issues, with one central issue being assumptions made about the functional form. This uncertainty regarding the true functional form is partially solved by estimating the temperature-mortality relationship using semi-parametric models (e.g Barreca, Clay, Deschênes, Greenstone, & Shapiro, 2016). However, estimating these models is very demanding on the data, and creates a tension between model flexibility and precision in estimation.

This paper develops a procedure that enables the use of machine learning for estimation of the causal temperature-mortality relationship. This has two improvements compared to previous studies. First, the use of machine learning solves the tension between allowing a flexible functional form while maintaining precision in estimation. While traditional regression analysis requires us to assume a functional form between our variables of interest, machine learning, on the other hand, uses an optimization criterion to discover the functional

form between the outcome variable and the covariates. This “data mining”-nature of machine learning might be perceived as worrying in traditional econometrics. However, one can argue that data mining is already occurring in empirical research, with researchers cherry-picking the models presented to their audience (Athey, 2018). At least in machine learning, there is an explicit and transparent criterion for choosing the preferred model specification.

Secondly, machine learning is essentially models for prediction; given an outcome variable and covariates, the algorithm finds the model that best explains the out-of-sample variation in the outcome variable.¹ Machine learning can thus improve the out-of-sample prediction power of the model of the temperature-mortality relationship. It is, however, not obvious how machine learning can be used when the goal is not prediction but causal inference. This remains largely an underexplored subject.²

The innovation in this paper is to demonstrate how machine learning can be used for causal inference in the presence of high-frequency panel data, as well as testing the statistical significance of the model chosen by the machine-learning algorithm. More specifically, this paper compares a semi-parametric OLS model with an method from machine learning known as regression trees. Although the regression tree is a popular method within machine learning, it is a novel method in empirical economics. The purpose of this paper is twofold. First, to introduce machine learning into the field of environmental economics; secondly, to compare the results from machine learning to those of traditional regression analysis. This will both clarify how machine learning can extend the climate literature, which is largely based on the use of OLS models, as well as clarifying the limitations of using machine-learning methods in causal inference.

Despite the fact that both the OLS model and the regression tree in this paper are models of the causal temperature-mortality relationship, they deliver opposite predictions

¹This is known as supervised learning, as opposed to unsupervised learning which are models used for classification.

²One exception being the work done on using machine learning in the causal estimation of treatment effects (see e.g. Athey, Imbens, Pham, & Wager, 2017; Wager & Athey, 2018; Athey & Imbens, 2016).

of the effect on mortality caused by increased temperatures in the future. The OLS model emphasizes the adverse effect on mortality by an increase in the occurrence of days with high temperatures, and thus predicts an increase in the number of temperature-induced fatalities. The regression tree, on the other hand, emphasizes the beneficial effect on mortality by a decrease in the occurrence of days with low temperatures, and thus predicts a decrease in fatalities.

I show that there are mainly two reasons for this discrepancy between the predictions made by the two models. First, the OLS model and the regression tree differ in their abilities of extrapolation outside of the temperature range used as the identifying variation in estimation; secondly, while the OLS model estimates the marginal effect of temperatures, the regression tree does not separate between the *marginal* effect and the *total* effect of temperatures on mortality.

The inability to extrapolate is a feature specific to regression trees. If the ability to extrapolate is an essential feature of the chosen model, other machine-learning algorithms can be used. The confounding of the marginal and total effect of covariates on the outcome variable is, however, a general feature of machine learning, and it is important to keep this in mind when using machine learning for causal inference.

The layout of this paper is as follows: section 2 gives an overview of the theory of the temperature-mortality relationship, as well as introducing the theory of machine learning and regression trees. The data used in this paper is from Barreca et al. (2016), and section 3 gives a brief explanation of their data. Section 4 offers a more in-depth explanation of how the semi-parametric OLS model and the regression tree are used for estimating the temperature-mortality relationship, while in section 5, the estimation results are presented. Section 6 presents the effect on the number of temperature-induced fatalities caused by an increase in mean temperatures, predicted by the two models. Sections 7 and 8 give a discussion of the findings in the paper and concludes.

2 Theory

The contribution in this paper is related to, but also differs from the previous research in climate economics. To clarify this, this section starts by discussing the existing methodology and empirical literature in this field, and then explains how new insights can be gained with the use of machine learning.

2.1 The new climate-economy literature

The idea that a country’s climate can affect its economic prosperity is not a new one. However, until recently, it was difficult to distinguish between the many confounding factors that are potentially correlated with both climate variables and economic activities. Although economists have long been aware of a strong correlation between climate variables and economic variables, because of omitted variable bias it has been hard to establish a causal relationship. That is, until the development of what Dell et al. (2014) has called the new climate-economy literature. According to Dell et al. (2014), in order to understand the impact of climate on the economy, we would ideally like to determine the following unknown functional relationship

$$\mathbf{y} = f(\mathbf{C}, \mathbf{X}) \tag{1}$$

which links vectors of climatic variables (\mathbf{C}) and other variables (\mathbf{X}) to outcomes, \mathbf{y} . The vector \mathbf{X} contains variables that are correlated with both \mathbf{y} and \mathbf{C} , and which we must control for in order to capture the causal relationship between our variable of interest and climatic variables.

Dell et al. (2014) argue that causal estimation of the model above is possible through the use of weather shocks as a source of exogenous variation in climatic variables. The regression model used in most of the climate-economics literature takes the following form

$$y_{it} = \beta \mathbf{C}_{it} + \gamma \mathbf{Z}_{it} + \mu_i + \theta_t + \varepsilon_{it} \tag{2}$$

where t indexes time and i indexes the spatial unit of analysis. The model above estimates the effect of climatic variables (\mathbf{C}) on the outcome of interest (y), while controlling for confounding factors (\mathbf{Z}). In addition, the model includes a unit-fixed effect, μ_i , which absorbs fixed characteristics for each spatial unit, and a time-fixed effect, θ_t , which absorbs common trends across spatial units. If correctly specified, the model exploits the random variation in weather over time within a given spatial unit, which is arguably orthogonal to other confounding factors.

In recent years, there has been a wave of empirical research exploiting high-frequency variation in temperature, precipitation, and other climatic variables over time, as a source of exogenous variation in climate. The outcome variable of interest ranges from agricultural output (Schlenker & Roberts, 2009; Feng, Krueger, & Oppenheimer, 2010), productivity (Jones & Olken, 2010; Cachon, Galliano, & Olivares, 2012), crime (Jacob, Lefgren, & Moretti, 2007; Ranson, 2014), health (Deschênes, Greenstone, & Guryan, 2009; Burgess, Deschenes, Donaldson, & Greenstone, 2017), and political stability (M. Burke, Miguel, Satyanath, Dykema, & Lobell, 2009; P. Burke & Leigh, 2010). The literature has documented the many channels through which human well-being and performance are affected by the climate, from reduced labour productivity (Niemelä, Hannula, Rautio, Reijula, & Railio, 2002) to a deterioration of our mental well-being, causing increased aggression among social media users (Baylis et al., 2018), and increased suicide rates (M. Burke et al., 2018).

It is conjectured that the largest share of the costs associated with global warming will come from the effect of increased exposure to hot temperatures on human morbidity and mortality (Hsiang et al., 2017). The epidemiology literature has documented the different channels that connects temperature exposure to human health, and it has found a strong association between extreme temperatures and excess mortality (see Deschênes (2014) for a review). This has spurred environmental economists to investigate and quantify the causal effect of temperatures on mortality by estimating versions of equation 2.

Studies have found that both extremely low and high temperatures have an adverse

effect on mortality rates (e.g. Deschênes & Greenstone, 2011; Barreca, 2012; Barreca et al., 2016; Burgess et al., 2017). In addition, the temperature-mortality relationship seems to be characterized by a considerable amount of age-heterogeneity, with the adverse effect of temperatures on mortality found largely among infants and the elderly, at least in developed countries (Deschênes & Greenstone, 2011). However, much of the estimation of the temperature-mortality relationship uses data from developed countries, which produces estimates that are unlikely to be representative for large part of the world. In fact, studies have found temperatures to have a much more severe impact on mortality in developing countries (Burgess et al., 2017; Carleton et al., 2018).

2.2 Machine learning

Machine-learning techniques are gaining traction among economists,³ especially in the field of prediction policy problems (Kleinberg, Ludwig, Mullainathan, & Obermeyer, 2015; Goel, Rao, & Shroff, 2016). When the problem at hand is a prediction problem, it is more or less straightforward how machine learning can be utilized in estimation. The same, however, cannot be said for problems of causal inference. When the goal is causal inference, it is a common perception that we should choose the preferred model specification prior to estimation, preferably based on economic theory and principles. Machine learning, on the other hand, has internalized data mining as a core principle of the field. In machine learning, the goal is to estimate and compare many models, and then pick the model that performs the best according to an explicit criterion.

