CARMA 2026

8th Int. Conf. on Advanced Research Methods and Analytics
1-3 July 2026, Valencia, Spain
DOI: https://doi.org/10.4995/CARMA2026.2026.21977

Clustered Dose–Response Functions for Continuous Treatments: the case of EU funds

Augusto Cerqua[1], Roberta Di Stefano[2], Raffaele Mattera[3]

[1]Department of Social Sciences and Economics, Sapienza University of Rome, Italy [2]Department of Economics, University of Molise, Italy, [3]Department of Mathematics and Physics, University of Campania "Luigi Vanvitelli", Italy.

How to cite: Cerqua, A.; Di Stefano, R.; Mattera, R. (2026). Clustered Dose–Response Functions for Continuous Treatments: the case of EU funds. In: 8th International Conference on Advanced Research Methods and Analytics (CARMA 2026). Valencia, 1-3 July 2026. https://doi.org/10.4995/CARMA2026.2026.21977

Abstract

When treatments are continuously distributed, non-randomly assigned, and generate heterogeneous responses across units, identifying causal dose–response relationships becomes particularly challenging. These difficulties are common in many policy settings, including EU Cohesion Policy, where regions receive substantially different amounts of funding. Existing approaches typically estimate a single population-level dose–response function, potentially masking important heterogeneity in the returns to funding. We propose a Clustered Dose–Response Function (Cl-DRF), a novel estimator that combines generalized propensity score adjustment with clustering techniques to identify groups of units characterized by distinct treatment–outcome relationships. Applying the method to EU-15 NUTS-2 regions over 1994–2015, we uncover substantial heterogeneity: higher funding increases growth in some regions without diminishing returns, while others appear unable to fully benefit.

Keywords: *Continuous treatment; Heterogeneous effects; Generalized propensity score; European Cohesion Funds.*

1. Introduction

European Union Cohesion Policy represents one of the largest place-based policy interventions in the world. Over the last decades, a substantial share of the EU budget has been devoted to structural and cohesion funds aimed at fostering economic convergence and reducing regional disparities across European regions. Because of its scale and policy relevance, Cohesion Policy has been at the center of a long-standing empirical debate on its effectiveness. The existing literature has produced mixed results, with some studies documenting positive growth effects while others find limited or heterogeneous impacts (see, for example, Bandeira Morais, 2025

Article in Press

This work is licensed under a Creative Commons License CC BY-NC-SA 4.0
Editorial Universitat Politècnica de València

for a recent overview). Much of this literature evaluates the effects of Cohesion Policy using binary treatment frameworks, focusing on the average treatment effect of being eligible for or receiving EU transfers. However, this perspective abstracts from an important feature of the policy: the intensity of funding varies substantially across regions. Regions receive different amounts of EU transfers relative to their population or economic size, implying that Cohesion Policy operates as a continuous treatment rather than a binary intervention. Understanding how economic outcomes respond to different levels of funding is therefore crucial for evaluating the effectiveness of the policy. A smaller strand of the literature addresses this issue by estimating the average dose–response function (ADRF), which describes how the expected outcome varies with treatment intensity. For example, Becker et al. (2012) and Cerqua and Pellegrini (2018) study how regional economic growth varies with the intensity of EU transfers and document patterns consistent with diminishing returns to higher funding levels. While these contributions provide important insights into the marginal effects of funding intensity, they typically estimate a single population-level dose–response function. Yet, the allocation process underlying Cohesion Policy suggests that regions receiving large amounts of funds may differ systematically from those receiving lower transfers in ways that may affect how they respond to funding. Differences in economic structure, institutional quality, and absorptive capacity may generate substantial heterogeneity in the returns to EU funds. In such settings, estimating a single population-level dose–response function may mask important variation in treatment responses across regions.

In this paper, rather than imposing a single population-level relationship between funding intensity and economic growth, we allow the dose–response function to vary across groups of regions. Importantly, these groups are not specified ex ante on the basis of observable characteristics. Instead, they are identified directly from the data as sets of regions that share similar treatment–outcome relationships. To do so, we implement a Clustered Dose–Response Function (Cl-DRF) approach. This approach combines generalized propensity score adjustment with clustering techniques. More specifically, we apply this approach to quantify the impact of EU Cohesion Policy funding on economic growth in EU-15 NUTS-2 regions over the period 1994–2015.

