UNIVERSITY OF CAMBRIDGE | Energy Policy Research Group

Causal Tree Estimation of Heterogeneous Household Response to Time-Of-Use Electricity Pricing Schemes

EPRG Working Paper   1906

Cambridge Working Paper in Economics   1865

Eoghan O'Neill and Melvyn Weeks

Both regulators and energy companies have recognised the need to understand the distributional implications of energy policies. A report produced by the Centre for Sustainable Energy for Ofgem (CSE2012, CSE2014) indicates how policy makers are interested in "going beyond the mean" and identifying groups of customers most strongly affected by a policy change.

This paper considers the example of the impact of Time of Use (TOU) tariffs on household electricity demand. Consumers in different socioeconomic groups may react in different ways to the introduction of (TOU) tariffs. Similarly, customers with distinct historical intra-day load profiles respond differently to the introduction of tariffs that charge different prices for electricity at different times of the day.

In this paper, we apply recently developed Machine Learning (ML) methods to determine how household demand response to Time of Use (TOU) electricity pricing schemes varies with survey variables and past consumption data. Heterogeneous response is described by estimates of Conditional Average Treatment Effects, which are the expected differences between treated and control households for subsets of the population defined by covariates. We use causal trees (Athey & Imbens 2016) to search across potential conditioning variables for aspects of heterogeneity that are possibly difficult to hypothesize a priori. We then obtain household-specific estimates from a causal forest (Wager & Athey 2017).

Household-specific estimates produced by a causal forest exhibit reasonable associations with covariates. For example, households that are younger, more educated, and that consume more electricity, are estimated to respond more to a new pricing scheme. In addition, variable importance measures suggest that some aspects of past consumption information may be more useful than survey information in producing these estimates. Furthermore, household response estimates exhibit some bimodality when past consumption information is available, in contrast to the distribution of estimates produced by using only survey covariates.

Contact	mw217@cam.ac.uk
Publication	January 2019