A main difference between traditional econometrics and machine learning is the focus on out-of-sample fit. In machine learning, there is always a fear of overfitting the model to the specific data used to estimate the model. Let $X \in \mathbb{R}^p$ denote a real valued random input vector and $Y \in \mathbb{R}$ a real valued random output vector, and assume that we

³Overviews on the use of machine learning in economics are given by Varian (2014); Athey (2017); Mullainathan and Spiess (2017).

have the model $Y = f(X) + \epsilon$ where $E(\epsilon) = 0$ and $Var(\epsilon) = \sigma_\epsilon^2$. We can then derive an expression for the expected prediction error of a regression fit $f(X)$ at an input point $X = x_0$, using squared-error loss:

$$\begin{aligned}
 Err(x_0) &= \mathbb{E}[(Y - \hat{f}(x_0))^2 | X = x_0] \\
 &= \sigma_\epsilon^2 + [\mathbb{E}\hat{f}(x_0) - f(x_0)]^2 + \mathbb{E}[\hat{f}(x_0) - \mathbb{E}\hat{f}(x_0)]^2 \\
 &= \sigma_\epsilon^2 + Bias^2(\hat{f}(x_0)) + Var(\hat{f}(x_0)) \\
 &= Irreducible\ error + Bias^2 + Variance
 \end{aligned}$$

There is generally a trade-off between the model bias and model variance in the prediction error of a model. When the model is more complex, the model will capture more of the structure between Y and X, reducing the bias in the model. However, a more complex model also runs the risk of being fitted too hard to the data used in the estimation of the model, and thus the model could have a high variance with regards to the specific data used for estimation.

In machine learning, the model complexity is chosen in order to strike a balance between model bias and model variance. An increase in model complexity will lead to an increase in the in-sample fit of the model. However, with too much fitting, the model adapts itself too closely to the specific data used in estimation, and it will not generalize well out-of-sample. This bias-variance trade-off has caused sample splitting to become a crucial part of machine learning, where models are estimated using only parts of the data, and then evaluated using the remaining data.

The complexity of the model is adjusted through a tuning parameter, α , and so the predicted functional form can be written as $\hat{f}_\alpha(X)$. In machine learning, the optimal model is always the model that produce the lowest out-of-sample prediction error. Many different methods have been developed in order to estimate a model's prediction error, but the simplest and most widely used method is cross-validation.

[Figure 1 here]

In cross validation, the data sample is split into a certain number of equally sized folds. In the case of five-fold cross-validation, the data sample is split into five parts. The model is estimated using four out of the five folds, while the fifth fold is used to calculate the prediction error of the model. This is done five times, each time using a different fold as the holdout sample used for testing the model’s prediction accuracy. The accuracy of the model is then given by the average of the prediction errors calculated in the five different holdout samples. Using cross-validation in order to determine the complexity of the model is an example of a data-driven way to balance the expressiveness of the model with the fear of overfitting the model to the specific data used to estimate the model. This will generally lead to increased external validity of the model as compared to a priori specification of the complexity of the model.

2.2.1 Regression trees

Because of their transparency, regression trees are very popular models in machine learning.⁴ The idea behind these trees is quite simple: partition the covariate space into mutually exclusive regions, and then fit a simple model within each region. The tree-building algorithm fits the following regression model

$$f(x) = \sum_{j=1}^J c_j \mathbb{I}\{x \in \mathbb{R}_j\}$$

where the covariate space has been partitioned into J mutually exclusive regions, with the partitions being found through the minimization of the sum of squares $\sum (y_i - f(x_i))^2$. In each of these partitions, the algorithm fit a simple model consisting of only a constant, c_j . This fitted constant will simply be the mean outcome across all observations that fall within that region, $\hat{c}_j = \text{ave}(y_i | x_i \in \mathbb{R}_j)$.

⁴For a thorough explanation of regression trees, see Hastie, Tibshirani, and Friedman (2001); James, Witten, Hastie, and Tibshirani (2013).

Finding the best partitions of the covariate space in terms of minimizing the residual sum of squares is generally infeasible. Instead, the partitions are found through a greedy algorithm known as recursive binary splitting. This algorithm successively splits the covariate space into sub-regions. This process of splitting the covariate space into smaller sub-regions can be illustrated as a tree. Figure 2 illustrates a simple regression tree consisting of nodes and branches, where each question is a node, and with questions linked through a hierarchy of branches. Each question in the tree leads you either to the next question or to a terminal node with the predicted outcome. These terminal nodes are known as the *leaves* of the tree.

[Figure 2 here]

The algorithm starts at the top of the tree with all of the observations in the root node, the algorithm searches through all covariates and all possible binary splits of each covariate, and then chooses the covariate and the corresponding split-point that leads to the largest reduction in the residual sum of squares. Depending on the value on the covariate chosen by the algorithm, observations are allocated either down the left branch or down the right branch to the next node. All observations now belong to one of two nodes, with the nodes representing mutually exclusive regions of the covariate space.

This process of splitting the covariate space into mutually exclusive regions continues until some stopping criterion is met, in which case, we have reached the leaves of the tree. Instead of asking a question, the tree now makes its prediction. Regression trees are a fully non-parametric estimation technique, but since it does not try to match observations on all of the covariates, it does not suffer from the curse of dimensionality as many non-parametric estimation techniques do (e.g. Horowitz, 2009). For a toy example illustrating the non-parametric nature of regression trees, see appendix C.

Building a tree through the process described above will usually lead to a tree with poor out-of-sample prediction accuracy. That is because the tree will overfit to the data used to grow the tree, and thus the tree will be too complex. A simpler tree with fewer splits will

often have a lower variance, and thus better out-of-sample prediction accuracy, despite the bias introduced in the simpler tree. In order to avoid overfitting, some restrictions must be put on the tree. This process of restricting the complexity of the tree is known as *pruning* the tree.

In this paper, the tree will be pruned by restricting how many times a region can be split into sub-regions (i.e. the depth of the tree). This is done by growing trees of different depths and then choosing the optimal depth through a 5-fold cross-validation. The optimal tree (and corresponding depth) is chosen as the tree with the lowest average prediction error across the five folds.⁵

3 Data

In order to demonstrate the use of machine learning, compared to OLS, in the estimation of the temperature-mortality relationship, this paper uses the data and baseline model from Barreca et al. (2016), henceforth referred to as BCDGS. In that study, the authors constructed monthly mortality rates and temperature variables for all US states over the period 1900-2004. By drawing on several sources, the authors compiled a comprehensive dataset on all-cause monthly mortality rates for all US states since 1900. The weather data was collected from the National Climatic Data Center (NCDC) Global Historical Climatology Network-Daily (GHCN-Daily), which is a database of daily climate summaries from land surface stations that are subject to a common set of quality assurance. The temperature variables were constructed by selecting only the weather stations that had no missing records in any given year, and then the station-level data was aggregated to the county level by taking an inverse-distance weighted average of all the selected stations that were located within a 300-kilometers radius of each county's centroid.⁶ The county-level variable was then aggregated up to the

⁵In addition, in order to be able to test the statistical significance of the leaves, each leaf is restricted to contain at least 40 state-year-month observations.

⁶The weights are the inverse of the squared distance to the county's centroid.

state-level by taking a weighted average of all the counties within a state, where each county was weighted by its population.

The temperature is defined as the mean daily temperature, and in order to capture potential nonlinearities in the temperature-mortality relationship, the temperature variable is divided into temperature *bins*. Each bin is defined as the number of days in a month where the daily mean temperature is within a certain temperature range. This allows the authors to construct a semi-parametric model where the marginal effect of temperatures on mortality has to be constant within each temperature range, but where each temperature range is allowed its own marginal effect on mortality.

The analysis in this paper will use only data from 1960 to 2004.⁷ Furthermore, although partitioning the temperature variable into many temperature bins allows for a very flexible estimation of the functional form of the temperature-mortality relationship, doing so is very demanding on the data and it becomes difficult to estimate the individual effect of each bin. Instead, in their main specification, BCDGS focus on the upper and lower tails of the daily temperature distribution by reducing the temperature variable to three “critical” bins and a reference bin. The three critical bins are the number of days with mean temperature below 40 °F, the number of days with mean temperature between 80 and 89 °F, and the number of days with mean temperature above 89 °F. The reference bin is thus the number of days with mean temperature between 40 and 79 °F. This paper follows BCDGS, and focuses on these three critical bins. For the remainder of the paper, these three bins will be referred to as the number of “cold”, “warm” and “hot” days, respectively, and with a day in the reference category referred to as a “normal” day.⁸

⁷The reason why is the sharp decrease in the temperature-mortality relationship after 1959 found by BCDGS in their original study, which they argue was caused by the increased use of residential air-conditioning.