2. The Cl-DRF estimator

Observational studies in the social sciences and economics often involve treatments that vary continuously in intensity and are assigned through non-random allocation mechanisms. Estimating their causal effects therefore represents a central challenge in causal inference, as it requires methods capable of addressing both confounding and the continuous nature of treatment. A growing literature tackles this problem using the GPS framework introduced by Hirano and Imbens (2004), which enables the estimation of the ADRF, describing how the expected outcome varies with treatment intensity. Identification in this framework relies on the

Article in Press

assumption of weak unconfoundedness—meaning that, conditional on observed covariates, the level of treatment received is independent of potential outcomes—and on positivity, which requires that units with similar characteristics have a positive probability of receiving different treatment intensities. Under these conditions, the GPS can be used to adjust for differences in observed covariates across treatment levels. Subsequent work has expanded this framework through alternative weighting strategies, matching approaches, and nonparametric estimators (Imai and van Dyk, 2004; Fong et al., 2018; Wu et al., 2024; Huling et al., 2023). Most existing approaches implicitly assume that treatment responses are homogeneous across units once observable characteristics are controlled for. In many policy contexts, however, this assumption may be restrictive. In such settings, estimating a single population-level ADRF over a wide range of treatment intensities may mask substantial heterogeneity in treatment responses. These concerns relate to a broader literature on heterogeneous treatment effects that has developed flexible data-driven approaches to detect and estimate heterogeneous responses, including causal trees and generalized random forests (Athey and Imbens, 2016; Wager and Athey, 2018; Athey, Tibshirani and Wager, 2019). However, this literature has primarily focused on binary or discrete treatments, while the estimation of heterogeneous dose–response relationships under continuous treatments remains relatively less explored.

We propose an alternative strategy based on local identification within groups of units characterized by similar treatment–outcome relationships, yielding empirically credible dose–response relationships. More precisely, we target a cluster-specific dose–response function, which constitutes the primary causal estimand of the proposed framework. Formally, for cluster c, the estimand is:

$$\mu_c(t) = E[Y(t) \mid C = c] \qquad (1)$$

which represents the expected outcome under treatment level t for units belonging to cluster c.

To estimate cluster-specific dose–response functions when cluster membership is unknown, we employ an iterative procedure that combines clustering techniques with GPS adjustment. The estimator extends the GPS framework of Hirano and Imbens (2004) to a clustered setting, drawing on ideas from grouped and clustered regression models (Demidenko, 2018; Sugasawa, 2021). This procedure simultaneously estimates cluster membership and cluster-specific DRFs, and will be referred to as the Cl-DRF estimator. In this framework, units assigned to the same cluster share common parameters governing the treatment–outcome relationship, while clusters are recovered through the iterative algorithm.

The estimation algorithm alternates between three main steps. We estimate cluster-specific dose–response functions using the proposed Cl-DRF estimator, which jointly recovers cluster membership and treatment–outcome relationships through an iterative procedure. First, we estimate a cluster-specific treatment assignment model to obtain the generalized propensity score (GPS), summarizing the conditional density of treatment given observed covariates.

Article in Press

Second, conditional on cluster membership, the outcome is modeled through a regression of the outcome on treatment intensity and the cluster-specific GPS. Third, cluster membership is updated by assigning each unit to the cluster that minimizes the squared residuals of the outcome model. These steps are iterated until convergence, sequentially updating the treatment model, the outcome regression, and the cluster assignments. Further methodological details are provided in Cerqua et al., 2026. Moreover, the Cl-DRF estimator builds on the k-means regression methodology and therefore requires specifying the number of clusters C in advance. To select C, we adopt a penalized empirical risk criterion inspired by information criteria commonly used in grouped regression models (Sugasawa, 2021). The resulting estimator identifies groups of units characterized by similar treatment–outcome relationships and estimates a cluster-specific dose–response function within each group.

3. Data and empirical setting

We focus on the relationship between EU Cohesion Policy funds and regional economic growth over the period 1994–2015. The main data source is the Annual Regional Database of the European Commission's Directorate-General for Regional and Urban Policy (ARDECO), which provides comprehensive socio-economic and demographic information for regions across the European Union.

The analysis is conducted at the NUTS-2 regional level for the EU-15. The outcome variable is the compound annual growth rate of GDP per capita, measured in Purchasing Power Standards (PPS), over the period 1994–2015. To account for differences in initial regional characteristics that may affect both economic growth and the allocation of funds, we include a set of pre-treatment covariates measured in 1994. These include GDP per capita in PPS, population density, the share of gross value added in the primary sector, average yearly hours worked per employee, average yearly compensation per employee, and the workplace employment rate.

The treatment variable measures the intensity of EU Cohesion Policy transfers received by each region. Specifically, we use yearly estimates of EU structural funds allocated per capita as the treatment intensity variable. Although all regions receive some level of funding, the magnitude of transfers varies substantially across regions, providing the variation necessary to estimate the dose–response relationship between EU funds and regional economic growth.

