Sapienza University of Rome

Department of Economics and Law

European PhD in Socioeconomic and Statistical Studies

XXXVIII cycle

Think Outside the Black Box!

Leveraging Machine Learning Predictions for Policy Insight

Candidate

Zanoni Angela

Thesis Advisor

Prof. Ventura Marco

Co-advisors

Prof. Paglialunga Elena

Prof. Resce Giuliano

Domestic animals expect food when they see the person who usually feeds them. [...] The man who has fed the chicken every day throughout its life at last wrings its neck instead, showing that more refined views as to the uniformity of nature would have been useful to the chicken.

— Bertrand Russell, **The Problems of Philosophy** (1912)

Contents

Introduction

When we think of good policy-making, we likely envision action that is both conscious of the status quo and sensible in simultaneously targeting multiple dimensions, considering numerous actors, and accounting for complex underlying trends and interrelations. A good policymaker is someone who has studied, learned, and formulated an informed hypothesis about the likely impacts of their actions. This being the case, it is easy to see why a "super learner" appeals in the context of many socioeconomic challenges.

Against the complexity of the real world, machine learning (ML) offers promising capabilities. By framing policy problems as prediction problems, ML algorithms mimic the iterative nature of human learning: formulating hypotheses, testing them against data, and refining understanding through repeated exposure to patterns. If properly trained, the result of an ML analysis reflects an accurate and nuanced understanding of the phenomenon—the cornerstone of all policy action.

However, using ML to inform policy encounters two fundamental challenges.

First, it involves a critical trade-off: predictive accuracy comes at the expense of clear causal claims. The reason this is problematic is, in essence, a manifestation of a well-known problem in epistemology—the problem of induction. Back in 1912, philosopher Bertrand Russell explained the conundrum through his parable of the inductivist chicken [1], who learns from repeated experience that the farmer brings food each morning, only to have this pattern catastrophically fail on the day it is

[1] Russell, B. (1912). The Problems of Philosophy. London: Williams and Norgate; New York: Henry Holt and Company

slaughtered. ML predictions, like Russell's chicken, excel at identifying regularities in data, but formulating policy suggestions based on past patterns requires assuming that observed patterns will hold under new conditions or policy interventions. The challenge, therefore, becomes one of carefully framing policy problems into settings where prediction offers valuable assistance despite this inherent limitation.

Second, ML is a powerful black box - mathematically robust yet often opaque in its assumptions and implications. Just as human learning is prone to bias and misjudgment, so too are algorithms - yet algorithmic biases can be harder to detect and interrogate. How is new information processed? What implicit assumptions govern pattern recognition? These questions are not merely technical but fundamentally shape the reliability and applicability of ML outputs. ML algorithms always "say something", but their output requires careful interpretation. How should the results be translated for good policy advise?

This thesis explores three distinct applications of ML to inform policy, each demonstrating how predictive algorithms can generate actionable insights while remaining cognizant of their limitations.

Chapter 1 employs an Elastic Net algorithm to construct "what-if" scenarios in the context of European green transition policies. The green transition represents a paradigmatic multifaceted process, requiring the simultaneous pursuit of multiple socioeconomic outcomes, many of which are collinear, across contexts marked by substantial local heterogeneity. Machine learning approaches are well-suited to this challenge: collinearity poses no obstacle to regularized regression methods, and these algorithms explicitly exploit heterogeneity to improve predictive performance. We extend the standard Elastic Net framework to accommodate spatial effects, accounting for spillover and agglomeration dynamics that have long been documented in regional development literature. This approach enables us to anticipate potential winners and losers of the transition and to compare different policy interventions based on their predicted effects across diverse regional contexts.

Chapter 2 trains an ensemble of ML algorithms to predict regional unemployment rates across the EU. Through this exercise, we explore both the relevance and functional form of relationships between diverse predictors and regional unemployment rates. The results provide data-driven insights into numerous questions raised by economic theory, from the role of inflation dynamics to the impact of environmental policies on labor market outcomes, paving the way for future impact assessments. Most significantly, we demonstrate that tree-based ML algorithms can predict unemployment with considerable precision, offering policymakers valuable early warning capabilities in detecting emerging regional vulnerabilities.

Chapter 3 explores the potential of ML in causal inference by delving into the Matrix Completion for Causal Panels (MCP) framework, a novel approach to counterfactual imputation. This method leverages ML's predictive capacity to impute counterfactuals for treated units, enabling a previously inaccessible level of disaggregation in treatment effect estimation. Crucially, MCP implicitly assumes interdependence among units, effectively overcoming the traditional independent and identically distributed (i.i.d.) assumption. This aspect is embedded in the algorithm's functioning, yet it represents a paramount shift compared to more traditional methods, entailing important economic implications. We apply MCP to assess the labor market impacts of minimum wage policies, demonstrating its value as an analytical tool while also highlighting important limitations. Thanks to the imputation of counterfactual quantities, we can test previously untestable hypotheses, including parallel trends, providing with additional powerful insights.

Together, these three studies underscore a crucial lesson. ML is a powerful tool, but its responsible application requires clearly delineating its appropriate scope, accompanying quantitative results with contextual understanding, and maintaining transparency about the trade-offs inherent in choosing prediction over explanation. Data-driven policy must be informed policy—guided by algorithms but tempered by judgment, domain expertise, and an awareness of what remains beyond the reach of prediction.

These studies also point toward a promising direction for future research. After "thinking outside the black box", we ought to "look inside the black box". Delving deeper into the nature of ML algorithm assumptions—a matter that has often passed under the radar— is necessary to verify their compatibility with economic theory.

The stance of this thesis is that ML offers genuine promise for policy analysis, but realizing that promise requires intellectual integrity: acknowledging both capabilities and constraints, and maintaining a careful eye on how algorithmic functionings compare with the reality of socioeconomic processes.

Chapter 1

More than just transition Uncovering heterogeneous socio-economic outcomes of climate policies

This essay is co-authored with Valeria Costantini (Roma Tre University) and Elena Paglialunga (Roma Tre University).

Abstract

Achieving sustainable human systems require transformative shifts to new socio-technical systems, boosted by far-reaching policy action. The green transition is a passage that entails complex and heterogeneous consequences across territories, bringing to potential negative distributional effects. Moreover, the web of linkages on which policy actions work escapes traditional analytical frameworks. Taking the European Union as a case study of a just transition process, we propose an empirical setting based on a machine learning technique (ML) with three main objectives. First, we develop a model that better captures the complexity of a sustainable and just transition process. Second, by comparing results with standard quantitative assessment models, we test whether the resulting insights derived from the implementation of an ML-based approach align with economic theory. Third, we explore the potential of a controlled ML framework in supporting context-specific policymaking for a just and sustainable transition.

Keywords: European Union; Inequality; Just transition; Machine learning; Regional disparities; Technological innovation.

J.E.L. classification: C31; D63; O33; Q55; R12.

1.1 Introduction

Achieving sustainable human systems requires more than incremental improvements; it demands transformative shifts to new socio-technical systems, *i.e.*, sustainability transitions - a concept first explored by Elzen (2004) and Grin et al. (2010). Such transitions are conceptualised as multidimensional, co-evolutive, and multi-actor processes (Kortetmäki et al., 2025; Köhler et al., 2019; Rosenbloom, 2025). The Just Transition framework (Newell and Mulvaney, 2013; Swilling, 2020), underscores the importance of pursuing these transformations in ways that ensure social acceptability and distributive justice, with particular attention to the spatial concentration of the benefits and burdens associated with the transition.

As a result, sustainability transitions research is called to answer "big picture" questions, including how transitions affect inequality and which strategies can help mitigate their adverse social impacts (Sovacool et al., 2021). Yet, these questions introduce unprecedented layers of uncertainty and complexity that challenge conventional economic analysis. Stern and Stiglitz (2021) argue that traditional economic frameworks fail to represent faithfully a reality where multiple market failures are present. While valuable for outlining theoretical frameworks, traditional theoretical and empirical settings struggle to integrate heterogeneity and externalities effectively, limiting their normative capacity. These limitations stem from the inherently multidimensional nature of sustainability transitions, which involve complex interactions between economic, social, and environmental factors that traditional linear models cannot adequately capture. Effective policy-making can be conceptualized as a predictive task, whereby a nuanced understanding of macroeconomic trends and local dynamics informs actions toward achieving desired outcomes. However, developing a deep knowledge of all factors and potential impacts driving society is far from straightforward. The challenge lies in the complex web of interactions among multiple stakeholders and institutional layers, where policy actions generate cascading effects across interconnected systems. This complexity is particularly pronounced in sustainability policies, which must simultaneously balance apparently contrasting elements such as economic prosperity, social equity, and environmental preservation (Purvis et al., 2019). Each of these dimensions introduces significant policy challenges, and their interconnectedness compounds the difficulty (Edmondson et al., 2019). Traditional economic models, with their reliance on simplified assumptions and linear relationships, are ill-equipped to handle this multidimensional complexity, creating a compelling case for alternative analytical approaches.

In addressing this gap, machine learning (ML) techniques offer promising capabilities, as they exploit data abundance and variety to draw accurate predictions through an iterative learning process that mimics human decision-making (Kumar et al., 2024). This being the case, regional heterogeneity and, therefore, the spatial nature of the sustainability transition are not seen as an obstacle to overcome, but rather as the very source of data variation that makes the prediction possible.

The analysis is based on an original database collecting a diversified set of socio-economic variables at the regional (NUTS2) level for the European Union in the time frame 2000-2019. In this context, we develop an empirical analysis with three main objectives. First, we exploit ML techniques to construct a quantitative model designed to capture the complexity of the sustainable transition process. Second, by comparing results with standard quantitative assessment models, we test whether the resulting insights derived from the implementation of an ML-based empirical analysis align with economic theory. Finally, we assess whether a controlled ML framework could better support policy-making for a just and sustainable transition. This research direction is motivated by ML's demonstrated capacity for handling high-dimensional data and complex non-linear relationships - characteristics that align with the multifaceted nature of sustainability transitions. By leveraging ML's predictive capabilities, policymakers can better anticipate the socioeconomic outcomes of interventions, enabling more informed and effective policy design. Our approach explicitly harnesses the multidimensional analytical power of ML to provide a comprehensive framework for supporting sustainability transitions. The policy implications of our work are substantial. In line with the idea that all studies on sustainability transitions should admit that models are more relevant as "learning machines" than as "truth machines" (Turnheim and Nykvist, 2019), leveraging ML's capabilities allows not only to identify in advance the winners and losers of the sustainability transition, but also to learn which characteristics of a region make it more prone to experience losses. Insights like this are crucial to predict possible drawbacks at the regional level and to identify regions that may be subject to development lock-ins.

The rest of the paper is structured as follows. In Section 1.2, we frame the main contributions drawing the complexity of the concept of the just and sustainable transition. Section 1.3 provides a full description of the empirical framework adopted, including the database developed for the analysis (Section 1.3.1), the characteristics of the empirical method (Section 1.3.2), and the policy scenarios tested (Section 1.3.3). In Section 1.4 we discuss the main results. Section 1.5 concludes by discussing the policy implications of our findings.

1.2 Literature review

The link between sustainable transition and social justice is well-documented in the literature. The interplay between sustainable transitions and social justice involves reconciling environmental goals with equity, inclusivity, and fairness in societal transformations. Indeed, a sustainable transition process in its core identity is mainly aimed at achieving environmental goals such as a deep decarbonization, resource efficiency or biodiversity protection, boosting ecological priorities on top of whatever strategy for economic growth. Such objectives may or may not have significant socio-economic and distributive consequences, where the outcome of the process could be significantly at risk if the negative impacts substantially affect the acceptability of the policies at play for such purpose (Avelino et al., 2024). If the society is particularly keen to a social justice perspective, in this case the policy-making process should also ensure that the sustainable transition is equitable, addressing the needs and rights of all stakeholders, particularly marginalized groups (Grossmann et al., 2022).

There are several aspects generated by the interaction between environmental protection goals and the related distributional impacts, with single or mutual causal relations and consequently different perspectives of analysis. For instance, Oueslati et al. (2017) find a positive correlation between inequality and energy tax reforms, while Faiella and Lavecchia (2021) emphasize the need for compensation mechanisms to alleviate the disproportionate impact of green energy taxation on poorer households. In this respect, the fiscal policies dedicated to supporting the adoption and diffusion of green energy technologies seem to be particularly effective in correcting negative distributional impacts of climate mitigation regulation (Costantini et al., 2025). This is in line with the more general findings that new technologies might contribute to addressing societal challenges and fostering social welfare only if they are supported by well-designed public policies directed to ensure equity in individuals' well-being (Castellacci, 2023).

Additionally, a growing body of research is focusing on the non-linearity of the relation between environmental issues and income distribution, revealing that both channels are jointly working, with mutually self-reinforcing mechanisms at play.

For example, there are several side effects associated with the development of new and sustainable technological trajectories, as the impact on the labour market through the creation of the so-called green jobs might also bring to negative effects to low-skilled workers (Consoli et al., 2016; Napolitano et al., 2022). Environmental regulations can shift sector compositions, leading to (geographically concentrated) job creation or loss (Vona, 2021), with local labour markets dependent on energy and carbon-intensive industries being particularly vulnerable (OECD, 2012). Other geographical factors may exacerbate job losses in these sectors, affecting reconversion and migration costs, and altering income distribution. Castellanos and Heutel (2021) suggest that the local impact of green transition policies varies based on labour mobility, green innovation and technologies (Costantini et al., 2018), market concentration and firm size distribution (Zaman and Borsky, 2021), and the demand for green skills (Marin and Vona, 2019). Markkanen and Anger-K raavi (2019) frame the inequality-transition relationship within the broader context of climate change mitigation policies, emphasizing the importance of spatial analysis for understanding local labour market effects and regional disparities.