⁸Over the period 1960-2004, an American experienced on average 73.6 number of cold days, 25.7 number of warm days, and 1.1 number of hot days each year. There is, however, substantial variation across states. While people in Washington state rarely experienced hot days during the year, people in Arizona experienced

4 Method

4.1 Baseline OLS model

In line with the new climate-economy literature, the effect of temperatures on mortality is identified by using the interannual variation in temperature within states. The baseline OLS model uses the same model specification as in BCDGS

$$y_{sym} = \sum_j \theta^j TMEAN_{sym}^j + \mathbf{X}_{sym}\beta + \alpha_{sm} + \rho_{ym} + \varepsilon_{sym} \quad (3)$$

where y_{sym} is the monthly mortality rate in state s , year y and month m . The vector of control variables, \mathbf{X}_{sym} , includes indicators for unusually high or low levels of precipitation,⁹ the share of a state’s population in one of four age categories, and the log of per capita income. The age shares and log of per capita income are interacted with month indicators in order to capture age- and income-specific seasonality effects that are common across all states. The vector also contains a quadratic time trend that is interacted with the state-month identifier, thus allowing each state to follow its own trend in seasonal mortality rates. The specification also includes state-by-month fixed effects, α_{sm} , and year-by-month fixed effect, ρ_{ym} . The state-by-month fixed effects are included to absorb unobserved, but permanent differences in the mortality rate between states, while allowing these differences to vary across month of the year. The year-by-month fixed effects captures the seasonality of mortality rates, and absorbs any random time shocks to the mortality rate that is constant across states.

The variables of interest are the temperature bins in $TMEAN_{sym}^j$. These bins capture the number of days in a state-year-month in which the mean temperature is in the j th bin. These bins have been restricted to three “critical” temperature bins and one reference bin, with the critical bins defined above as the number of cold, warm and hot days on average 37.1 hot days during the year. In return, Arizona experienced on average only 11.5 cold days during, while people in Washington state were exposed to 65 such days during the year.

⁹Which indicates whether a state during a month experienced precipitation below the 25th or above the 75th percentile of their state-month specific distribution.

for each state-year-month observation. Estimating the model using temperature bins creates a functional form that is generally very flexible. However, the model is assuming a linear effect within each temperature bin, and it is not allowing for interaction effects between the temperature bins.

4.2 Regression tree

The advantage of the regression tree compared to the OLS model is that it does not require us to specify the particular form of the potential nonlinearities in the temperature-mortality relationship.¹⁰ This paper uses regression trees to estimate the ideal model from Dell et al. (2014) in equation 1 where the vector of climatic variables contains the temperature bins, the outcome variable is the mortality rate, and the vector \mathbf{X} contains the control variables and fixed effects from the baseline OLS model. The difference from the OLS model is that the functional form of the relationship between the temperature bins and the mortality rate is not specified, but instead left up to the algorithm to decide. The temperature bins are as before the number of cold, warm and hot days in a month.

The temperature-mortality relationship is characterized by the presence of many confounding factors that are correlated with both temperature and mortality rates. In addition, many of these confounding factors have a much higher correlation with mortality than temperatures, e.g. there is a strong seasonal cycle to mortality rates. Furthermore, even when the tree identifies the number of cold, warm or hot days as important predictors of mortality rates, it is not obvious whether the tree is capturing the causal effect of temperatures, or if it is simply capturing the confounding correlation between temperatures and, for example, income. A regression tree is just as plagued by omitted variable bias and the

¹⁰Whether a regression tree will provide a higher out-of-sample fit compared to an OLS model depends on the true underlying functional form. If the relationship between the covariates and the outcome variable is characterized by nonlinearities, then a regression tree will outperform OLS in terms of capturing the true functional form.

presence of confounding factors as an OLS model; the tree makes no contributions concerning problems of identification. This paper proposes that in the case of high-frequency panel data, by isolating the exogenous shocks to the variables of interest, it becomes possible to use machine-learning algorithms for causal inference. In the case of regression trees, the tree can identify the causal temperature-mortality relationship by first removing the influence of the confounding factors from the temperature and mortality variables. This can be accomplished by regressing the mortality and temperature variables on the confounding factors as specified in equation 3. This means estimating the following models

$$y_{sym} = \mathbf{X}_{sym}\beta + \alpha_{sm} + \rho_{ym} + \varepsilon_{sym} \quad (4)$$

$$TMEAN_{sym}^j = \mathbf{X}_{sym}\beta + \alpha_{sm} + \rho_{ym} + e_{sym} \quad (5)$$

where y_{sym} is the mortality rate, and $TMEAN_{sym}^j$ is the j th temperature bin. In these models, the mortality rate and the temperature bins are regressed on all of the control variables and fixed effects from the baseline OLS model. If we believe that the baseline OLS model is able to control for all of the confounding factors in the temperature-mortality relationship, then the residuals from estimating the models above for the mortality rate and temperature bins can be interpreted as the variation in these variables that is orthogonal to the confounding factors.

The regression tree-algorithm can then be used to estimate the following non-parametric model

$$\tilde{y}_{sym} = f(\widetilde{TMEAN}_{sym}^j) + u_{sym} \quad (6)$$

where \tilde{y}_{sym} and $\widetilde{TMEAN}_{sym}^j$ are the residuals in the mortality rate and the temperature bins from estimation of equations 4 and 5.¹¹ Using the residuals in these variables, the task of deciding the functional form on the effect of cold, warm and hot days on mortality is left up to the tree-building algorithm. Although regression trees are initially models for prediction, in the absence of confounding factors, correlations can be given a causal interpretation.

¹¹This procedure is inspired by the the Frisch-Waugh-Lovell theorem (Lovell, 2008).

Using the same data sample for both building the tree and estimating the outcome within each leaf is by Athey and Imbens (2016) labeled as adaptive estimation, and this procedure will not provide valid confidence bands on the predicted outcome in the leaves of the tree. In this paper, I use the procedure of honest estimation, developed by Athey and Imbens (2016), to place valid confidence intervals on the estimated outcome in each leaf. This requires splitting the data into two sub-samples. This is done through a randomization algorithm that allocates 75% of the states in the sample to a training set,¹² with the remaining 25% of states allocated to a test set.¹³ The confounding factors of the temperature-mortality relationship are then removed separately from the training sample and test sample.

5 Estimation

5.1 Baseline OLS model

Table 1 shows the estimated coefficients from the regression of the monthly mortality rate on the number of cold, warm and hot days experienced during a month. The coefficients in the table are the one-month accumulated effect of an additional day in a temperature bin, compared to experiencing an additional day with normal temperatures instead. The coefficients are interpreted as the change in the number of deaths per 100 000 people in a state caused by an additional cold, warm or hot day during the month. However, since the estimation is done using the random fluctuations in the number of cold, warm and hot days, the coefficients should be interpreted as the effect on mortality of having an additional

¹²States in training set: Alabama, Arizona, Arkansas, California, Connecticut, Delaware, Florida, Georgia, Indiana, Iowa, Kansas, Kentucky, Louisiana, Maine, Maryland, Massachusetts, Michigan, Minnesota, Mississippi, Missouri, Montana, Nebraska, New Hampshire, New Jersey, Oregon, Pennsylvania, South Carolina, South Dakota, Texas, Utah, Vermont, Virginia, Washington, West Virginia, Wisconsin and Wyoming.

¹³States in test set: Colorado, Idaho, Illinois, Nevada, New Mexico, New York, North Carolina, North Dakota, Ohio, Oklahoma, Rhode Island and Tennessee.

unexpected day in a temperature bin.

[Table 1 here]

Cold, warm and hot temperature shocks all have a statistically significant effect on the monthly mortality rate, with an additional day in each of the three temperature bins causing an increase in mortality. An additional unexpected hot day has the largest effect on mortality by increasing the mortality rate by an additional 0.33 deaths per 100 000 people, while an additional unexpected cold or warm day increases mortality by an additional 0.187 and 0.138 deaths per 100 000, respectively.