4. Results

The Cl-DRF algorithm identifies three clusters. Cluster 1 includes regions from Ireland, Southern Germany, and Southern England, while Cluster 2 predominantly consists of regions in Spain, Northern Germany, the Netherlands, Sweden, and Finland. Cluster 3, instead, is composed of many French regions as well as nearly all regions in Italy and Greece. Figure 1

Article in Press

reports the cluster-specific dose–response functions. Cluster 1 exhibits a steep positive relationship between EU funds and economic growth, although the estimates are less precise due to the relatively small number of observations in this cluster. Cluster 2 shows a mildly positive relationship between EU funds and growth, which gradually flattens and eventually turns slightly negative for the most highly subsidized regions. By contrast, regions in Cluster 3 display a slightly negative relationship between funding intensity and economic growth. However, the composition of this cluster provides additional context. Many of the regions in Cluster 3 belong to countries that were among the most severely affected by the Great Recession. This may help explain the counterintuitive finding that higher levels of EU funding are associated with lower economic growth in these regions. Another important feature of the results concerns the support of the treatment distribution. The common support for Cluster 1 is substantially smaller and does not include the most highly subsidized regions. For some of the wealthiest EU regions in this cluster, our analysis therefore finds no evidence supporting the hypothesis that a maximum funding threshold exists beyond which additional resources fail to stimulate further economic growth. These findings have important policy implications. They suggest that policymakers could reduce deadweight losses by tailoring subsidy schemes to the heterogeneous marginal returns identified across clusters, rather than applying uniform funding caps across regions. At the same time, recognizing that some regions may not fully benefit from additional funding highlights the importance of investing in the development of local absorptive capacity.

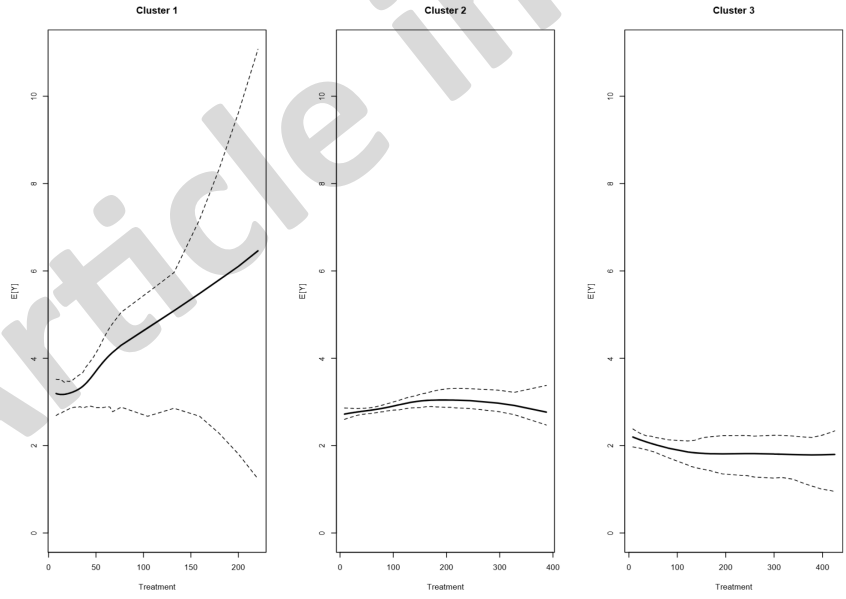


Figure 1. Clustered DRFs with the proposed approach. 95% confidence intervals are computed using 200 within-cluster bootstrap replications, including GPS estimation at each iteration.

Article in Press

References

Athey, S., & Imbens, G. W. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences, 113*(27), 7353–7360.

Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *The Annals of Statistics, 47*(2), 1148–1178.

Bandeira Morais, M. (2025). The Economics of Place-Based Policies in the EU: A Review of the Literature and Economic Impact. *Luxembourg: Publications Office of the European Union.*

Becker, S. O., Egger, P. H., & von Ehrlich, M. (2012). Too much of a good thing? On the growth effects of the EU's regional policy. *European Economic Review, 56*(4), 648–668.

Cerqua, A., Di Stefano, R., & Mattera, R. (2026). *Heterogeneous responses to continuous treatments: A cluster-based causal framework* (arXiv:2409.08773).

Cerqua, A., & Pellegrini, G. (2018). Are we spending too much to grow? The case of structural funds. *Journal of Regional Science, 58*(3), 535–563.

Fong, C., Hazlett, C., & Imai, K. (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics, 12*(1), 156–177.

Hirano, K., & Imbens, G. W. (2004). The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives* 226164, 73–84.

Huling, J. D., Greifer, N. & Chen, G. (2024), 'Independence weights for causal inference with continuous treatments', *Journal of the American Statistical Association* **119**(546), 1657–1670.

Imai, K., & Van Dyk, D. A. (2004). Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association, 99*(467), 854–866.

Sugasawa, S. (2021), 'Grouped heterogeneous mixture modeling for clustered data', *Journal of the American Statistical Association* **116**(534), 999–1010

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.

Wu, X., Mealli, F., Kioumourtzoglou, M.-A., Dominici, F. & Braun, D. (2024), 'Matching on generalized propensity scores with continuous exposures', *Journal of the American Statistical Association* **119**(545), 757–772.

Article in Press