O'Sullivan et al. (2020) further highlight the interrelation between spatial justice and energy justice in the context of energy system transformations. Knowledge and innovative capacity emerge as a pivotal factor in the so-called twin (*i.e.*, green and digital) transitions (Damioli et al., 2024), but they are also unevenly distributed across regions (Balland et al., 2019).

From a global societal perspective, income distribution and climate change are strictly connected in a vicious cycle, since the most vulnerable groups across and within countries are more prone to suffer from climate-related disasters, and, at the same time, climatic challenges reinforce the weakness of the poorest (Cappelli et al., 2021; Paglialunga et al., 2022).

Overall, it emerges that, on the one hand, environmental stress is responsible for a significant increase in the divergence in income distribution. On the other side, those policies designed with the intent of reducing the sources of environmental impacts (*e.g.*, climate and energy policies or green innovation) might also produce further pressure on more vulnerable groups.

From a political perspective, these drawbacks partly originate from the fact that environmental policies are often decided at the national level, even though local implementation may encounter different obstacles. The European Union (EU) is an excellent case study of an area where environmental protection has been at the core of the development strategy for decades. Policy objectives are decided at the EU level of aggregation (supranational approach) while the practical decisions from a policy instrument perspective are taken at the national and regional levels. Moreover, the socioeconomic conditions at the sub-national level are diversified, and several EU policies have been designed and implemented to support the development process within a social equity framework. In this vein, the European Green Deal (EGD) outlined during Von der Leyen's mandate follows the Keynesian idea that a "big push" in terms of environmentally virtuous public investments could lead the EU to an economic renaissance. Nonetheless, the specific policy instruments embedded in the EGD are not always designed to respect a social justice criterion (Crespy and Munta, 2023), suggesting that some adjusting mechanisms are strongly required to reduce the negative impacts perceived by citizens.

According to Vona (2023), green deals such as the European one are designed to address three types of market failures: i) negative environmental externalities; ii) knowledge externalities that can be internalized in the green economy; iii) network externalities inherent to green technologies. This brings about the compulsory requirement of a multidimensional approach that should integrate fairness in resource allocation, participatory governance, cultural recognition, and inter-generational responsibility. Only by addressing these aspects is it possible to ensure that the transition process is not only environmentally sustainable, but also socially inclusive, creating a more just and resilient society. Accordingly, to sustain the policy-making process, a complexity perspective must be adopted (van Bommel and Höffken, 2023). Standard methods of analysis based on linear models might be able to capture only a portion of the web of causal relations and should therefore be at least complemented by more flexible analytical solutions such as the ML techniques. This avenue is gaining momentum (Biancalani et al., 2024; Donti and Kolter, 2021; Yao et al., 2023).

1.3 Empirical framework

This section describes the dataset and the methodology used to investigate the complex dynamics of sustainable transitions in the European Union. The analysis aims at bridging the gap between economic and policy perspectives by leveraging ML techniques to explore the intricate relationships between socioeconomic, environmental, and innovation factors. By integrating regional-level (NUTS 2) data from 2000 to 2019, we seek to capture the interactions shaping regional disparities and their implications for sustainability transitions.

The dataset used in the analysis includes a comprehensive array of indicators across socioeconomic, environmental, and innovation dimensions. This allows us to analyse regional dynamics across the three pillars of sustainability (*i.e.*, economic prosperity, environmental protection, and social equity) in a context of heterogeneous regional challenges and opportunities. The data is leveraged to study regional disparities and simulate outcomes of alternative policy interventions under different scenarios.

The empirical ML framework allows us to develop a quantitative model that captures the complexity and interdependencies inherent in sustainable transition processes. This approach forms the basis for a robust analysis of policy trade-offs and synergies, providing valuable insights to support equitable and effective sustainability transitions.

1.3.1 Data

Our dataset incorporates regional data for EU NUTS 2 regions, supplemented with national data on policy settings, over a 20-year time frame (2000 to 2019). The dataset includes 38 variables, mirroring the three-pillar structure in Purvis et al. (2019) (see Table 1.1 for a detail of variables' sources and codification in our dataset).[1]

Our scenarios will be evaluated on four socioeconomic indicators. First, the regional unemployment rate, expressed as a percentage of the labour force. Because employment was at the core of the Cohesion policy planning 2014-2020 (thematic objective 8, see European Commission (2023)), this dimension has strategic policy implications. Unemployment rates have been widely used to measure the societal impacts of socio-technical transitions (Castellanos and Heutel, 2021; Vona, 2021). Secondly, we look at Gross Domestic Product (GDP) in NUTS 2 regions, measured in constant million euros, in line with studies that see economic growth as a substantial outcome to determine disparities in green transition deployment, and development traps along climate policies (Diemer et al., 2022; Iammarino et al., 2019). Third, and complementarily, we consider disposable income in NUTS 2 regions, measured in million purchasing power standards[2]. Disposable income is more directly linked to households' consumption behaviour and it is therefore used to explore uneven household effects of climate policies (Rodríguez-Pose and Bartalucci, 2023). Finally, we explore the effect on the gender employment gap by NUTS 2 regions (as a percentage). This serves as a proxy for intersectionality and inclusivity in socio-economic systems and aligns with extensive literature that uses gender as a proxy for systemic inclusivity (see Johnson et al. (2020) for a review). Gender representation in climate governance is crucial to ensure equitable transitions (Kro-

[1] To preserve as much information as possible, breaks in time series are interpolated. Observations for region i at time t that still present missing values for certain variables are dropped.

[2] Both GDP and disposable income are considered in per capita terms, *i.e.*, we divide the absolute value by the yearly regional population

nsell, 2013). Evidence suggests that women are significantly underrepresented in both carbon-intensive sectors and in high-technology or (green) innovation-driven sectors (Maldonado et al., 2024). This structural segmentation of the labour market helps explain why shocks to innovation and/or sectoral restructuring may not see corresponding improvements in gender equality.

The features considered for the prediction task are inspired by the literature review in Section 1.2. Patenting activity is a common proxy for the innovative capacity of a territory (Barbieri et al., 2023). Similarly, eco-patents capture eco-innovative propensity, often linked to successful green transitions in the literature (Costantini et al., 2017, 2018; Napolitano et al., 2022).

The OECD REGPAT (Regional Patent) Database collects information on patent applications at the European Patent Office (EPO) as well as within the Patent Co-operation Treaty (PCT). Only the first case is considered in our analysis. We use the Spring 2023 version of the dataset and focus on EPO applications. From the database, we extract information on the technological class of patents by cooperative patent classification codes (CPC). We then assign patent applications to each NUTS 2 region, counting how many have priority year between 2000 and 2019 by an inventor located in the region. This step is then repeated after filtering out patent applications with a CPC code equal to Y02, which indicates inventions dedicated to climate change adaptation and mitigation. In this way, we compute an eco-innovative patent count. Figure 1.1 presents an overview of the results of our counting[3].

However, patenting activity alone may not capture the full extent of the innovative landscape in regions. Following the understanding that regional support for R&D plays a crucial role in preparedness against climate change and complex societal changes (Cappellano et al., 2024), we include variables of knowledge capital building as collected in the Regional Innovation Scoreboard (RIS). The RIS database focuses on research and innovation-related metrics. Out of those available, we select three variables: employment in technology and knowledge-intensive sectors, partic-

[3]Due to mismatches between OECD regional classifications and Eurostat, and to NUTS 2 updates over the time frame considered, some recoding was necessary. Find details on the process in Appendix 1.A.1.

Figure 1.1: Number of patent applications in the EU, by year

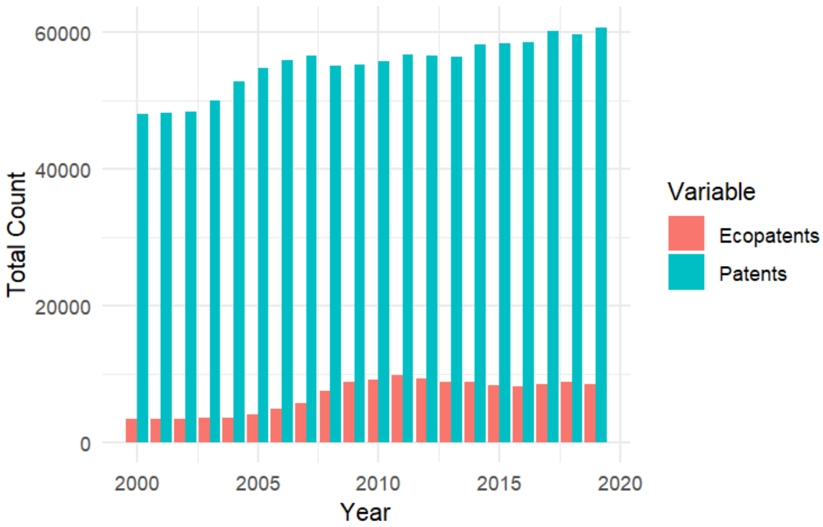

ipation rate in education and training, and tertiary educational attainment in the age group 25-64. The latter two come as percentages, while the former is expressed in thousands of persons, therefore we transform it into relative terms by dividing the absolute value by the yearly regional population.

Moreover, we consider relative employment in sectors vulnerable to climate transition (hereafter "vulnerable sectors"). Rodríguez-Pose and Bartalucci (2023); Vandeplas et al. (2022); Vona (2021) identify reliance on labour in "brown" sectors as a key determinant of regional vulnerability to the distributive effects of climate policies. To incorporate this research strand, we include employment data by economic activity for NUTS 2 regions across NACE sectors A through F, measured in thousands of persons [4].

We complement information on sectoral employment and innovation dynamics

[4] Our classification of vulnerable sectors represents a calibration between data availability and existing literature on sectors most likely to be significantly affected by sustainable transition. Following Rodríguez-Pose and Bartalucci (2023), brown sectors typically include agriculture, forestry, and fishing (A); mining and quarrying (B); specific manufacturing subsectors such as coke and refined petroleum products (C19), chemicals (C20), non-metallic mineral products (C23), and basic metals (C24); utilities including electricity, water supply, and waste management (D, E); and transport (H). However, we make several adjustments to this framework. First, we exclude transport sector H due to its aggregation with sectors G and I in Eurostat data. Second, unlike Rodríguez-Pose and Bartalucci (2023), we retain the entirety of NACE sector C (manufacturing). Third, we include construction (NACE F), given evidence from Maldonado et al. (2024) that this sector is particularly prone to transformation due to high future demand for green jobs and in light of recent discussions to include the construction sector within the Emission Trading Scheme framework.

by including the median age of the population in the region, to control for the possible relevance of demographic dynamics.

Greenhouse gas emissions are a common indicator of the environmental footprint of human activities. Even if they are not exhaustive of the overall impact, these are widely used to assess the extent and success of green transitions and, more broadly, sustainable development (Bianchini et al., 2023). The EDGAR (Emissions Database for Global Atmospheric Research) collects emission data from different sources and harmonizes this information to generate detailed estimates of emissions for a wide range of pollutants. The measurements available at the NUTS 2 level, which we include in our dataset, are carbon dioxide (CO2), methane (CH4), nitrous oxide (N2O), and fluorinated greenhouse gases (F-gases), expressed in Kt of CO2eq. For these variables, we consider a per capita value by dividing each variable by the yearly value of the regional population.

Lastly, the OECD environmental policy stringency (EPS) indicators offer insights into countries' regulatory landscape surrounding environmental protection and sustainability. These indicators gauge a country's efforts to mitigate pollution and foster sustainable development practices. We assign to every region the corresponding national score. This composite index tracks climate change and air pollution mitigation policies. Increased stringency, indicative of favourable institutional and political landscapes, is a strong predictor of environmental performance (Frohm et al., 2013) and eco-innovation (Ghisetti and Pontoni, 2015; Xie et al., 2023). We assign to every region the corresponding national score.