5.2 Regression tree

In order to place unbiased confidence interval on the predicted outcome in the leaves found by the regression tree, the residuals in temperature and mortality in the training sample are used for growing the tree, while the residuals in the test sample are used for testing the statistical significance of the temperature-mortality relationship found by the tree.

5.2.1 Growing the tree

Figure 3 illustrates the optimal tree for the temperature-mortality relationship found by the algorithm.¹⁴ It was argued that even though regression trees are models of prediction, in absence of confounding factors, mortality and temperatures are correlated only through a causal relationship. Furthermore, choosing the model complexity in order to balance the model bias with the model variance, we are limiting the possibility that the tree is capturing spurious correlations between temperatures and mortality specific to the data used to grow the tree (i.e. the training set).

¹⁴As explained in the method section, the optimal tree is found through a grid search using 5-fold cross-validation, and where optimal means the tree that minimizes the mean prediction error of the tree in the holdout across the five different holdout sets.

[Figure 3 here]

The optimal tree found by the algorithm is three levels deep, and it allocates the 19 440 state-year-month observations in the training sample to a total of seven leaves. The tree starts by asking whether a state in any given month had more than 1.323 unexpected cold days. If yes, then the observation is allocated down the right branch from the root node, if no, then the observation is allocated down the left branch. In this first split of the tree, 16 761 state-year-month observations are allocated down the left branch, while 2 679 state-year-month observations are allocated down the right branch.

We start by looking at the state-year-month observations that went down the left branch after the initial split in the root node. These observations are allocated to four different leaves, with each leaf representing a unique partition of the covariate space. Since the allocation of an observation to a specific leaf depends on the number of unexpected cold, warm and hot days experienced by that state during a month, these leaves can be thought of as temperature leaves. There are in total 631 state-year-month observations allocated to temperature leaf 1. These are the states that experienced at least 3.543 *fewer* cold days than expected, and for these states the tree predicts a decrease in mortality of 0.987 fewer deaths per 100 000. The states in temperature leaf 2 are the states that experienced between -3.543 and 1.323 unexpected cold days, and at most 0.597 unexpected warm days. There are 13 422 observations in this leaf, and the tree predicts that these states will experience 0.145 fewer deaths per 100 000.

States that experience at most 1.323 unexpected cold days and 0.071 unexpected hot days, but more than 0.597 unexpected warm days are allocated to temperature leaf 3. There are in total 1 925 observations in this leaf, and the tree predicts that these states will experience an additional 0.176 deaths per 100 000. The states experiencing at most 1.323 unexpected cold days, but experiencing more than 0.597 unexpected warm days and 0.071 unexpected hot days, are allocated to temperature leaf 4. There are 783 state-year-month observations in this leaf, and the tree predicts that these states will experience an even larger

increase in mortality with an additional 0.742 deaths per 100 000.

We then look at the state-year-month observations that went down the right branch after the initial split in the root node. These states experienced more than 1.323 unexpected cold days, and the tree allocates these observations to three mutually exclusive temperature leaves. There are 426 state-year-month observations in temperature leaf 5; these states experienced at most 7.495 unexpected cold days and -0.023 unexpected hot days. In other words, these states are experiencing some unexpected cold days, but they are experiencing fewer hot days than expected. For these states, the tree predicts an increase in the mortality rate of 1.048 additional deaths per 100 000.

There are 2 212 state-year-month observations allocated to temperature leaf 6. States in this leaf are experiencing both some unexpected cold days and some unexpected hot days, and while the tree predicts an increase in mortality for these states as well, the increase of 0.545 additional deaths per 100 000 is smaller than the predicted increase in leaf 5. In other words, although unexpected hot days might have an adverse effect on mortality, it can also work to mitigate the harmful effects of cold shocks on mortality.

In temperature leaf 7, we find the states that experienced more than 7.495 unexpected cold days during a month. Given the magnitude on the temperature shock experienced by states in this leaf, there are only 41 state-year-month observations in the training sample allocated to this leaf. The tree predicts the largest increase in mortality for states in this leaf with an additional 2.472 deaths per 100 000.

5.2.2 Statistical significance of the tree

When the goal is causal inference, using the same data to grow the tree and estimate the outcome in each leaf (i.e. adaptive estimation), will often lead to an overestimation of the statistical significance of the leaves found by the regression tree (Athey & Imbens, 2016). In honest estimation, valid confidence bands for the leaf outcome are found by estimating the outcome in each leaf using a test sample, but where the leaves are found from the regression

tree grown using the training sample.

Honest estimation can be implemented by using the test sample and estimating an OLS model where the mortality rate is regressed on dummy variables for the temperature leaves found by the tree-building algorithm. Intuitively, since the test sample was not used in the tree-building process when the leaves of tree were uncovered, the test data will provide us with unbiased estimates on the effect of the temperature leaves on mortality. Table 2 shows the regression result from estimating the mortality rate on the temperature leaves using adaptive and honest estimation.

[Table 2 here]

Column (1) in table 2 shows the results from estimating the effect of the temperature leaves on the mortality rate, when the OLS model is estimated using the training data, (i.e. adaptive estimation). From the table we can see that the estimates for each leaf are equal to the estimates predicted by the regression tree within each leaf. This is not surprising since the regression tree uses the mean outcome across all observations within a leaf as the predicted outcome of belonging to a leaf, which corresponds to fitting a simple OLS model consisting of only a constant term within each leaf.

Column (2) in table 2 shows the results from estimating the same model, but where the estimation is done using the test sample (i.e. honest estimation). Honest estimation delivers slightly different estimates of the effect of the temperature leaves on mortality than adaptive estimation. For instance, the model estimated using the training set finds that states in temperature leaf 2 will experience 0.145 fewer deaths per 100 000, while estimating the same model using the test data delivers an estimate of 0.085 fewer deaths per 100 000 of being in temperature leaf 2. Nevertheless, for most of the temperature leaves, the difference is small between honest and adaptive estimation of the effect of the temperature leaves on mortality. Furthermore, for most of the temperature leaves, the effect on mortality remains statistically significant even when the leaf outcome is estimated using honest estimation.

The only exception is temperature leaf 3, which becomes slightly negative and statistically insignificant, when the model is estimated using the test sample.

Column (1) and (2) show the result from unweighted regressions. This means that the estimates can be interpreted as the effect on mortality for the average *state* from being in a certain temperature leaf during a month. However, what we often are interested in is the effect of temperatures on mortality for the average *person*. This interpretation of the estimates can be achieved through weighting each state-year-month observation with the state population.¹⁵ Column (3) shows the results from honest estimation of the effect of the temperature leaves on mortality, but where each state-month-year observation is weighted by the state population.

Comparing honest estimation of the temperature leaves between the unweighted and weighted OLS model, we see that for most of the temperature leaves there is only a slight change in the magnitude of the coefficients. For instance, in the unweighted regression, it was estimated that being in temperature leaf 5 caused an additional 0.614 deaths per 100 000, while in the weighted regression, the increase in mortality has fallen to 0.387 additional deaths per 100 000. This decrease in the coefficient on temperature leaf 5 causes this temperature leaf to lose its statistically significant effect on mortality, while temperature leaf 3 remains statistically insignificant in both regressions. All other temperature leaves maintain their statistically significant effect on mortality between the unweighted and weighted regression. Given the interpretation of the regression results in column (3) in table 2, this is the preferred model specification of the temperature leaves found by the regression tree.

¹⁵When building the regression tree, it is possible to weight each state-month observation by the state-population. It is, however, not obvious that observations should be weighted when the goal is to uncover the functional form of the temperature-mortality relationship, as opposed to estimating the leaf effects.

6 Prediction

This paper has argued that the baseline OLS model and the regression tree are both able to capture the causal relationship between temperatures and mortality. However, OLS models and regression trees are fundamentally different ways of estimating the temperature-mortality relationship, with the former requiring us to make a priori assumptions on the functional form, while the latter being essentially data mining. In order to assess the implication of using the two methods for estimating the temperature-mortality relationship, this section compares the predicted effect of increased temperatures caused by climate change on mortality using the results from the baseline OLS model and the regression tree.

Climate projections are acquired from the Global Daily Downscaled Projections (GDDP) data created by NASA Earth Exchange (Thrasher, Maurer, McKellar, & Duffy, 2012). This data is compiled from the General Circulation Model runs conducted under the Coupled Model Intercomparison Project Phase 5, with projections provided for two of the Representative Concentration Pathways (RCP4.5 and RCP8.5). The NEX-GDDP data consists of high-resolution ($0.25^\circ \times 0.25^\circ$)¹⁶ bias-corrected climate projections throughout the year 2100, generated from 21 different climate models. These projections provide the daily minimum and maximum temperature at each grid point for the two standardized emission scenarios.