1.3.2 Elastic net model for spatial regression

ML is a bundle of techniques to detect patterns hidden in complex data structures, using an iterative process called cross-validation. Cross-validation is a statistical method used to evaluate the performance of a model by testing it on data that were not used during the training phase. The process involves splitting the dataset into multiple subsets, or "folds"; the model is trained on some of these folds (the training set) and validated on the remaining fold(s) (the validation set). This process is repeated several times, with each fold taking a turn as the validation set. The

Table 1.1: List of Variables

Source	Name	Description	Unit of Measure
Cohesion Open Data Platform	coh	Cohesion fund recipient	1= yes, 0= no
EDGAR – Emissions Database for Global Atmospheric Research	ch4	Methane emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	co2	Carbon dioxide emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	fgas	Fluorinated gases emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	n2o	Nitrous oxide emissions	Kt CO2eq
Environmental Policy Stringency (EPS) Index Dataset	c_co2ets	CO2 Trading Scheme	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_co2tax	Carbon dioxide (CO2) Tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_diesel	Diesel tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_enepol	Adoption of solar and wind energy support measures	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_lowrd	Low-carbon R&D expenditures	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_market	Market-based instruments	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_nonmark	Non-market based instruments	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_noxlim	Emission limit value NOx	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_noxtax	NOx taxation	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_pmlim	Emission limit value PM	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_renets	Renewable Energy Trading Scheme	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_solar	Solar Energy support (Auctions & FITs)	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sox	Emission limit value SOx	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sulflim	Emission limit value sulphur	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sulftax	Sulphur Oxides (SOx) Tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_tech	Technology support policies	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_wind	Wind Energy support (Auctions & FITs)	Integer
Eurostat	disp	Disposable income	Million purchasing power standards
Eurostat	egap	Gender employment gap	Percentage
Eurostat	emp	Employment in NUTS2 regions	Thousands of persons
Eurostat	gdp	Gross domestic product (GDP)	Million euros
Eurostat	htec	Employment in technology and knowledge-intensive sectors	Thousands of persons
Eurostat	lllearn	Participation rate in education and training	Percentage
Eurostat	median	Median age of the population	Years
Eurostat	nacea	Employment in Agriculture, forestry and fishing	Thousands of persons
Eurostat	nacebe	Employment in Industry (except construction)	Thousands of persons
Eurostat	nacef	Employment in Construction	Thousands of persons
Eurostat	pop	Population by NUTS 2 regions	Thousands of persons
Eurostat	ted	Tertiary educational attainment, age group 25-64	Percentage
Eurostat	unemp	Unemployment rate	Percentage
REGPAT (2024)	ecopat_count	Number of eco-innovative patent applications (EPO)	Integer
REGPAT (2024)	pat_count	Number of patent applications (EPO)	Integer
Own elaboration on Eurostat data	vuln	Employment in vulnerable sectors (NACE A to F)	Thousands of persons

Note: All variables, except those from the EPS dataset (which are at the country level), are at the NUTS 2 level. Employment data by economic activity are classified according to NACE Rev. 2. For consistency, in the EN model we convert all absolute values to relative measures using two approaches: (1) dividing the variable by regional population, or (2) dividing by the total number of employed persons in the region, in the case of sectoral employment variables. Variables that undergo this transformation will be marked with the suffix "_*rel*".

performance of the model (or "fit") is evaluated on different measures of the error between predicted and observed values in the validation step.[5]

Unlike traditional statistical methods, which rely on predefined assumptions and equations describing relationships between variables, ML models derive predictions based on observed trends and interrelations, in an assumption-light way, letting the algorithm formulate and develop hypotheses as cross-validation moves forward.

In our case, we structure the prediction task as predicting different response variables, at the regional level, for the last year of our panel. The algorithm is trained on data from 2000 to 2018 and tested on out-of-sample data from 2019. We perform the training stage by regressing the $\mathrm{Y}_{i,t}$, on all the features in our dataset $\mathrm{features}_{i,t}$, for all t except the last. In addition, to account for agglomeration and geographical spillovers, we consider a spatial lag term Wy_t, computed in turn for each target variable Y. W is an adjacency matrix with value 1 when two regions are neighbours and 0 otherwise. The out-of-sample prediction is evaluated in the last year (2019).

$$\mathrm{Y}_{i,t} = f(features_{i,t}, Wy_t)) \tag{1.1}$$

$features$ is the vector of 36 predictors. After including Wy_t, we use 37 features to predict Y. $f(.)$ is the best-fitted Elastic Net (EN) model (Taylor and Tibshirani, 2015; Tibshirani, 1996; Zou and Hastie, 2005)). The EN solves the following optimization problem:

$$\hat{\beta} = \arg\min_{\beta} \left(\frac{1}{2n} \sum_{i=1}^{n} \left(y_i - \mathbf{x}_i^\top \beta \right)^2 + \lambda \left(\alpha \|\beta\|_1 + \frac{1-\alpha}{2} \|\beta\|_2^2 \right) \right) \tag{1.2}$$

Here:

[5] To evaluate the performance of our models, we rely on common error metrics widely used for regression tasks, like described in Cerulli (2023): Root Mean Squared Error (RMSE), Standardized RMSE, the Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). RMSE measures the square root of the average squared differences between the observed and predicted values. To allow comparability across different datasets, the standardized RMSE normalizes RMSE by the standard deviation of the observed errors. In the same spirit, MAE calculates the average absolute error, while the MAPE expresses the MAE as a percentage of the actual values. It is commonly used for its interpretability in cases where a continuous variable is predicted.

- y_i and $\mathbf{x}_i$ are the observed response and predictor vector for the i-th observation, respectively.
- $\|\beta\|_1 = \sum_{j=1}^{p} |\beta_j|$ is the LASSO penalty.
- $\|\beta\|_2^2 = \sum_{j=1}^{p} \beta_j^2$ is the Ridge penalty.
- $\lambda > 0$ controls the overall regularization strength, while $\alpha \in [0, 1]$ determines the balance between LASSO and Ridge penalties.

The L1 penalty ($\|\beta\|_1$) encourages sparsity by shrinking some coefficients exactly to zero, effectively performing feature selection. In contrast, the L2 penalty ($\|\beta\|_2^2$) distributes shrinkage more evenly across coefficients, favouring models that handle correlated predictors more effectively. By blending these approaches, Elastic Net inherits the advantages of both methods, balancing variable selection and stability in regression estimates. The Elastic Net is particularly valuable in high-dimensional settings, where traditional regression methods struggle due to multicollinearity or the curse of dimensionality. Our analysis, whose results are presented in Section 1.4, is carried out using the Caret package in R (Kuhn, 2008).

1.3.3 Scenarios

Once we retrieve the optimal model, we use it to predict three scenarios; one for each pillar of the sustainability transition. In practice, we use the estimated best-fit model to generate alternative predictions for the value of Y in 2019, and compare the observed values of selected independent variables with these newly generated values. The details of the scenarios are described below.

Scenario 1 - Enhanced innovative capacity

Grillitsch and Hansen (2019), argue that a pre-existing industrial capability is a condition for the development of local green economies. The innovative activity of firms mediates the relationship, as confirmed by Perruchas et al. (2020) and Barbieri et al. (2023). This scenario simulates uniform growth in patent applications across all regions, calibrated to match the average growth rate observed in regions not receiving cohesion funds during 2000-2018. The objective is to assess how a sustained enhancement in innovative capacity - known to be strongly associated with the green transition - might influence the socioeconomic performance of regions.

Scenario 2 - Emission cap

The Fit for 55 Package establishes a target of reducing greenhouse gas (GHG) emissions by 55% by 2030, relative to 1990 levels. To simulate this scenario, we use total GHG emissions data for NUTS2 regions in 1990 from the EDGAR database. We constrain 2019 emission levels to represent a 55% reduction relative to their 1990 baseline values, and simulate a scenario where GHG emissions decline at the average yearly rate necessary to achieve this target.

Scenario 3 - Labour market policies

Vandeplas et al. (2022) find a positive and significant correlation between employment in vulnerable sectors and unemployment, at the NUTS 2 regional level. At the same time, Rodríguez-Pose and Bartalucci (2023) argue that regions reliant on brown energy generation and consumption are experiencing the harshest consequences of green policies. The Clean Planet for All initiative in 2018 foresaw particular backlash on the energy and industrial sectors. This scenario simulates a uniform reduction in employment in vulnerable sectors, calibrated to match the average rate observed in regions not receiving cohesion funds during 2000-2018.

Scenario 4 - Policy mix

Rosenbloom (2017) and Turnheim and Nykvist (2019) point out that sustainability transitions consist of a broad range of economic, technological, and behavioural changes that are often overlapping. In sum, in a real-life context, the three scenarios above could as well co-exist. To address this possibility, we construct a "policy mix" scenario, where the previous three are combined.

The parameters for the simulations are carefully chosen following an in-depth analysis of historical trends to avoid the trap of extreme counterfactuals (King and Zeng (2006)). Additionally, the evolutions we simulate are within the historical trends exposed in Table 1.2.

For each of the three scenarios, we evaluate the effects of variations of certain features on the four response variables: unemployment, GDP, disposable income, and gender employment gap. Our choice of unemployment as the primary dependent variable in the prediction task aligns with a vast body of literature that sees unemployment as a crucial indicator of the acceptability of green policies at the local level (Vona, 2023). In this strand of literature, unemployment is regarded as a proxy for social cohesion and well-being, two of the three pillars of sustainable development (Purvis et al., 2019). Moreover, we also account for GDP, disposable income, and the gender employment gap as response variables, as they provide complementary insights into economic performance, household welfare, and labour market inclusivity, all of which are critical dimensions of sustainable development.

Variable	**Average yearly growth (%)**	**SD**
`Number of patent applications`	3.06	0.47
`Employment in vulnerable sectors`	-0.77	0.06

Table 1.2: Average yearly growth for features chosen for the scenarios. Note: The table reports the average yearly percentage variation and standard deviation observed in more advanced regions.

1.4 Results

1.4.1 Descriptive statistics and baseline econometric results

As a preliminary step in the empirical analysis, we explore the correlations between key variables with the aim of examining the patterns emerging both within and across groups of variables, namely, economic performance indicators, social factors, emissions, innovation, and environmental policy measures. The results, presented in the heatmap in Figure 1.2, reveal that variables capturing economic performance, such as GDP and disposable income, are positively correlated. On the other hand, and consistent with expectations, unemployment shows a negative correlation with both GDP and disposable income. Moreover, the gender employment gap is positively correlated with unemployment, suggesting that regions with higher unemployment also tend to exhibit greater disparities between men and women in labour market participation. Innovation variables, including general patents and eco-innovation, are positively correlated with GDP and disposable income, highlighting the role of technological development in driving economic performance.

Turning to indicators related to the labour market, the correlation map shows a positive relationship between employment in knowledge-intensive sectors and educational attainment. Conversely, unemployment is negatively correlated with educational attainment. Taken together, this evidence suggests the crucial role of education in facilitating access to the labour market and, in particular, to skilled labour opportunities.

In the environmental dimension, we observe strong positive correlations among variables representing GHG emissions (*e.g.*, CO2, methane, and nitrous oxide), indicating that regions with high emissions of one pollutant often exhibit elevated levels also of others. Similarly, Environmental Policy Stringency (EPS) measures (*e.g.*, CO2 taxes and CO2 trading schemes) display strong positive within-group correlations, reflecting the tendency of regions with high environmental commitments to adopt multiple complementary policies. Notably, the heatmap also shows negative correlations between EPS measures and emissions, indicating that regions in countries with more stringent environmental policies tend to have lower levels of emissions, which points to a degree of environmental policy effectiveness.

These observed patterns highlight the importance of considering interdependencies among variables, particularly in policy analysis. Measures targeting one variable or group of variables are likely to have indirect effects on others, reinforcing the need for an integrated approach to policy design that accounts for these complex relationships, both to enhance the effectiveness of the policies and to avoid unintended consequences across economic, social, and environmental dimensions.

This evidence is corroborated by the results from the econometric exercise reported in Table 1.3, although the interpretation is different, and some nuances emerge. More specifically, Table 1.3 provides the results obtained from Ordinary least Square (OLS) econometric estimation, where we test the determinants of our four response variable of interest, *i.e.*, (from column 1 to 4) unemployment, GDP, disposable income, and gender employment gap. In all models, we include regional

Figure 1.2: Heatplot of Correlation Between Variables in the Dataset

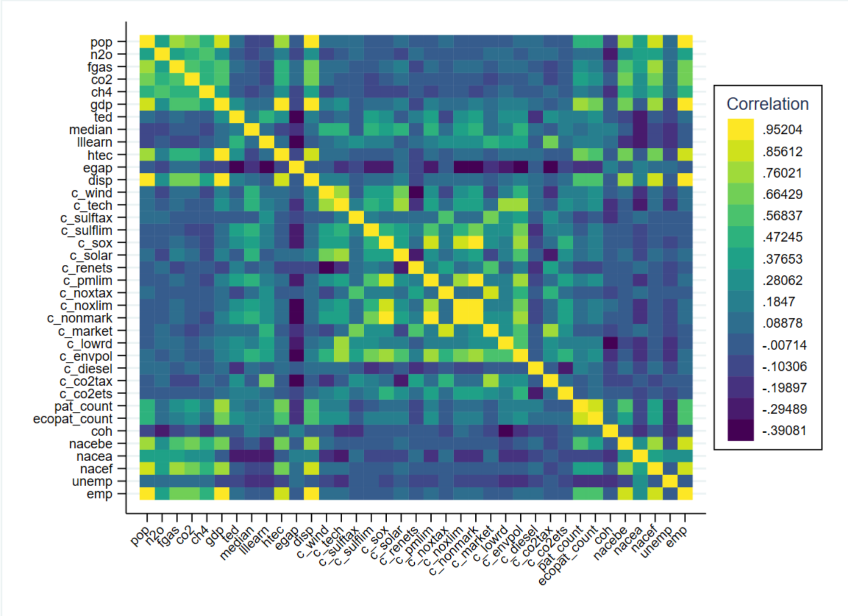

(NUTS 2) and year fixed effects to control for time-invariant characteristics of regions and common time trends. We also perform all estimates clustering the standard errors at the regional NUTS 2 level to control for heteroskedasticity and autocorrelation.

It should be noted that the purpose of this exercise is not to directly replicate the ML framework under a traditional linear model, but rather to explore the linear associations among features and variables of interest and to assess the extent to which these correlations would be problematic in a standard linear setting. For a closer counterpart to the ML exercise, which includes a spatial lag, we refer to the Durbin-type regression presented in Appendix 1.A.3.

The results in Column 1 show that GDP is negatively associated with unemployment levels, suggesting that relatively wealthier regions tend to have lower unemployment rates. CO2 emissions also have a negative and statistically significant association with unemployment, implying that, everything else being equal, higher emissions levels are linked to higher employment. Innovation, as measured by EPO patents, exhibits a positive and statistically significant coefficient, potentially pointing to a trade-off between innovation performances and job creation. While the coefficients for median age, education participation, and both market-based and non-market-based EPS measures are statistically significant, their magnitudes are negligible.