Given that the goal of this paper is not to produce an estimate on the impact of increased temperatures on mortality, but instead to compare different estimation techniques of the temperature-mortality relationship, this paper uses climate projections from only one of the 21 available models in the NEX-GDDP data.¹⁷ The daily climate projections from this model is acquired for the RCP8.5, an emission scenario with intense growth in fossil fuel emission throughout the century, which corresponds to a mean global surface warming of

¹⁶This resolution equals approximately 25×25 km at the equator.

¹⁷The model used is the INMCM4.

approximately $4 - 6\text{ }^{\circ}\text{C}$ ¹⁸ above pre-industrial levels by 2100. Daily temperature projections are acquired for each grid point starting in the year 2080 and throughout the year 2099. At each grid point, the number of cold, warm and hot days is counted for each month during this 20-year period. The data is then aggregated up to the county-level by taking a simple average across all grid points that fall within a county. In order to aggregate the data from the county-level to the state-level, a population-weighted¹⁹ average is taken across all counties within a state. The final data set of climate projections is then the number of cold, warm and hot days for each state in each month in the period 2080-2099.

In order to estimate the effect of climate change on mortality in the future, we need an estimate on the change in the number of cold, warm and hot days experienced by states in the future. Climate change is defined as the change in the number of cold, warm and hot days, for each state-month observation in the 2080-2099 period, compared to the state-month average found for the baseline period 1960-2004.²⁰

6.1 Climate change and mean temperatures

Using the climate projections, we can compare how climate change under RCP8.5 affects the average number of cold, warm and hot days in the period 2080-2099, compared to the period 1960-2004. Figure 4 illustrates the difference for each state in the average annual number of cold, warm and hot days between the two periods. According to the climate projections, each state will experience a decrease in the average annual number of cold days, with the decrease being larger for states that are currently experiencing many cold days during the

¹⁸Which corresponds to $39.2 - 42.8\text{ }^{\circ}\text{F}$

¹⁹The weight given is the share of the US population for each county in 2004, using population data from the National Cancer Institute (2017), which means that it is assumed that the distribution of people in the US remains the same after 2004.

²⁰Thus, it is assumed that although climate change causes an increase in mean temperatures, expectations towards temperatures remain the same; meaning that the increase in the number of warm and hot days caused by climate change will come unexpectedly.

year.

[Figure 4 here]

When it comes to the change in the average annual number of warm days, the pattern across states is less clear; however, almost every state is predicted to experience an increase in the number of warm days. The only exception being Arizona, which is predicted to experience a decrease in the number of warm days. At the same time, Arizona is also the state predicted to experience the largest increase in the number of hot days. When it comes to the number of hot days, all states are predicted to experience an increase in the average annual number of hot days, with the increase being larger for the states that are already experiencing relatively many hot days during the year.

[Figure 5 here]

Using the estimated temperature changes, figure 5 illustrates the average annual distribution of state-month observations across the temperature leaves defined by the regression tree from the previous section. Given the projections from the climate model, states will in the future find themselves more frequently in temperature leaf 1, while they will find themselves less frequently in temperature leaf 2, compared to the period 1960-2004. The sharp increase of state-month occurrences in temperature leaf 1 is caused by the decline of the number of cold days experienced by all states in the future, which causes a shift of states out of leaf 2 and into leaf 1.

There will also be an increase in the number of state-month occurrences in temperature leaves 3 and 4 caused by the increase in the number of warm and hot days in the future, and a decrease of occurrences in temperature leaves 5, 6 and 7 caused by the reduction in the number of cold days. Furthermore, the change in the number of occurrences of the different temperature leaves will have counteracting effects on mortality. E.g. from column (3) in table 2 we see that the regression tree predicts a decrease in mortality for both temperature

leaves 1 and 2, thus the increase in the occurrence of the former will reduce mortality, while the reduction in the occurrence of the latter will increase mortality.

6.2 Climate change and mortality

Given the estimated change in the number of cold, warm and hot days, we can predict the impact on the mortality rate for each state-month observation using the estimates on the temperature-mortality relationship found from the baseline OLS model and the regression tree.²¹ For a state-level analysis of the average annual number of temperature-induced fatalities, see appendix D.

Table 3 compares the average annual number of temperature-induced fatalities predicted by the two models. For the baseline OLS model, the change in mortality can be broken down across the estimated change in the number of cold, warm and hot days. The regression tree, on the other hand, does not provide us with the marginal effect on mortality from each of the temperature bins; instead, it estimates the effect on mortality from being in a particular temperature leaf. The change in number of temperature-induced fatalities predicted by the regression tree is the result of a change in the distribution of state-month observations across the leaves as shown in figure 5.

[Table 3 here]

The baseline OLS model illustrates the fact that climate change will have both a beneficial and adverse effect on mortality. The beneficial effect comes from a reduction in the number of cold days, while the adverse effect comes from the increase in the number of warm and hot days. Which of these two effects dominate is an empirical question. According to the baseline OLS model, the increase in the number of fatalities caused by warm and hot

²¹Using population estimates for total US population in 2080-2099 (United Nations, 2017), the predicted impact on the mortality rate can then be used to predict the impact on the number of temperature-induced fatalities.

temperatures heavily outweighs the beneficial effect on mortality from the decrease in the number of cold days, and thus the model predicts an increase in the average annual number of temperature-induced fatalities in the US in the period 2080-2099.

The regression tree, on the other hand, finds that the beneficial effect of reduced number of cold days outweighs the adverse effect of increased number of warm and hot days, and thus there will be a decrease in the number of temperature-induced fatalities in 2080-2099. The increase in the number of warm and hot days will cause an increase in the occurrence of state-month observations in leaves with a predicted increase in mortality (e.g. leaf 3). However, this adverse effect on mortality will be outweighed by the increased occurrence of states in leaves with a predicted decrease in mortality caused by the decrease in the number of cold days (e.g. leaf 1).

7 Discussion

This paper has argued that even though regression trees are models of prediction, in absence of confounding factors, mortality and temperatures are correlated only through a causal relationship. While the OLS model requires us to assume a functional form between temperatures and mortality, the regression tree is a fully non-parametric estimation technique without the curse of dimensionality. Furthermore, choosing the optimal model specification while considering the out-of-sample fit of the model, we avoid overfitting the model to the specific data used for estimation. In fact, in their study, BCDGS found a sharp decline in the US temperature-mortality relationship after 1959, which means that the estimated temperature-mortality relationship for the first half of the previous century had limited external validity even for the second half of the same century.

While the procedure developed in this paper allows us to use regression trees to model the causal temperature-mortality relationship, the OLS model and the regression tree deliver very different predictions of the effect of climate change on the number of temperature-

induced fatalities in the future. There are two main reasons for this discrepancy between the predictions of the models. First, the OLS model allows extrapolation outside of the temperature range used as the identifying variation in the estimation; regression trees, on the other hand, are not suited for extrapolation. Secondly, while the OLS model estimates the marginal effect of cold, warm and hot days on mortality, the regression tree is free to choose its own model specification. When choosing the optimal model, the tree makes no separation between the marginal effect of cold, warm and hot days, and the total effect of such days on mortality.

In order to illustrate the effect of extrapolation, I re-run the prediction analysis in the previous section, but where I remove state-month observations in the future that are outside of the range of temperature shocks observed in the baseline period 1960-2004.²² Table 4 shows the predicted changes in the number of temperature-induced fatalities by the two models. Removing observations that are not within the range of variation that the model has been estimated on has very little effect on the predictions made by the regression tree. For most temperature leaves, the predicted change in the number of temperature-induced fatalities remains approximately the same, and there is only a small increase in the total number of temperature-induced fatalities made by the tree. For the OLS model, on the other hand, there is a large drop in the predicted number of temperature-induced fatalities. The number of temperature-induced fatalities caused by cold days stays more or less the same, while there is a substantial reduction in the number of temperature-induced fatalities caused by warm and hot days.

[Table 4 here]

²²From the training data, I find the lower and upper limit on the number of unexpected cold, warm and hot days (i.e. the residuals) observed in the period 1960-2004. State-month observations in the 2080-2099 period with an estimated change in the number of cold, warm or hot days above the upper limit or below the lower limit are dropped. This causes a reduction in the number of state-month observations from 11 520 to 10 339.