In Column 2, participation in education and training, and patent activity are both positively associated with GDP, highlighting the key role of innovation and human capital in driving economic growth. Conversely, unemployment, median age, and EPS measures show negative and statistically significant coefficients, suggesting a detrimental influence of these factors on economic performance.

For average disposable income (Column 3), most results align with those for GDP. Specifically, unemployment, median age, and non-market-based EPS measures have significant negative effects, while the variable *lllearn* has a positive and significant association, supporting the idea that higher involvement in education and training fosters income growth. However, unlike the previous results, EPO patents negatively affect disposable income, while CO2 emissions show a positive and significant coefficient. These findings align with the evidence in Column 1: higher emissions levels, which are associated with increased employment, also contribute to higher disposable income. On the other hand, the trade-off between innovation and job creation may also lead to a negative effect on average income, reflecting potential disparities in income distribution.

Finally, Column 4 focuses on the gender employment gap, identifying two significant predictors: the unemployment rate and employment levels in technology and knowledge-intensive sectors. Both coefficients are positive and, while the magnitude of the latter is nearly zero, the former indicates that higher unemployment exacerbates gender disparities in the labour market.

This evidence highlights the complex interplay between economic, social, and environmental factors, revealing trade-offs and synergies that policymakers should take into account to jointly foster inclusive growth, advance technological systems, environmental sustainability and equitable development. Notably, the complexity of these relationships and the mutual influences across the dimensions of interest further justify our reliance on ML techniques.

Table 1.3: OLS regression model with regional and year fixed effects

	(1) Unemployment	(2) GDP	(3) Avg. disp. income	(4) Gender empl. gap
GDP	-0.00133***			0.00009
	(0.00022)			(0.00013)
Unemployment		-33.14539***	-43.57921***	0.04985**
		(5.88857)	(6.03629)	(0.01839)
Median age	-0.00000***	-0.00103***	-0.00101***	0.00000
	(0.00000)	(0.00008)	(0.00012)	(0.00000)
lllearn	0.00000***	0.00011***	0.00008*	-0.00000
	(0.00000)	(0.00002)	(0.00003)	(0.00000)
htec	-0.00000	0.00001	-0.00001	0.00000***
	(0.00000)	(0.00001)	(0.00001)	(0.00000)
CO2 emissions	-0.00379***	0.04963	0.41385***	0.00031
	(0.00035)	(0.06910)	(0.06570)	(0.00022)
Patent (EPO)	0.04429***	5.53345***	-7.76134**	0.00604
	(0.00641)	(1.48728)	(2.46167)	(0.00555)
EPS market	-0.00001***	-0.00150***	-0.00016	0.00000
	(0.00000)	(0.00013)	(0.00021)	(0.00000)
EPS non-market	0.00000***	-0.00118***	-0.00156***	-0.00000
	(0.00000)	(0.00009)	(0.00017)	(0.00000)
N	4160	4160	4160	4160
adj. R^2	0.666	0.953	0.973	0.662
NUTS 2 FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
lllearn is the Participation rate in education and training;
htec is the Employment in technology and knowledge-intensive sectors.

1.4.2 ML model performance and feature importance

The results in Table 1.3, while informative as a baseline, are likely inflated by multicollinearity among predictors, resulting in unstable and unreliable coefficients. In contrast, the Elastic Net method incorporates L1 and L2 penalties, performing feature selection and coefficient shrinkage to filter out redundant or irrelevant information. This approach results in more robust and interpretable coefficients, effectively clearing out the "noise" that complicates causal inference. Although Elastic Net sacrifices a direct claim to causality, its advantages in improving model stability and interpretability outweigh this trade-off for the scope of our analysis.

The results, illustrated in Figures 1.3 to 1.6, highlight the importance of different variables in predicting, respectively, unemployment, GDP, average disposable income, and the gender employment gap at the regional level, as identified by the EN model. The performance of the fitted models is in Table 1.4.

Table 1.4: Model performance across target variables

	Unemployment	GDP	Disposable income	Gender employment gap
R^2	0.86853	0.78002	0.73820	0.55797
RMSE	1.90575	0.00732	0.00193	5.56332
Normalized RMSE	0.39969	0.48830	0.52652	0.96535
MAE	1.37613	0.00420	0.00150	4.17179
MAPE	0.29566	0.13772	0.09322	0.49441

Notes: Performance variables are computed on the test set.

With an EN approach, the R^2 improves relative to OLS in all but one outcome, reflecting the regularization benefits of the EN model in a setting with highly collinear covariates. Unemployment, GDP, and disposable income are predicted with accuracy, with R^2 of 0.87, 0.78, and 0.73, respectively. For unemployment and gender employment gap, which are expressed in percentage points, the MAE provides the most interpretable measure of fit. In their case, the MAE indicates a 1.3% and a 4% average mismatch between predicted and actual values[6]. The gender employment gap proves the most difficult to predict, suggesting that its regional variation is driven by factors that are not fully captured by the available predictors. Overall, the model performs well across most outcomes, providing a reliable basis for the scenario analysis that follows.

Figures 1.3 - 1.6 present a feature importance analysis for the four outcome variables of interest. Feature importance analysis measures the relative contribution of each predictor to the model's forecasting performance. In the context of an EN, importance is assessed through the magnitude of the estimated coefficients after penalization – larger absolute coefficients indicating greater predictive relevance. While this reflects statistical association rather than causal impact, identifying which inputs carry the most information for forecasting remains a useful exercise that sheds light on which variables the model relies on most, allowing us to assess whether the results align with economic intuition or whether unexpected patterns emerge. Beyond this exploratory aim, feature importance also serves a diagnostic purpose, providing a direct window into the model's variable selection outcome.

[6]For these variables, an elevated MAPE despite a reasonable R^2 likely reflects disproportionately large proportional errors in regions where the rates are low, rather than a general model failure.

Figure 1.3 shows that, aside from the regional employment rate (*emp*), the most influential feature explaining unemployment is its spatial lag (*spatial_ lag_ unemp*), *i.e.*, the unemployment levels registered in neighbouring regions, highlighting that proximity and concentration effects have a key role in shaping local labour market outcomes. EPS indicators show notable relevance, although with different signs. While some indicators, like the renewable energy trading scheme (*c_ renets*), the public R&D expenditure in low carbon technologies (*c_ lowrd*), the tax on sulphur oxides (*c_ sulfax*), and wind energy support policies(*c_ wind*) show a negative link with unemployment, other indicators go in the opposite direction. Overall, these results support the hypothesis that national environmental protection strategies play a crucial role in shaping local unemployment patterns. However, the impact of each specific measure seems widely heterogeneous, warranting further research. Participation in lifelong learning (*lllearn*) and tertiary education attainment (*ted*), in percentage, are associated with higher unemployment rates. Coupled with the negative sign of *median*, we can interpret this finding as an indication that demographically older regions tend to show higher unemployment, requiring reskilling progammes to take place. Other key predictors include the share of employment in the construction sector (*nacef_ rel*) and in vulnerable sectors (*vuln_ rel*), suggesting a strong link between regional employment levels and sectoral composition. Interestingly, patent applications (*pat_ count_ rel*) show a relatively weaker direct influence of innovation-related variables on unemployment.

Figure 1.3: Feature importance plot - Unemployment

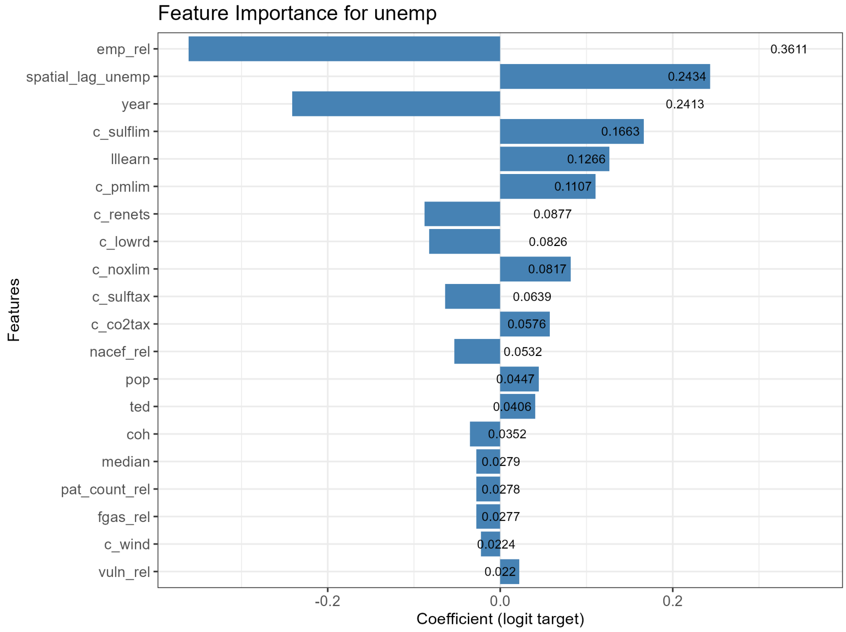

Notes: Feature importance plot indicating the relative contribution of predictors to the fitted model. The intercept ranked highest but was omitted for readability. The dependent variable was logit-transformed to ensure EN predictions remain within the 0-100 range.

Similarly, for regional GDP (Figure 1.4), the spatial lag (*spatial_ lag_ gdp*)—representing GDP in neighbouring regions— emerges as the most influential factor, further suggesting the spatially interconnected nature of regional economic activities. As it is to be expected, overall employment levels (*emp_ rel*) serve as a key driver of regional GDP. Among sectoral employment variables, employment in the high-tech sector (*htec_ rel*) positively affects regional GDP, while the opposite can be said about employment in the industrial sector (*nacebe_ rel*) and in vulnerable sectors (*vuln_ rel*). This contrasting effect reflects the importance of high-value-added employment in regional economies. The presence of tertiary education attainment (*ted*) and eco-innovative patents (*ecopat_ count_ rel*) among the top drivers, underscores the critical role of human capital and social factors in determining economic output. In contrast to unemployment, environmental regulation indicators from the EPS dataset play a less prominent role in explaining regional GDP variations, suggesting that environmental regulation may have more direct effects on labour market outcomes than on overall economic output. On the other hand, pro capita CO_2 emissions (*co2_ rel*) are associated with higher GDP levels, indicating how eco-

Figure 1.4: Feature importance plot - Regional GDP

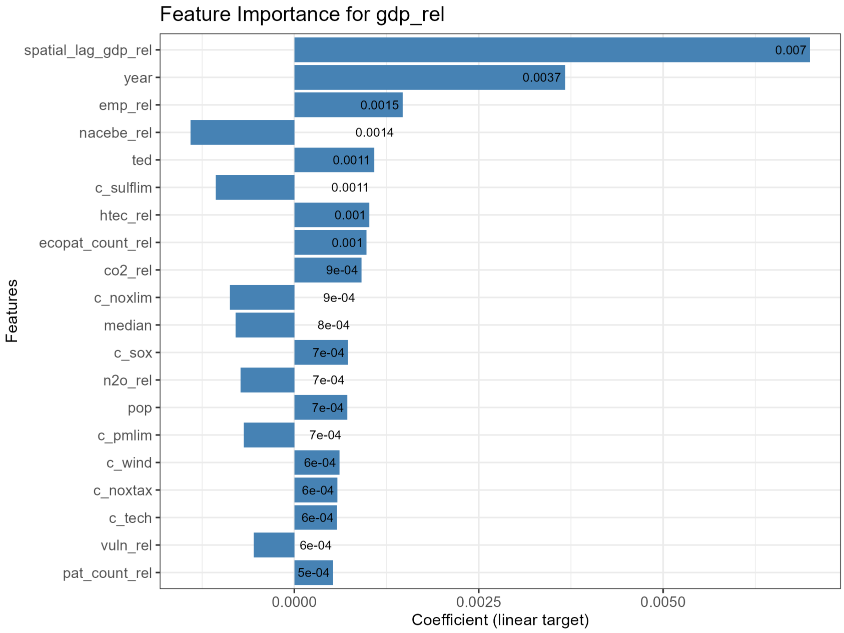

Notes: Feature importance plot indicating the relative contribution of predictors to the fitted model. The intercept ranked highest but was omitted for readability.

nomic performance comes at the price of a higher environmental footprint in most European regions.