From table 4 we can see that the majority of the difference in the predictions made by the OLS model and the regression tree is caused by the extrapolation of the OLS model. Since the OLS model provides a marginal effect for each temperature bin, it is easy to extrapolate and estimate the effect of climate change on mortality far outside the range of variation where the marginal effects have been identified. In the tree, however, the thresholds for the covariates used to allocate observations to their respective leaves are constrained by the range of variation in the data used to grow the tree. Whether an observation is slightly above a threshold or very far above it, the observation will still end up in the same leaf, and thus receive the same predicted change in mortality.

When using machine learning, it is important to consider whether the ability to extrapolate is an important feature of the chosen model. The inability to extrapolate is not a general feature of machine learning; instead, it is a feature specific to regression trees. When the research question at hand requires extrapolation, other machine-learning algorithms can still be used.

Even when extreme observations are removed, the regression tree still predicts a large decrease in the number of temperature-induced fatalities in the future. This is because the regression tree emphasizes the harmful effect of cold temperatures on mortality, as opposed to the OLS model that estimates the marginal effect of cold, warm and hot days on mortality. Regression trees are allowed to choose which covariates to use in building the model, and if allowed to, they will use all of the covariates and model a complex dynamic between them. In order to avoid overfitting to the training sample, trees are pruned, causing the algorithm to choose only the covariates with the highest correlation with the outcome variable. This means that the choice of covariates by the tree is a function of both the marginal effect and the total effect of a covariate on the outcome variable.

Figure 6 illustrates the temperature-mortality relationship estimated by the OLS model and the average annual distribution of days across different temperature ranges. According to the OLS model, the marginal effect of a hot day is almost twice the marginal effect

of a cold day. However, from the bottom-panel in figure 6, we can see that most Americans experience very few hot days during the year, while they experience significantly more days with low temperatures. This means that although the marginal effect of a hot day might be large, the total effect of hot days on mortality is low compared to the total effect of cold days. In OLS, the total effect of covariates on the outcome variable affects the model only by affecting the precision of which the marginal effects can be estimated with.

[Figure 6 here]

The total effect of covariates on the outcome variable has a much more complex effect on the regression tree. Figure 7 gives a 2-dimensional illustration of the functional form of the temperature-mortality relationship uncovered by the tree. The covariate space is partitioned into seven mutually exclusive regions, and within each region, the tree predicts an outcome for the mortality rate. Figure 7 illustrates the fact that the regression tree does not estimate marginal effects; instead, the tree relies on the correlation between temperatures and mortality to decide on the partitions of the covariate space. The low total effect of hot temperatures on mortality means that there is a low correlation between hot temperatures and mortality, thus leading the tree to highlight the role of the number of cold days on mortality, as opposed to the role of the number of warm and hot days.

[Figure 7 here]

When using regression trees for causal inference, it is important to keep in mind the fact that the regression trees does not separate between the marginal and the total effect of covariates when deciding the partitions of the covariate space. Because of pruning, the regression tree may omit covariates that have a causal effect on the outcome variable, either because the total effect of the covariate on the outcome is small, or because covariates are correlated with each other, and thus the algorithm might only choose one of them. Furthermore, this confounding of the marginal and total effect of a covariate on the outcome

variable is not a problem specific to regression trees; instead, it is an issue general to machine learning. When the selection of the functional form is outsourced to an algorithm, the algorithm will choose only the covariates with the highest correlation with the outcome variable of interest.

8 Conclusion

This paper has illustrated the use of two very different estimation techniques, a traditional OLS model and the popular machine-learning algorithm known as regression trees, in the estimation of the temperature-mortality relationship. Regression trees are, as most machine-learning algorithms, models for prediction; however, it was argued that if the exogenous variation in temperature and mortality could be isolated, then the correlation discovered by the tree must in fact be the result of a causal relationship.

We can learn mainly two things from using machine-learning techniques in climate economics. First, regression trees allow us to estimate the temperature-mortality relationship without specifying any a priori assumptions on the functional form; secondly, regression trees are grown while considering the out-of-sample fit of the model, where the optimal complexity of the model is chosen in order to balance the model bias with the model variance. Although research questions in climate economics are often treated as problems of causal inference, in fact, in many cases, the objective of the research is to predict the effect of future climate change on our outcome variable of interest. Given the fact that regression trees are grown while considering the out-of-sample fit of the model means that machine learning could be better suited estimation technique when the goal is to predict the effect of future climate change.

When it comes to predicting the effect of an increase in temperatures caused by climate change on the number of temperature-induced fatalities, the OLS model and the regression tree make opposite predictions. While the OLS model predicts an increase in

the number of temperature-induced fatalities in the future, the regression tree predicts a decrease. One of the reasons for this discrepancy between the predictions made by the two models is the fact that the OLS model allows extrapolation outside the range of variation where the marginal effects have been estimated. Removing this channel of extrapolation has a considerable effect on the predictions made by the OLS model, while the predictions made by the regression tree remain largely unchanged.

Furthermore, while the OLS model estimates the marginal effects, the regression tree does not separate between the marginal effect and the total effect of the covariates on mortality when choosing the optimal model specification. Since the average American has a higher exposure to cold temperatures than to hot temperatures, the total effect of hot days on mortality is low compared to the total effect of cold days. This leads the regression tree to emphasize the harmful effect of cold temperatures on mortality as opposed to the effect of hot temperatures, and thus the increase in temperatures caused by climate change is found by the regression tree to be overall beneficial for mortality.

While the inability to extrapolate is not a general feature of machine learning, the confounding of the marginal and total effect of covariates on the outcome variable is. In machine learning, when the model selection is outsourced to an algorithm, the algorithm will choose to include in its model only the covariates with the highest correlation with the outcome variable. This means that the algorithm may omit covariates that have a high marginal effect on the outcome when the total effect of these covariates is low.

To conclude, this paper has demonstrated a procedure that enables the use of machine learning for causal inference, and it opens up an important venue for future research. In order to truly exploit the power of machine learning in causal inference, machine-learning techniques must be further developed to distinguish between the marginal and total effect of covariates on the outcome variable. When solved, the applicability of these techniques will improve climate economics, as well as find important applications well beyond this field.

Acknowledgements

This paper has greatly benefited from discussions with Costanza Biavaschi, Tamma Carleton, Colin Green, Amir Jina and Ragnar Torvik. I also thank workshop participants at the Norwegian University of Science and Technology and at the University of Chicago for their valuable comments. Parts of this research was conducted while visiting the University of Chicago Harris School of Public Policy.

References

- Athey, S. (2017). Beyond prediction: Using big data for policy problems. *Science*, *355*(6324), 483–485.
- Athey, S. (2018). The Impact of Machine Learning on Economics. In Agrawal Ajay K., J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*. University of Chicago Press.
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, *113*(27), 7353–7360.
- Athey, S., Imbens, G. W., Pham, T., & Wager, S. (2017). Estimating Average Treatment Effects: Supplementary Analyses and Remaining Challenges. *American Economic Journal*, *107*(5), 278–281.
- Barreca, A. (2012). Climate change, humidity, and mortality in the United States. *Journal of Environmental Economics and Management*, *63*(1), 19–34.
- Barreca, A., Clay, K., Deschênes, O., Greenstone, M., & Shapiro, J. (2016). Adapting to Climate Change: The Remarkable Decline in the U.S. Temperature-Mortality Relationship over the 20th Century. *Journal of Political Economy*, *124*(1), 105–159.
- Baylis, P., Obradovich, N., Kryvasheyev, Y., Chen, H., Coviello, L., Moro, E., ... Fowler, J. H. (2018). Weather impacts expressed sentiment. *PLOS ONE*, *13*(4).

- Burgess, R., Deschenes, O., Donaldson, D., & Greenstone, M. (2017). Weather and Death in India: Mechanisms and Implications of Climate Change. *EPIC Mimeograph*.
- Burke, M., González, F., Baylis, P., Heft-Neal, S., Baysan, C., Basu, S., & Hsiang, S. (2018). Higher temperatures increase suicide rates in the United States and Mexico. *Nature Climate Change*, 8(8), 723–729.
- Burke, M., Miguel, E., Satyanath, S., Dykema, J. A., & Lobell, D. B. (2009). Warming increases the risk of civil war in Africa. *Proceedings of the National Academy of Sciences*, 106(49), 20670–20674.
- Burke, P., & Leigh, A. (2010). Do Output Contractions Trigger Democratic Change? *American Economic Journal: Macroeconomics*, 2(4), 124–157.
- Cachon, G. P., Galliano, S., & Olivares, M. (2012). Severe Weather and Automobile Assembly Productivity. *Columbia Business School Research Paper*, 12(37).
- Carleton, T., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., . . . Zhang, A. T. (2018). Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits. *University of Chicago, Becker Friedman Institute for Economics Working Paper*(51).
- Dell, M., Jones, B. F., & Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature*, 52(3), 740–798.
- Deschênes, O. (2014). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, 46, 606–619.
- Deschênes, O., & Greenstone, M. (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review*, 97(1), 354–385.
- Deschênes, O., & Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4), 152–185.
- Deschênes, O., Greenstone, M., & Guryan, J. (2009). Climate Change and Birth Weight.