Figure 1.5 shows that the main predictor of regional average disposable income is its spatial lag (*spatial_lag_disp_rel*), suggesting once again that concentration effects are a substantial component of regional economic development and ought not to be overlooked. A higher disposable income is driven upwards by a higher employment rate and downwards by a higher share of employment in the industrial sector and the construction sector (*nacebe_rel* and *nacef_rel*). In this context, regions that are better equipped in terms of eco-innovative activities (*ecopat_count_rel*) have higher disposable income. Note that the same cannot be said of overall patenting activity (*pat_count_rel*). These results jointly reflect the importance of both labour market conditions and innovation activities. EPS indicators give a few interesting indications. We remark that emission limits (*c_noxlim* and *c_pmlim* generally point in the direction of a lower disposable income. The same applies for carbon tax (*c_co2tax*) and for solar energy support (*c_solar*). Conversely, ex-

Figure 1.5: Feature importance plot - Average disposable income

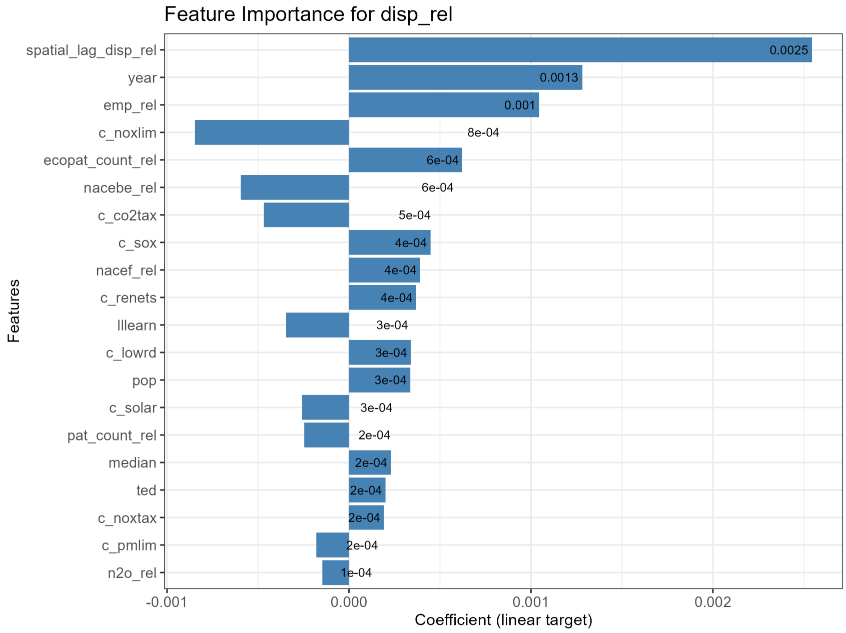

Notes: Feature importance plot indicating the relative contribution of predictors to the fitted model. The intercept ranked highest but was omitted for readability.

penditure on low emission R&D (*c_ lowrd*) and renewable energy trading systems (*c_ renets*) point in the opposite direction. These contrasting effects, coupled with the role played by *ecopat_ count_ rel*, suggest that regions investing proactively in green transition capabilities—rather than merely complying with environmental restrictions—experience measurable economic benefits in terms of disposable income.

Finally, for the gender employment gap (Figure 1.6), the most important predictor is still the spatial lag *spatial_ lag_ egap* (*i.e.*, the gender employment gap in neighbouring regions), suggesting that labour market characteristics are spatially concentrated, with pockets of EU areas where women are more persistently excluded from employment, both within their own region and in adjacent ones. The presence of employment variables (*nacef_ rel*, *htech_ rel*, and the aggregated vulnerable sectors *vuln_ rel*), reinforces the idea that sectoral composition not only plays a critical role in influencing overall unemployment levels, but also shapes gender disparities in employment. Interestingly, three out of four GHG emissions rank among the main predictors. In particular, we note how nitrous oxide emissions (*n2o_ rel*), and

Figure 1.6: Feature importance plot - Gender employment gap

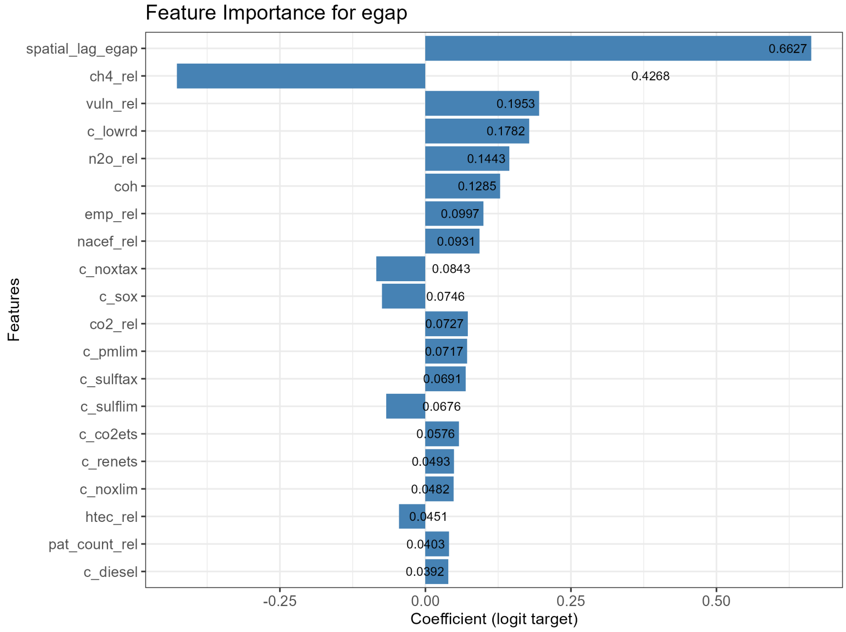

Figure 1.7: Notes: Feature importance plot indicating the relative contribution of predictors to the fitted model. The intercept ranked highest but was omitted for readability. The dependent variable was logit-transformed to ensure EN predictions remain within the 0-100 range.

CO_2 emissions (*co2_rel*) show a positive correlation with the gender employment gap. This pattern suggests that women are not uniformly distributed across sectors vulnerable (or less vulnerable) to the sustainability transition, and that their lower/higher presence in certain industries may amplify gendered disparities in labour market outcomes.

Taken together, these results show that all models are strongly driven by spatial spillovers and employment levels across sectors. While innovation and education are critical drivers of long-term economic growth and welfare, their effects on employment levels may be more indirect, *e.g.*, potentially mediated through industrial transformation and regional preparedness. Overall, the results suggest that regions are strongly influenced by the socio-economic and employment conditions of their neighbours, highlighting the spatially interconnected nature of these dynamics. In contrast, national-level environmental protection measures exhibit limited predictive power for regional outcomes. This finding suggests that, although national institutional capacity plays a role in determining socioeconomic outcomes, regional heterogeneity and spatial spillover effects are more decisive factors. Moreover, all the features we use to construct the scenarios consistently emerge among the most significant predictors. These empirical patterns provide robust validation for our analytical framework. First, they substantiate the theoretical justification for our variable selection in the shock scenarios. Second, they demonstrate that adopting a regional focus coupled with explicit modeling of spatial spillover effects yields substantially different insights compared to conventional regression approaches, while corroborating *a priori* theoretical expectations.

Finally, a comparative analysis of Figures 1.3–1.6 and Table 1.3 reveals several noteworthy discrepancies that warrant examination in light of theoretical expectations. Regarding unemployment, patent applications demonstrate negligible predictive importance despite exhibiting statistically significant associations in the OLS regression. The same applies to CO_2. Conversely, in the ML analysis of GDP, employment in high-tech and knowledge intensive sector emerges as an important predictor in the feature importance analysis, notwithstanding its lack of significance in the OLS framework. This result can be attributed to the EN approach to managing multicollinearity: the high-tech variable effectively captures the variability associated with human capital endowments while maintaining a more direct correlation with GDP outcomes compared to other variables. The case of average disposable income yields particularly intriguing findings. Environmental regulations (EPS) appear to exert greater influence on regional income levels than GHG emissions, contrasting with the OLS results where EPS market measures failed to reach the 0.1 significance threshold, while the coefficient of CO_2 emissions is statistically significant (and of non-negligible magnitude). Gender employment gap analysis presents the most striking divergences between methodological approaches. While the OLS regression identifies unemployment rates as significant predictors, the EN feature importance analysis assigns minimal weight to this variable. Instead, EPS indicators and greenhouse gas emissions emerge as prominent determinants, supporting the hypothesis that women are not uniformly distributed across sectors with different levels of vulnerability to the sustainability transition, thereby linking environmental performance and protection mechanisms to gender equity outcomes in the labour market. These discrepancies highlight the value of a comparative approach, as different methods uncover complementary dimensions of the complex relationships between environmental performance, sectoral structures, and labour market outcomes. At the same time, they also highlight the limitations of traditional regression approaches (whose linear framework struggles to capture the multidimensional interactions underpinning labour market outcomes) and the advantages of ML methods, which, by accommodating complex and non-linear relationships, provide valuable insights into the intertwined dynamics involved in the just and sustainable transition.

1.4.3 What if?

This section discusses the regional outcomes of the four simulated scenarios, highlighting the differentiated effects on unemployment, GDP, disposable income, and gender employment gap across European regions. The results are provided in Figures 1.8 to 1.10, which represent, respectively, the scenarios: 1) Enhanced innovative capacity; 2) Emission cap; 3) Labour market policies; 4) Policy mix scenario.

Starting from Scenario 1 *"Enhanced innovative capacity"* (Figure 1.8), while innovation is often seen as a driver of economic growth and competitiveness, our findings suggest that focusing exclusively on boosting innovative capacity may lead to unintended negative consequences. Unemployment changes vary across regions: many regions, especially in Northern Europe, including large areas of Sweden, Denmark, and Belgium, display reductions in unemployment, whereas most of Eastern Europe, along with pockets of Southern and Western Europe—including areas in Greece and Spain—show increases. The employment gains in Northern Europe suggest that these regions capitalize on enhanced innovation to expand job opportunities. However, in Eastern Europe, structural labour market rigidities or sectoral mismatches might prevent innovation from translating into widespread employment benefits.

On the other hand, under this scenario, the effects on GDP are more scattered, with very heterogeneous outcomes coexisting within the same country. One notable exception is Poland, which is to benefit uniformly. Improvements are expected also in Mediterranean regions of Spain and France, and in Southern Italy. These positive changes in GDP align with the expectation that stronger innovation fosters economic growth, especially in regions lagging behind or characterised by less mature innovation ecosystems which have stronger learning scope (Barabuffi et al., 2025). In contrast, the negative effects in certain regions might signal trade-offs due to resource reallocation or competitive disadvantages due to disparities in absorptive capacity.

Compared to GDP outcomes, disposable income displays a few noticeable divergences: many Northern regions (*e.g.* Ireland, Benelux and Denmark) benefit, whereas most of France, and parts of Spain and Finland see limited or negative changes. Regions characterised by lower income levels (such as in Eastern Europe, Greece, regions in Central Italy and Southern Spain) seem to benefit relatively more. In several instances, we observe a divergence between GDP and disposable income outcomes, notably in regions such as Occitanie (France), Basilicata (Italy), and Southern Finland. This disparity suggests that innovation-driven economic growth does not always translate into widespread welfare improvements, as unequal access to opportunities can prevent many from benefiting. Since knowledge and capital accumulation may disproportionately favour certain groups (exacerbating inequality rather than reducing it), without inclusive policies, innovation risk deepening existing socio-economic disparities rather than fostering welfare for all.

The results on gender employment gap show a mix of positive and negative effects across Europe, with significant improvements in some regions (*i.e.*, strong reduction in the gender employment gap), but also signs of widening disparities elsewhere. This variation suggests that fostering innovation (and, more broadly, innovation-led transitions) can support gender equality in employment when they generate opportunities in high-skill, knowledge-intensive sectors that tend to be more gender-inclusive. However, in regions where the labour market remains dominated by male-intensive or lower-skill sectors, or where structural barriers persist (*e.g.*, limited access to STEM careers or care-related constraints), these benefits may not materialize. This evidence seems to align with recent Eurostat data, which show that while the EU gender employment gap has slightly narrowed in recent years, significant differences remain in female employment rates in technology-intensive sectors (*i.e.*, men accounting for over two-thirds of the total, on average)[7]. This underscores the need for complementary gender-sensitive policies to ensure that innovation-driven transitions do not reinforce existing inequalities.

Moving to Scenario 2 *"Emission cap"* (Figure 1.9), our findings highlight a more pronounced country effect, particularly evident in Hungary, Sweden, and Poland. This finding underscores the critical role of national measures, including early implementation of policy stringency and a proactive approach to green growth, in

[7] Eurostat (2023)

Figure 1.8: Scenario 1 - Enhanced innovative capacity

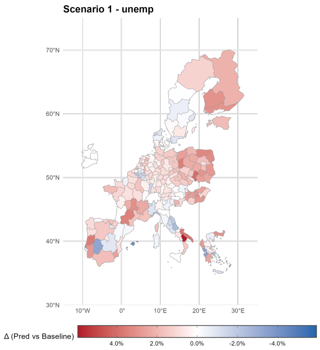

(a) Unemployment

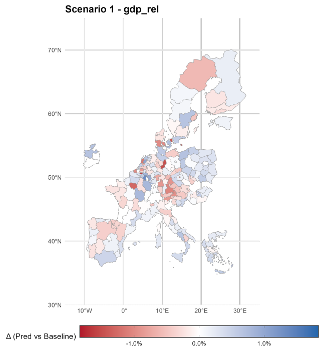

(b) GDP

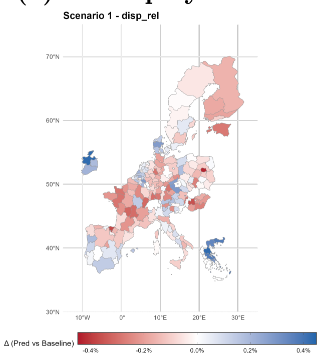

(c) Disposable income

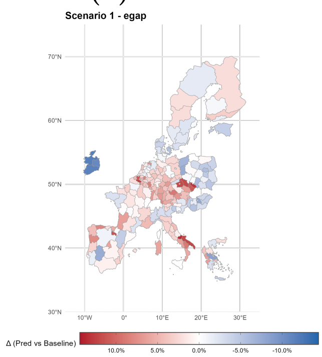

(d) Gender employment gap

Notes: White indicates no change; shades of blue show improvements over the status quo (*i.e.*, higher GDP and disposable income, lower unemployment or gender gap) and shades of red show worsening, using a per-map symmetric scale around zero; numbers are percentage-point changes from baseline.

shaping regional outcomes under an emissions cap. While the labour market effects are predominantly negative—manifested in rising unemployment rates consistent with the literature discussed in Section 1.2—it is noteworthy that many regions exhibit a reduction in unemployment. Indeed, our analysis indicates that 71 out of 199 regions are projected to experience employment growth.