- American Economic Review*, 99(2), 211–217.
- Feng, S., Krueger, A. B., & Oppenheimer, M. (2010). Linkages among climate change, crop yields and Mexico–US cross-border migration. *Proceedings of the National Academy of Sciences*, 107(32), 14257–14262.
- Goel, S., Rao, J. M., & Shroff, R. (2016). Precinct or prejudice? Understanding racial disparities in New York City’s stop-and-frisk policy. *The Annals of Applied Statistics*, 10(1), 365–394.
- Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The Elements of Statistical Learning - Data Mining, Inference, and Prediction* (2nd ed.). New York: Springer.
- Horowitz, J. L. (2009). *Semiparametric and Nonparametric Methods in Econometrics*. New York, NY: Springer US.
- Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., . . . Houser, T. (2017). Estimating economic damage from climate change in the United States. *Science*, 356(6345), 1362–1369.
- Jacob, B., Lefgren, L., & Moretti, E. (2007). The Dynamics of Criminal Behavior: Evidence from Weather Shocks. *The Journal of Human Resources*, 42(3), 489–527.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning - with Applications in R* (1st ed.). New York: Springer.
- Jones, B. F., & Olken, B. A. (2010). Climate Shocks and Exports. *American Economic Review*, 100(2), 454–459.
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Obermeyer, Z. (2015). Prediction Policy Problems. *American Economic Review*, 105(5), 491–495.
- Lovell, M. C. (2008). A Simple Proof of the FWL Theorem. *The Journal of Economic Education*, 39(1), 88–91.
- Mullainathan, S., & Spiess, J. (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives*, 31(2), 87–106.
- National Cancer Institute. (2017). *U.S. Population Estimates, 1969–2016*. Retrieved from

<http://seer.cancer.gov/popdata/download.html>

- Niemelä, R., Hannula, M., Rautio, S., Reijula, K., & Railio, J. (2002). The effect of air temperature on labour productivity in call centres—a case study. *Energy and Buildings*, *34*(8), 759–764.
- Ranson, M. (2014). Crime, weather, and climate change. *Journal of Environmental Economics and Management*, *67*(3), 274–302.
- Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, *106*(37), 15594–15598.
- Thrasher, B., Maurer, E. P., McKellar, C., & Duffy, P. B. (2012). Technical Note: Bias correcting climate model simulated daily temperature extremes with quantile mapping. *Hydrology and Earth System Sciences*, *16*(9), 3309–3314.
- United Nations. (2017). *World Population Prospects 2017*. Retrieved from <https://population.un.org/wpp/Download/Standard/Population/>
- Varian, H. R. (2014). Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, *28*(2), 3–28.
- Wager, S., & Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, *113*(523), 1228–1242.

A Figures




Figure 1: Illustration of 5-fold cross-validation. Code for figure adapted from: https://github.com/amueller/introduction_to_ml_with_python/blob/master/mglearn/plot_cross_validation.py.

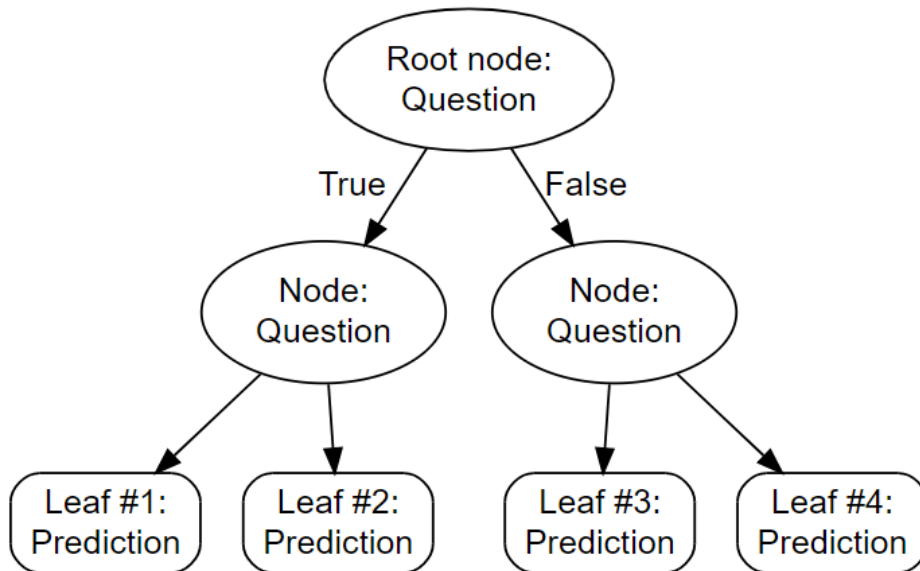


Figure 2: A simple regression tree.

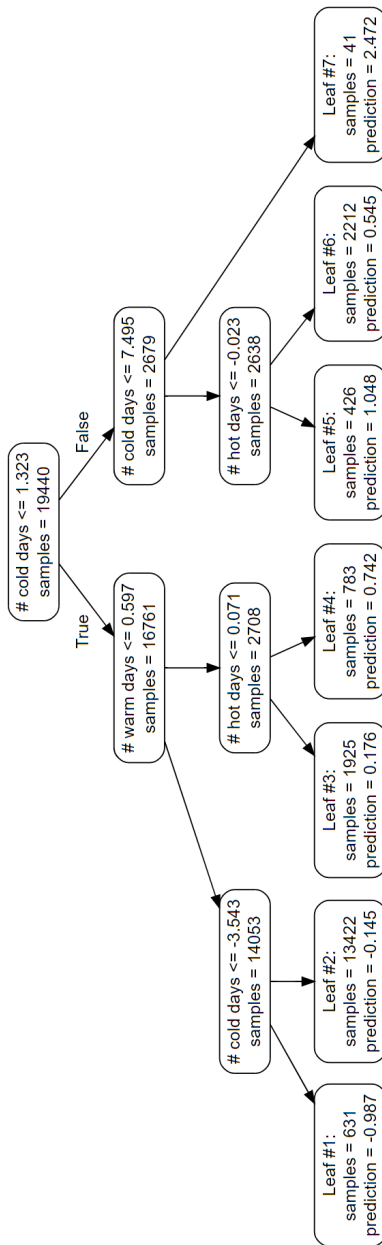


Figure 3: Regression tree for the US temperature-mortality relationship.




Figure 4: Predicted change in average annual number of (a) cold days, (b) warm days, and (c) hot days, between the periods 1960-2004 and 2080-2099.




Figure 5: Average annual distribution of state-month observations across the leaves defined by the regression tree (see figure 3) for the periods 1960-2004 and 2080-2099.




Figure 6: In (a), the dose-response function estimated by the baseline OLS model (see table 1). In (b), the population-weighted average annual distribution of days per year in ten different temperature ranges in the period 1960-2004.




Figure 7: 2-dimensional representation of the partitions of the covariate-space made by the regression tree. Number of cold days on the x-axis and number of warm days on the y-axis. In (a), the covariate space is illustrated when $\# \text{warm days} > 0.597$; while in (b), $\# \text{warm days} \leq 0.597$. The colorbar corresponds to the predicted effect on mortality in each leaf (see figure 3).

B Tables

Table 1: Estimates from the baseline OLS model of the impact of cold, warm and hot days on monthly mortality rate.

	(1)
Cold day	0.187*** (0.016)
Warm day	0.138*** (0.017)
Hot day	0.330*** (0.079)
N	26 460

Note: The dependent variable is the monthly mortality rate. The model includes all control variables and fixed effects specified in the baseline OLS model. State-year-month observations are weighted by state population, and standard errors are clustered by state. ***, **, * indicate statistical significance at 1%, 5% and 10%, respectively.

Table 2: Estimation from the regression tree of the impact of the temperature leaves on monthly mortality rate.

	(1)	(2)	(3)
Leaf 1	-0.987*** (0.121)	-0.988*** (0.328)	-1.042*** (0.159)
Leaf 2	-0.145*** (0.017)	-0.085*** (0.026)	-0.067*** (0.013)
Leaf 3	0.176*** (0.043)	-0.063 (0.089)	0.139 (0.098)
Leaf 4	0.742*** (0.081)	0.586*** (0.197)	0.462** (0.189)
Leaf 5	1.048*** (0.187)	0.614** (0.282)	0.387 (0.290)
Leaf 6	0.545*** (0.064)	0.571*** (0.124)	0.441*** (0.075)
Leaf 7	2.472*** (0.778)	2.430*** (0.529)	2.306*** (0.517)
Honest estimation:	No	Yes	Yes
Weighted:	No	No	Yes
N	19 440	7 020	7 020

Note: The dependent variable is the residual variation in the monthly mortality rate. Standard errors are clustered by state, and degrees of freedom are adjusted for the removal of confounding factors in previous regression. In weighted regression, state-year-month observations are weighted by state population. ***, **, * indicate statistical significance at 1%, 5% and 10%, respectively.

Table 3: Predicted average annual number of temperature-induced fatalities, 2080-2099.

	(1)		(2)
	Baseline OLS		Regression tree
<i>Temperature bin:</i>		<i>Temperature leaf:</i>	
Cold days	-17 979	Leaf 1	-12 072
Warm days	12 839	Leaf 2	-1 403
Hot days	22 681	Leaf 3	1 043
		Leaf 4	5 455
		Leaf 5	0
		Leaf 6	332
		Leaf 7	209
Total:	17 541		-6 436

Note: Column (1) and column (2) show the average annual number of temperature induced-fatalities for the US as a whole, predicted by the baseline OLS model and regression tree, respectively. For the OLS model, the prediction can be broken down across each temperature bin, while for the regression tree, the prediction can be broken down across each temperature leaf.

Table 4: Predicted average annual number of temperature-induced fatalities, 2080-2099, where observation outside range of temperature shocks observed in baseline period have been removed.

	(1)		(2)
	Baseline OLS		Regression tree
<i>Temperature bin:</i>		<i>Temperature leaf:</i>	
Cold days	-17 008	Leaf 1	-11 691
Warm days	9 666	Leaf 2	-1 288
Hot days	8 194	Leaf 3	993
		Leaf 4	4 165
		Leaf 5	0
		Leaf 6	332
		Leaf 7	209
Total:	852		-7 279

Note: Column (1) and column (2) show the average annual number of temperature induced-fatalities for the US as a whole, predicted by the baseline OLS model and regression tree, respectively. For the OLS model, the prediction can be broken down across each temperature bin, while the the regression tree, the prediction can be broken down across each temperature leaf.

C Toy example: Who survived the Titanic?

Using a data set with information on passengers on the Titanic, we can grow a regression tree in order to predict who survived the Titanic²³. In this case, the outcome variable is binary: 0 if you died and 1 if you survived²⁴. The algorithm is given the outcome variable and covariates describing characteristics of each individual. The characteristics the algorithm can choose from are age, sex, class, number of siblings/spouses aboard, number of parents/children aboard, passenger fare and port of embarkment.

There are in total 1 043 observations in the data set, with 80 % used to train the model, while the remaining 20 % are used to test the prediction accuracy of the model. In order to avoid that the algorithm overfit the tree to the training sample, the depth of the tree is restricted to be at most three levels deep. Figure 8 shows the tree grown by the algorithm.

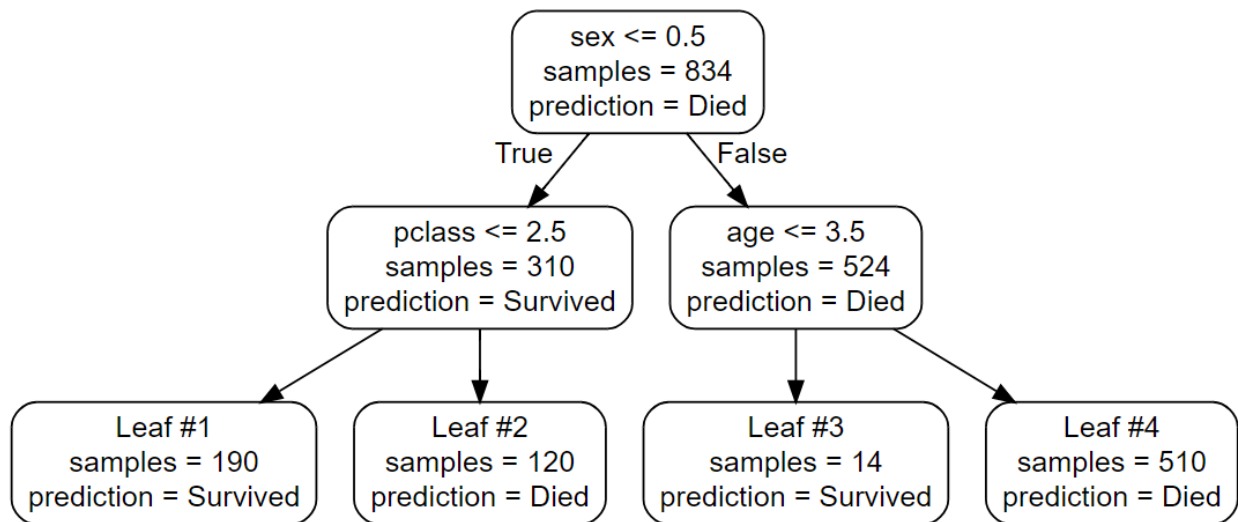


Figure 8: Regression tree for who survived the Titanic.

The root node contains all of the 834 observations in the training sample, and the

²³Data can be found at: <http://biostat.mc.vanderbilt.edu/wiki/pub/Main/DataSets/titanic3.xls>.

²⁴The term regression tree usually refers to trees with a continuous outcome variable. When the outcome variable is binary, the tree is usually referred to a decision tree

algorithm has chosen a tree that is three levels deep. The first variable the tree splits on is the sex of the passengers. The covariate sex takes the value 1 for males and 0 for females. This means that the tree starts by allocating all of the males on the Titanic to the right region, while all the female passengers are allocated to the left region. In the second level of the tree, the algorithm chooses different covariates to split on, depending on whether the passenger is male or female.

For the males the tree splits on the age, where boys younger than four years old are allocated to the left sub-region, while everyone four years or older are allocated to the right sub-region. For the female passengers, the tree instead splits on the class variable, which takes the values 1, 2 and 3. The tree allocates the female passengers that travelled third class to the right sub-region, while the female passengers that travelled first or second-class to the left sub-region. All of the observations in the training sample have now reached their leaves. The tree predicts that young boys under the age four and females that traveled first or second-class will survive; the rest will die.

In this toy example, the tree discovers a heterogeneity in the probability of surviving the Titanic that we would only have observed in regression analysis if we included interaction terms between sex and class, and sex and age. The tree discovered this dynamic between the covariates without us having to specify any functional form. However, there is always the risk that the tree has overfitted the data. This can be ascertained by comparing the prediction accuracy of the model in both the training and test sample. While the model achieved a prediction accuracy of 79.5 % in the training sample, the prediction accuracy remained 78.5 % in the test sample. The small difference in these scores implies that the tree is in fact capturing a dynamic between the covariates and the probability of survival common to all passengers aboard the Titanic.

D Predictions: State-level heterogeneity

Figure 9 shows the state-average annual number of temperature-induced fatalities in the period 2080-2099, compared to the period 1960-2004, predicted by the baseline OLS model and the regression tree.

The baseline OLS model shows a clear pattern in the effect of increased temperatures on the number of temperature-induced fatalities in the future across US states. According to the model, many states will experience a decrease in the average annual number of temperature-induced fatalities; this is especially the case for states in the Northern part of the country, while states in the Southern part will experience an increase in temperature-induced mortality. In fact, the large increase in the number of temperature-induced fatalities given by the baseline OLS model in table 3 is driven mainly by a handful of states: 13 331 of the predicted 17 541 additional temperature-induced fatalities each year comes from only four states, namely California, Arizona, Texas and Florida.

According to the regression tree, most states will in the period 2080-2099 experience a decrease in the average annual number of temperature-induced fatalities. For most states, this decrease will be small, while for states in the Northeast, the decrease will be much larger, with e.g. New York experiencing 1095 fewer temperature-induced fatalities each year. As opposed to the baseline OLS model, the regression tree finds that most states in the Southern part of the country will be unaffected by the increase in temperatures, while the adverse effects of increased temperatures is felt mainly by California and Florida.




Figure 9: Predicted change in average annual number of temperature-induced fatalities predicted by (a) the baseline OLS model, and (b) the regression tree, between the periods 1960-2004 and 2080-2099.

Centre for Applied Macroeconomics and Commodity Prices (CAMP)

will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

<http://www.bi.no/camp>