The projection on GDP draws a mostly positive picture, with widespread growth patterns observed in Eastern and Northern Europe, as well as France, suggesting that the emission cap may be aligned with these regions' green transition goals (and current energy mix). In contrast, negative effects are more pronounced in Spain and other Central continental regions. However, the results in terms of disposable income show a more nuanced picture. While some European regions experience significant income gains (most notably in Ireland and Greece), most regions in France, Germany, Poland, Hungary, Sweden, and Finland are characterized by very negative outcomes. The positive changes may reflect the policy's alignment with regional industrial transformation strategies, leading to income growth; negative effects may signal a higher cost of compliance with environmental policies. Occasionally, a divergence emerges between GDP growth and disposable income outcomes. For instance, while most regions in France are projected to experience GDP growth, they simultaneously face a decline in disposable income. Conversely, the opposite pattern emerges in Southern and Eastern Ireland, where an emission cap is associated with higher disposable income but a reduction in GDP. This divergence suggests that immediate improvements in households' capacity to consume, invest, and save may not align with broader economic growth. One possible explanation lies in a socio-political landscape that prioritizes corporations and capital accumulation in the first case, while favouring household benefits in the second. Additionally, in certain regions (such as in Southern Spain) an emission cap is projected to lead to increases in both GDP and disposable income, while simultaneously leading to higher unemployment rates. This finding provides an intriguing perspective on the complex interplay between economic performance and inequality, suggesting that positive macroeconomic indicators can sometimes mask underlying issues of social exclusion.

Similar to the previous figure, the implications in terms of gender employment gap are quite heterogeneous across regions in this scenario, too.

The results for Scenario 3, *Labour market policies* (Figure 1.10), underline the complex relationship between labour market restructuring and socio-economic outcomes. In this scenario, unemployment declines in most regions, more markedly than in previous cases, though significant within-country disparities remain. These findings suggest that gradual yet sustained worker reallocation away from vulnerable sectors is critical to limiting labour market disruptions. These results strengthen the rationale for policy frameworks prioritizing reskilling and upskilling, in line with the aim and targets of the Just Transition Fund. Regarding GDP, our findings show that regions shifting away from vulnerable, low-value-added sectors toward services and high-tech industries generally experience widespread GDP gains. How-

Figure 1.9: Scenario 2 - Emission cap

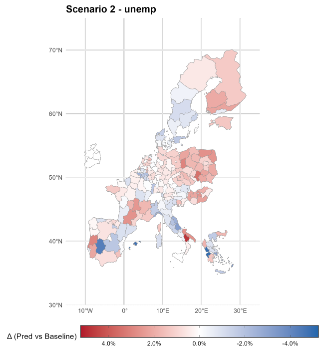

(a) Unemployment

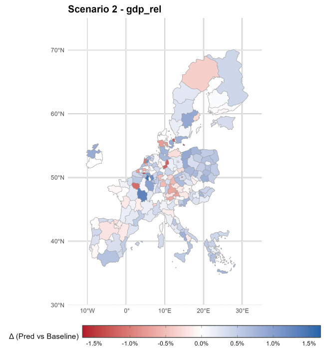

(b) GDP

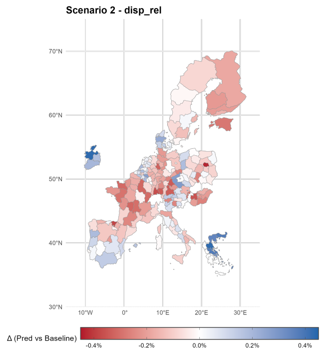

(c) Disposable income.

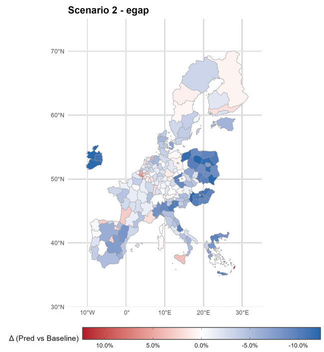

(d) Gender employment gap

Notes: White indicates no change; shades of blue show improvements over the status quo (*i.e.*, higher GDP and disposable income, lower unemployment or gender gap) and shades of red show worsening, using a per-map symmetric scale around zero; numbers are percentage-point changes from baseline.

ever, these gains are not consistently reflected in disposable income outcomes. In several areas (of Poland and Germany, namely in Ślaskie and in the area of North Rhine-Westphalia), disposable income growth remains relatively weak despite significant GDP increases. In the Spanish regions of Navarre and Madrid, and in the Stockholm region, disposable income decreases, contrary to GDP growth. This divergence signals a potential risk that the green transition, while boosting overall economic output, may not translate into commensurate benefits for workers and households. This further underscores the critical need for policies that support upskilling and re-skilling to prevent social distortions that extend beyond economic performance. Furthermore, the gender employment gap follows a notably more negative trajectory, with the widest disparities observed so far, particularly across all regions of Germany and Eastern Europe. This finding raises equity concernsmarket's and suggests that the labour market resilience to the green transition may disproportionately exclude vulnerable groups, exacerbating existing inequalities.

Figure 1.10: Scenario 3 - Labour market policies

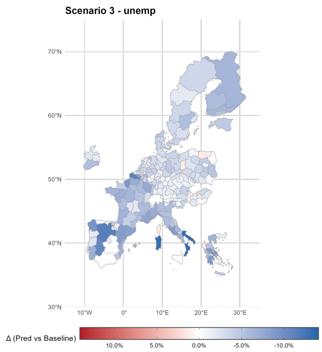

(a) Unemployment

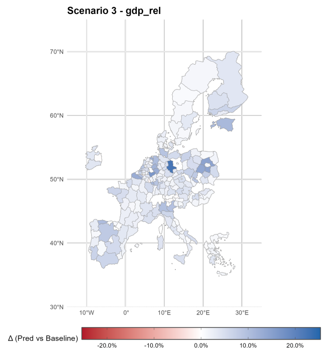

(b) GDP

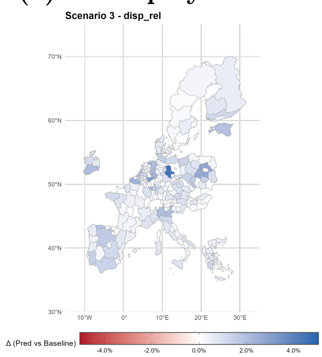

(c) Disposable income.

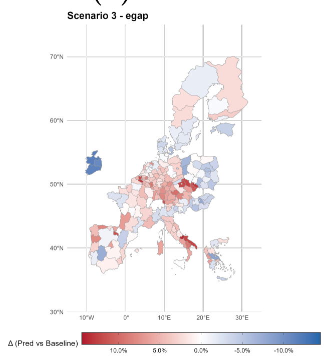

(d) Gender employment gap

Notes: White indicates no change; shades of blue show improvements over the status quo (*i.e.*, higher GDP and disposable income, lower unemployment or gender gap) and shades of red show worsening, using a per-map symmetric scale around zero; numbers are percentage-point changes from baseline.

The results of Scenario 4, *Policy mix* (Figure 1.11), closely resemble those of Scenario 3. From a methodological point of view, this similarity can be explained by the prominent role of sectoral composition variables in our predictions (see Section 1.4.2), but it also captures a well-established insight from the literature: labour market dynamics and the transition away from traditional employment represent the most disruptive aspect of the sustainability transition.

However, we observe that combining all three policy measures (*i.e.*, enhanced innovation capacity, GHG emission reduction, and labour market transition) helps mitigate the negative effects on gender employment gaps, although it shows limited impact on other target indicators. This finding is particularly significant, as it suggests that comprehensive policy mixes can compensate for the disruptive effects of individual measures, thereby promoting greater inclusion and social justice during the green transition.

Figure 1.11: Scenario 4 - Policy mix

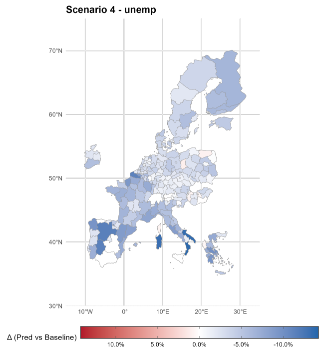

(a) Unemployment

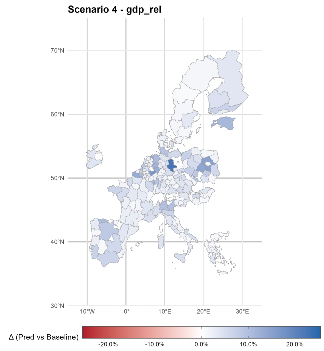

(b) GDP

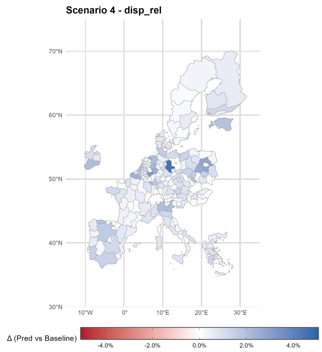

(c) Disposable income.

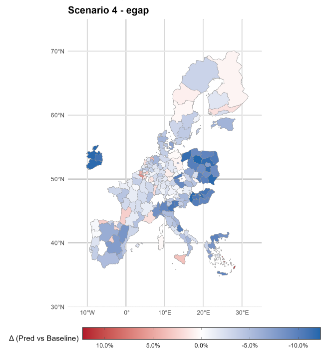

(d) Gender employment gap

Notes: White indicates no change; shades of blue show improvements over the status quo (*i.e.*, higher GDP and disposable income, lower unemployment or gender gap) and shades of red show worsening, using a per-map symmetric scale around zero; numbers are percentage-point changes from baseline.

To provide a more concise assessment of the results across the different scenarios and to better capture distributive outcomes, we summarize our findings in Table 1.5. This further analysis reveals many relevant insights. First, we notice how scenario 4 produces, on average, the biggest improvement in all target variables (*i.e.*, lower unemployment, higher GDP, higher income and lower gender gap), reinforcing the idea that a comprehensive policy action is the most beneficial when dealing with the multidimensional challenges of the sustainability transition. At the same time however, scenario 4 also yields the most detrimental distributional effects, as captured by the Gini index. In contrast, Scenarios 1 and 2 yield more equally distributed benefits (or costs), as shown by the lower Gini indices for unemployment, GDP, and disposable income, with an improvement compared to the observed values for 2019. It is worth highlighting that the gender employment gap, on the other hand, faces unfavorable distributional consequences in all scenarios. These findings hint at a possible equity-efficiency trade-off in sustainability transition policies, and act as a warning of possible conflicts between environmental goals and cohesion in the EU.

Lastly, we conduct a sensitivity analysis to test the robustness of our findings by re-estimating all three scenarios under four alternative shock intensities, *i.e.*, ±5% and ±10% changes to the shocked variable (see Table 1 in Appendix). Across these specifications, the main patterns of regional heterogeneity and outcome remain consistent and the fluctuations in model results are substantially smaller than the changes applied to scenario parameters, providing strong evidence of the model's stability and the robustness of our comparative assessment across policy scenarios.

	Unemployment		**GDP_rel**		**Disp_rel**		**Egap**	
	Mean	Gini	Mean	Gini	Mean	Gini	Mean	Gini
Actual	6.59495	0.36870	0.03210	0.25218	0.01719	0.12174	11.14394	0.28164
Predicted	7.33976	0.32795	0.03172	0.20107	0.01672	0.10956	13.12435	0.30782
Scenario 1 — Predicted	7.44429	0.32228	0.03131	0.19762	0.01692	0.11291	12.82930	0.31290
Scenario 2 — Predicted	6.90786	0.33080	0.03345	0.18494	0.01675	0.10928	7.68602	0.32009
Scenario 3 — Predicted	2.69753	0.59374	0.07219	0.31786	0.02493	0.19528	13.12435	0.30782
Scenario 4 — Predicted	2.56977	0.59151	0.07351	0.31069	0.02515	0.19512	7.50262	0.32553

Table 1.5: Mean and Gini index of actual vs. predicted outcomes under the four scenarios. Notes: Baseline refers to the EN baseline prediction for 2019, before applying scenario-specific shocks.A sensitivity analysis is provided in the Appendix.

1.5 Conclusions and further research

The analysis presented in this study highlights the complex and heterogeneous dynamics of sustainability transitions across subnational regions, demonstrating that policies, even if uniformly designed, can produce divergent socio-economic outcomes at the local level. While some areas benefit from GDP growth and higher income, others experience higher unemployment, increasing income disparities, and widening gender employment gaps. These findings reinforce the importance of including social justice and equity considerations into policy frameworks to ensure that sustainability transitions are not only environmentally effective but also socially inclusive.

ML techniques, as employed in this study, offer a valuable tool for identifying local vulnerabilities and providing evidence-based recommendations to guide policy strategies. The results from the simulated scenarios (enhanced innovative capacity, emissions cap, labour market restructuring, the combined policy mix) suggest that sustainability transitions are neither linear nor uniformly beneficial across space. Innovation, a key driver of economic growth, does not inherently lead to inclusive development. Unequal access to opportunities and the concentration of benefits among certain groups may result in rising inequalities rather than fostering welfare improvements for all. Similarly, while policies aimed at emissions reduction can generate economic gains, they may simultaneously lead to rising unemployment, particularly in regions dependent on carbon-intensive industries. This suggests that positive macroeconomic indicators can sometimes mask underlying vulnerabilities, making it essential to consider local context when designing policy interventions. The results also highlight the potential risks that transition patterns pose to labour markets, and moving away from vulnerable sectors can have particularly adverse consequences in regions more exposed to the transition risks (where job displacement is likely to be more pronounced). Without targeted measures, the labour market adjustments necessary for a green transition risk exacerbating rather than alleviating social disparities.

These insights have significant policy implications, as our results suggest that a uniform, top-down approach to sustainability transitions may not be able to ensure equitable outcomes.

Indeed, from one side, the impact of climate and energy policies (typically coordinated at the national level) is highly context-dependent, reinforcing the idea that one-size-fits-all solutions are insufficient. At the same time, this appears in line with the intertwined evolution of the EU environmental and cohesion agenda, which jointly support the green transition but also the need for territorially differentiated policy strategies that consider both EU-level coordination and region-specific needs and capacities. In this context, the Just Transition Fund (JTF) emerges as a relevant instrument in the EU policy, aimed at supporting the economic diversification and reconversion of the regions most negatively affected by the net-zero transition (*e.g.*, through up- and reskilling of workers, or job-search assistance). While the JTF, which aim to support investments through territorially targeted measures in areas exposed to transition risks, provides much-needed resources, this fund alone cannot fully address all the challenges associated with the sustainable transition. Its success will depend on the ability of national and regional governments to effectively implement transition strategies that go beyond short-term economic compensation and focus on long-term structural change. This requires strengthening labour market policies, investing in education and innovation, and ensuring that green industrial policies do not reinforce existing inequalities. Ultimately, achieving a just and sustainable transition is not only a technological or economic challenge, but a deeply social one.

The heterogeneity of gender gap effects, which we treat as a proxy for the broader structural inclusivity of the labour market, represents an important yet often overlooked dimension of the just transition debate. Our results reveal pronounced spatial variation in transition outcomes, underscoring that policy schemes must be explicitly designed to prevent exacerbating pre-existing inequalities at the subnational level. The relevance of gender employment balance in this context is further reinforced by its identification as a key component of regional labour market efficiency (Borsekova and Korony, 2023). Our findings suggest that no single policy measure can uniformly guarantee inclusive outcomes. Rather, addressing inclusivity in transition processes hinges critically on the meaningful incorporation of local contexts and the active participation of regional stakeholders in the policy design process.

Another dimension that tends to be underappreciated in policy design is the role of spatial spillovers in shaping and anticipating regional outcomes. Our analysis finds that spatial interdependencies carry substantial predictive information about what unfolds in a given region, and yet these signals rarely inform policy action. We argue that greater attention should be devoted to the agglomeration and spatial diffusion effects of transition policies — whether these manifest as positive externalities, such as knowledge spillovers, or negative ones, such as lock-in effects. One promising avenue in this direction is the promotion of inter-regional coordination mechanisms that grant subnational actors greater autonomy in managing cross-boundary interdependencies. The Interreg programmes represent a valuable step in this direction, fostering local consortia and cross-border collaboration within the EU framework. More broadly, our results caution against space-blind policy approaches that disregard the complexity of interrelated local contexts, and instead point toward the adoption of network-aware policy frameworks that explicitly account for spatial interconnectedness — an approach whose effectiveness has already been documented in the regional development literature (Ramajo et al., 2008).

While this study provides important contributions to understanding the socio-economic effects of sustainability transitions, it also opens several avenues for future research. Notably, our analysis has focused on the isolated effects of each policy intervention, and has treated the interactions among them only marginally. Future studies should explore the dynamic interplay across these dimensions, *i.e.*, how emissions mitigation, innovation incentives, and labour market policies interact over time, leading to either reinforcing or counteracting effects. Future research should also explicitly incorporate additional policy drivers (such as welfare systems and education), as well as examine how the interaction between different levels of governance influences transition outcomes. Understanding how these policy domains intersect, overlap, or conflict will be key to designing coherent and just transition strategies across diverse territorial contexts. The success of such efforts, however, hinges on the availability of high-quality data capable of faithfully capturing local contexts.

Bibliography

F. Avelino, K. Wijsman, F. van Steenbergen, S. Jhagroe, J. Wittmayer, S. Akerboom, K. Bogner, E. F. Jansen, N. Frantzeskaki, and A. Kalfagianni. Just sustainability transitions: politics, power, and prefiguration in transformative change toward justice and sustainability. *Annual Review of Environment and Resources*, 49, 2024.

P.-A. Balland, R. Boschma, J. Crespo, and D. L. Rigby. Smart specialization policy in the european union: relatedness, knowledge complexity and regional diversification. *Regional Studies*, 53(9):1252–1268, 2019.

S. Barabuffi, V. Costantini, V. Leone Sciabolazza, and E. Paglialunga. Knowledge spillovers through high-skilled migration network: evidence from oecd countries. *Industry and Innovation*, pages 1–31, 2025.

N. Barbieri, D. Consoli, L. Napolitano, F. Perruchas, E. Pugliese, and A. Sbardella. Regional technological capabilities and green opportunities in europe. *Journal of Technology Transfer*, 48:749–778, 4 2023. ISSN 15737047.

F. Biancalani, G. Gnecco, R. Metulini, and M. Riccaboni. The impact of the european union emissions trading system on carbon dioxide emissions: a matrix completion analysis. *Scientific reports*, 14(1):19676, 2024.

S. Bianchini, G. Damioli, and C. Ghisetti. The environmental effects of the "twin" green and digital transition in european regions. *Environmental and Resource Economics*, 84:877–918, 4 2023. ISSN 15731502.

K. Borsekova and S. Korony. Resilience and vulnerability of regional labour markets: principal component analysis of labour market efficiency in the eu. *Regional Studies*, 57(12):2373–2390, 2023.

F. Cappellano, A. M. Santos, and N. F. Dotti. Regional ri ventures to tackle climate change: A new geography of challenge-oriented innovation landscape. *Papers in Regional Science*, 103:100052, 10 2024. ISSN 1056-8190.

C. Cappelli, V. Costantini, and D. Consoli. The trap of climate change-induced "natural" disasters and inequality. *Global Environmental Change*, 70:102329, 2021.

F. Castellacci. Innovation and social welfare: A new research agenda. *Journal of Economic Surveys*, 37(4):1156–1191, 2023.

J. Castellanos and G. Heutel. Unemployment, labor mobility, and climate policy. Technical Report 25797, NBER, 2021.

G. Cerulli. Fundamentals of supervised machine learning. *Cham: Springer*, 2023.

Cohesion Open Data Platform. Cohesion open data platform. URL `https://cohesiondata.ec.europa.eu/`. Accessed: September 2025.

D. Consoli, G. Marin, A. Marzucchi, and F. Vona. Do green jobs differ from non-green jobs in terms of skills and human capital? *Research Policy*, 45(5):1046–1060, 2016.

V. Costantini, F. Crespi, G. Marin, and E. Paglialunga. Eco-innovation, sustainable supply chains and environmental performance in european industries. *Journal of Cleaner Production*, 155:141–154, 7 2017. ISSN 0959-6526.

V. Costantini, F. Crespi, and E. Paglialunga. The employment impact of private and public actions for energy efficiency: Evidence from european industries. *Energy Policy*, 119:250–267, 2018.

V. Costantini, C. Martini, B. Mina, and M. Zoli. The welfare impact of climate action: A distributional analysis for italy. *Energy Economics*, 143:108181, 2025. ISSN 0140-9883.

A. Crespy and M. Munta. Lost in transition? social justice and the politics of the eu green transition. *Transfer: European Review of Labour and Research*, 29(2): 235–251, 2023.

G. Damioli, S. Bianchini, and C. Ghisetti. The emergence of a 'twin transition'scientific knowledge base in the european regions. *Regional Studies*, pages 1–17, 2024.

A. Diemer, S. Iammarino, A. Rodríguez-Pose, and M. Storper. The regional development trap in europe. *Economic Geography*, 98:487–509, 10 2022. ISSN 19448287.

P. L. Donti and J. Z. Kolter. Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, 46(1):719–747, 2021.

EDGAR – Emissions Database for Global Atmospheric Research. Edgar – emissions database for global atmospheric research. URL `https://edgar.jrc.ec.europa.eu/`. Accessed: January 2024.

D. L. Edmondson, F. Kern, and K. S. Rogge. The co-evolution of policy mixes and socio-technical systems: Towards a conceptual framework of policy mix feedback in sustainability transitions. *Research Policy*, 48(10):103555, 2019.

G. F. W. G. K. Elzen, Boelie. System innovation and the transition to sustainability: Theory, evidence and ... - google libri, 2004.

Environmental Policy Stringency (EPS) Index Dataset. Environmental policy stringency (eps) index dataset. URL `https://www.oecd.org/environment/indicators-modelling-outlooks/environmental-policy-stringency.htm`. Accessed: January 2024.

European Commission. Summary report 2023 on the implementation of the european structural and investment funds. Technical report, European Commission, Directorate-General for Regional and Urban Policy, 2023. URL `https://ec.europa.eu/regional_policy/sources/reports/annual_2023/2023_Annual_Summary_Report.pdf`. Accessed: [insert date].

Eurostat. Eurostat regional statistics database. URL `https://ec.europa.eu/eurostat/data/database`. Accessed: January 2024.

Eurostat. Which EU regions employ more women in high-tech? Eurostat News, Oct. 2023. URL `https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20231030-1`. Accessed: 2026-03-05.

I. Faiella and L. Lavecchia. Households' energy demand and the effects of carbon pricing in italy. Technical Report 614/2021, Bank of Italy, 2021.

E. Frohm, F. M. D'arcangelo, T. Kruse, and M. Pisu. Environmental policy stringency and co2 emissions: Evidence from cross-country sector-level data. 2013.

C. Ghisetti and F. Pontoni. Investigating policy and rd effects on environmental innovation: A meta-analysis. *Ecological Economics*, 118:57–66, 10 2015. ISSN 0921-8009.

M. Grillitsch and T. Hansen. Green industry development in different types of regions. *European Planning Studies*, 27(11):2163–2183, 2019.

J. Grin, J. Rotmans, and J. Schot. Transitions to sustainable development: New directions in the study of long term transformative change. *Transitions to Sustainable Development: New Directions in the Study of Long Term Transformative Change*, pages 1–397, 1 2010.

K. Grossmann, J. J. Connolly, M. Dereniowska, G. Mattioli, L. Nitschke, N. Thomas, and A. Varo. From sustainable development to social-ecological justice: Addressing taboos and naturalizations in order to shift perspective. *Environment and Planning E: Nature and Space*, 5(3):1405–1427, 2022.

S. Iammarino, A. Rodriguez-Pose, and M. Storper. Regional inequality in europe: evidence, theory and policy implications. *Journal of Economic Geography*, 19: 273–298, 3 2019. ISSN 1468-2702.

O. W. Johnson, J. Y. C. Han, A. L. Knight, S. Mortensen, M. T. Aung, M. Boyland, and B. P. Resurrección. Intersectionality and energy transitions: A review of gender, social equity and low-carbon energy. *Energy Research Social Science*, 70: 101774, 12 2020. ISSN 2214-6296.

G. King and L. Zeng. The dangers of extreme counterfactuals. *Political Analysis*, 14(2):131–159, 2006.

T. Kortetmäki, C. Timmermann, and T. Tribaldos. Just transition boundaries: Clarifying the meaning of just transition. *Environmental Innovation and Societal Transitions*, 55:100957, 2025.

A. Kronsell. Gender and transition in climate governance. *Environmental Innovation and Societal Transitions*, 7:1–15, 6 2013. ISSN 2210-4224.

M. Kuhn. Building predictive models in R using the caret package. *Journal of Statistical Software*, 28(5):1–26, 2008. doi: 10.18637/jss.v028.i05.

S. Kumar, S. Datta, V. Singh, D. Datta, S. K. Singh, and R. Sharma. Applications, challenges, and future directions of human-in-the-loop learning. *IEEE Access*, 2024.

J. Köhler, F. W. Geels, F. Kern, J. Markard, E. Onsongo, A. Wieczorek, F. Alkemade, F. Avelino, A. Bergek, F. Boons, L. Fünfschilling, D. Hess, G. Holtz, S. Hyysalo, K. Jenkins, P. Kivimaa, M. Martiskainen, A. McMeekin, M. S. Mühlemeier, B. Nykvist, B. Pel, R. Raven, H. Rohracher, B. Sandén, J. Schot, B. Sovacool, B. Turnheim, D. Welch, and P. Wells. An agenda for sustainability transitions research: State of the art and future directions. *Environmental Innovation and Societal Transitions*, 31:1–32, 6 2019. ISSN 2210-4224.

J. E. Maldonado, A. Vandeplas, I. V. Vigani, and A. Turrini. Assessing green job dynamics in the eu: A comparison of alternative methodologies, 7 2024.

G. Marin and F. Vona. Climate policies and skill-biased employment dynamics: Evidence from eu countries. *Journal of Environmental Economics and Management*, 98:102253, 2019.

S. Markkanen and A. Anger-K raavi. Social impacts of climate change mitigation policies and their implications for inequality. *Climate Policy*, 19(7):827–844, 2019.

L. Napolitano, A. Sbardella, D. Consoli, N. Barbieri, and F. Perruchas. Green innovation and income inequality: A complex system analysis. *Structural Change and Economic Dynamics*, 63:224–240, 2022.

P. Newell and D. Mulvaney. The political economy of the 'just transition'. *The Geographical Journal*, 179:132–140, 6 2013. ISSN 1475-4959.

OECD. The jobs potential of a shift towards a low-carbon economy: Final report for the eu commission, dg employment, 2012.

K. O'Sullivan, O. Golubchikov, and A. Mehmood. Uneven energy transitions: Understanding continued energy peripheralization in rural communities. *Energy Policy*, 138:111288, 2020.

W. Oueslati, V. Zipperer, D. Rousselière, and A. Dimitropoulos. Energy taxes, reforms and income inequality: An empirical cross-country analysis. *International Economics*, 150:80–95, 2017.

E. Paglialunga, A. Coveri, and A. Zanfei. Climate change and within-country inequality: New evidence from a global perspective. *World Development*, 159: 106030, 2022.

F. Perruchas, D. Consoli, and N. Barbieri. Specialisation, diversification and the ladder of green technology development. *Research Policy*, 49(3):103922, 2020. ISSN 0048-7333.

B. Purvis, Y. Mao, and D. Robinson. Three pillars of sustainability: in search of conceptual origins. *Sustain Sci*, 14:681–695, 2019.

J. Ramajo, M. A. Marquez, G. J. Hewings, and M. M. Salinas. Spatial heterogeneity and interregional spillovers in the european union: Do cohesion policies encourage convergence across regions? *European Economic Review*, 52(3):551–567, 2008.

REGPAT. Regpat database – oecd regional patent database, 2024. URL `https://www.oecd.org/sti/inno/regpat-database.htm`. Accessed: January 2024.

A. Rodríguez-Pose and F. Bartalucci. The green transition and its potential territorial discontents. *Cambridge Journal of Regions, Economy and Society*, 17(2): 339–358, 11 2023. ISSN 1752-1378.

D. Rosenbloom. Pathways: An emerging concept for the theory and governance of low-carbon transitions. *Global Environmental Change*, 43:37–50, 2017.

D. Rosenbloom. Building momentum for a 'policy turn'in sustainability transitions: Lessons from canada to consolidate strengths and bridge science-policy divides. *Environmental Innovation and Societal Transitions*, 55:100956, 2025.

B. K. Sovacool, B. Turnheim, A. Hook, A. Brock, and M. Martiskainen. Dispossessed by decarbonisation: Reducing vulnerability, injustice, and inequality in the lived experience of low-carbon pathways. *World Development*, 137:105116, 2021.

N. Stern and J. E. Stiglitz. *The social cost of carbon, risk, distribution, market failures: An alternative approach*, volume 15. National Bureau of Economic Research Cambridge, MA, USA, 2021.

M. Swilling. *The age of sustainability: Just transitions in a complex world.* Taylor & Francis, 2020.

J. Taylor and R. J. Tibshirani. Statistical learning and selective inference. *Proceedings of the National Academy of Sciences*, 112(25):7629–7634, 2015.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58:267–288, 1 1996. ISSN 1369-7412.

B. Turnheim and B. Nykvist. Opening up the feasibility of sustainability transitions pathways (stps): Representations, potentials, and conditions. *Research Policy*, 48 (3):775–788, 2019.

N. van Bommel and J. I. Höffken. The urgency of climate action and the aim for justice in energy transitions – dynamics and complexity. *Environmental Innovation and Societal Transitions*, 48:100763, 2023. ISSN 2210-4224.

A. Vandeplas, I. Vanyolos, M. Vigani, and L. Vogel. The possible implications of the green transition for the eu labour market, 12 2022.

F. Vona. Managing the distributional effects of environmental and climate policies: The narrow path for a triple dividend, 2021.

F. Vona. Managing the distributional effects of climate policies: A narrow path to a just transition. *Ecological Economics*, 205:107689, 2023. ISSN 0921-8009.

P. Xie, Y. Xu, X. Tan, and Q. Tan. How does environmental policy stringency influence green innovation for environmental managements? *Journal of Environmental Management*, 338:117766, 7 2023. ISSN 0301-4797.

Z. Yao, Y. Lum, A. Johnston, L. M. Mejia-Mendoza, X. Zhou, Y. Wen, A. Aspuru-Guzik, E. H. Sargent, and Z. W. Seh. Machine learning for a sustainable energy future. *Nature Reviews Materials*, 8(3):202–215, 2023.

B. Zaman and S. Borsky. The impact of supply structure on solar home system installations in rural off-grid areas. *Environmental Innovation and Societal Transitions*, 40:625–644, 2021.

H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67:301–320, 4 2005. ISSN 1369-7412.

1.A Appendix

1.A.1 Recoding of Regions to 2021 NUTS2

Croatia

In 2021, the region HR04 Continental Croatia (HR04) was split into Pannonian Croatia (HR02), City of Zagreb (HR05), and Northern Croatia (HR06). For this case, our approach involved creating a proportional relationship between the population of the new regions and HR04, based on population. In particular, after performing a linear extrapolation of the population of the smaller regions (available since 2013), the patent counts of HR02, HR05, and HR06 were proportionally assigned based on their share of the total population in HR04.

France

In the case of France, Regpat reports the code of the second-level metropolitan departments, under standard ISO 3166-2. In this case, it was sufficient to verify in which NUTS2 region the metropolitan department was located and to substitute the appropriate NUTS2 classification. Note that more than one metropolitan department can be contained in a NUTS2 region. Specifically:

1. FR21 was replaced by FRC1.
2. FR22 was replaced by FRH0.
3. FR23 and FR24 were both replaced by FRI1.
4. FR25 was replaced by FRC1.
5. FR26 was replaced by FRK2.
6. FR30 was replaced by FRJ1.
7. FR41 was replaced by FRB0.
8. FR42 and FR43 were replaced by FRK2.
9. FR51 and FR52 were replaced by FRF2.

10. FR53 was replaced by FRG0.
11. FR61 was replaced by FRD1.
12. FR62 was replaced by FRE1.
13. FR63 was replaced by FRK2.
14. FR71 was replaced by FRC1.
15. FR72 was replaced by FRG0.
16. FR81 and FR82 were replaced by FRJ2.
17. FR83 was replaced by FRL0.

Hungary

In Hungary, applications coming from the regions of Budapest (HU11) and Pest (HU12) are reported jointly under the code HU10. In this case, too, the patent counts for HU11 and HU12 were assigned according to population share.

Ireland

In Ireland, a similar approach was taken for regions IE05 and IE06, which replaced IE02 after the latest NUTS revision in 2016. The population of IE05 and IE06 was extrapolated when not available (before 2012); then the patent counts were proportionally scaled to reflect their share of the total population in IE02.

Lithuania

For Lithuania, Regpat does not distinguish between regions LT01 (Sostines regionas) and LT02 (Vidurio ir vakary Lietuvos regionas). Instead, it codifies all applications as coming from LT00. The patent counts for LT01 and LT02 were adjusted based on their share of the total Lithuanian population.

Poland

Several regions in Poland were recoded following the 2018 revision.

1. PL11 was replaced by PL71.
2. PL31 was replaced by PL81.
3. PL32 was replaced by PL82.
4. PL33 was replaced by PL72.
5. PL34 was replaced by PL84.

Another couple of regions in Poland, PL91 and PL92, resulted from the disaggregation of the previous PL12 region. In this case, the patent counts were scaled according to their population share (extrapolation was necessary before 2014).

1.A.2 Features: descriptive statistics

In Table 1.A.1 we present the descriptive statistics for the features included in our analysis. Note that these are computed over the whole panel 2000 to 2019.

1.A.3 Durbin type regression

As a complement to the OLS analysis presented in Section 1.4, we replicate the same regression exercise augmenting each equation with the spatial lag of the dependent variable, following a Durbin-type specification. While an exact replication of the EN is not feasible in an OLS setting — including the full set of variables used in the EN model would result in a computationally singular system — this approach offers a closer comparison by combining spatial dependence with the same set of regressors reported in Table 1.3.

The Durbin regression results reported in Table 1.A.2 allow us to assess whether the divergences observed between the OLS and elastic net frameworks persist once spatial dependence is accounted for in a linear setting. For the most part, the Durbin specification does not resolve the discrepancies identified above, and in some cases sharpens them.

Table 1.A.1: Descriptive Statistics

Variable	**Mean**	**Median**	**SD**	**Min**	**Max**
coh	0.834	1.000	0.372	0.000	1.000
ch4_rel	0.00144	0.00122	0.00106	0.000110	0.00931
co2_rel	0.00869	0.00721	0.00505	0.00144	0.04440
fgas_rel	0.000224	0.000191	0.000166	0.0000245	0.004683
n2o_rel	0.000609	0.000458	0.000482	0.0000333	0.003496
c_co2ets	1.426	1.000	1.253	0.000	4.000
c_co2tax	0.704	0.000	1.569	0.000	6.000
c_diesel	3.936	4.000	0.813	0.000	6.000
c_envpol	2.799	2.889	0.738	0.528	4.722
c_lowrd	2.319	2.000	1.469	0.000	6.000
c_market	1.434	1.167	0.804	0.333	4.167
c_nonmark	4.744	5.250	1.059	1.250	6.000
c_noxlim	4.377	5.000	1.461	0.000	6.000
c_noxtax	1.278	0.000	1.931	0.000	6.000
c_pmlim	4.376	6.000	1.950	0.000	6.000
c_renets	0.548	0.000	1.402	0.000	6.000
c_solar	2.066	2.000	1.896	0.000	6.000
c_sox	4.673	5.000	0.823	0.000	6.000
c_sulflim	5.549	6.000	0.500	4.000	6.000
c_sulftax	0.709	0.000	1.320	0.000	6.000
c_tech	2.221	2.250	1.210	0.000	6.000
c_wind	2.180	2.000	1.797	0.000	6.000
disp_rel	0.01449	0.01463	0.00424	0.000445	0.03008
egap	13.887	12.700	8.170	0.000	104.200
emp_rel	0.000426	0.000430	0.0000547	0.000155	0.000539
gdp_rel	0.02568	0.02539	0.01238	0.003165	0.10160
htec_rel	0.03660	0.03385	0.01996	0.000	0.13232
lllearn	9.428	7.500	6.597	0.000	36.000
median	41.377	41.200	3.286	28.000	51.700
nacea_rel	0.05368	0.03395	0.05835	0.000883	0.37501
nacebe_rel	0.19333	0.18673	0.07216	0.032598	0.38995
nacef_rel	0.07571	0.07250	0.02219	0.020574	0.32706
pop	1,954,886	1,480,048	1,635,841	118,861	12,252,917
ted	24.214	24.200	11.037	0.000	105.200
unemp	8.874	7.600	5.388	1.200	36.300
ecopat_count_rel	0.0000172	0.00000741	0.0000250	0.000	0.000239
pat_count_rel	0.000132	0.0000726	0.000171	0.000	0.001298
vuln_rel	0.3227	0.3166	0.09411	0.10152	0.58843
Spatial lags					
spatial_lag_unemp	8.879	8.115	3.689	3.653	21.300
spatial_lag_gdp_rel	0.02534	0.02701	0.009344	0.006668	0.04536
spatial_lag_disp_rel	0.01444	0.01515	0.002853	0.006252	0.01960
spatial_lag_egap	13.962	13.183	5.957	3.142	32.563

Notes: Statistics computed over the full panel (2000–2019).

Table 1.A.2: OLS regression model with regional and year fixed effects

	(1) Unemployment	(2) GDP	(3) Avg. disp. income	(4) Gender empl. gap
Spatial lag (dep. var.)	−33.78965***	−38.44451***	−38.37997***	−23.59836***
	(5.23059)	(5.58094)	(6.40773)	(4.03114)
Median age	−0.33591*	−0.00011	0.00006	−0.04794
	(0.13726)	(0.00020)	(0.00011)	(0.20801)
lllearn	0.16160***	0.00001	−0.00006**	−0.15804***
	(0.03270)	(0.00003)	(0.00002)	(0.04783)
htec	−2.67441	−0.01903*	−0.03183***	14.42917
	(12.10897)	(0.00868)	(0.00940)	(16.01831)
CO_2 emissions	−431.92209***	0.02587	0.09764*	299.30849*
	(103.84783)	(0.10608)	(0.04911)	(143.36255)
Patent (EPO)	474.04271	1.29712	0.22540	5551.61073*
	(1018.92874)	(2.82120)	(1.45268)	(2526.20455)
EPS market	−0.42406	−0.00086***	−0.00000	1.62505***
	(0.24194)	(0.00022)	(0.00012)	(0.31846)
EPS non-market	0.57678***	−0.00078***	−0.00071***	−0.26737
	(0.12197)	(0.00014)	(0.00010)	(0.17407)
N	4160	4160	4160	4160
NUTS 2 FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
adj. R^2	0.57218	0.58630	0.66178	0.53628

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. *lllearn*: participation rate in education and training; *htec*: employment in technology and knowledge-intensive sectors.

An interesting finding is that the insignificance of patents for unemployment in the EN is mirrored in the Durbin, where patents also fail to reach significance. This suggests that their significance in the OLS was likely an artefact of omitted spatial correlation rather than a genuine predictive relationship. The CO_2–unemployment association, by contrast, remains negative and highly significant in the Durbin, indicating that this is a robust linear correlation, even if the EN does not regard it as a sizeable source of predictive information.

For GDP and disposable income, the Durbin confirms the negative and significant role of non-market EPS instruments, partially bridging the gap between the OLS and EN results. However, the prominence of high-technology employment in the elastic net's GDP feature importance rankings is not recovered by the Durbin. Instead, *htec* reaches only marginal significance. For the gender employment gap, the Durbin corroborates the positive and significant association with market-based EPS instruments, consistent with the EN's emphasis on environmental policy variables as key determinants. However, emission variables, which were identified as prominent predictors by the EN, do not survive in the Durbin framework.

A noteworthy finding concerns model fit. The adjusted R^2 declines considerably under the Durbin specification relative to the OLS, particularly for GDP and disposable income — precisely the two outcomes for which the EN achieves its best out-of-sample predictions. This deterioration is indicative of multicollinearity between the spatial lag and the remaining regressors, making it harder for the linear model to attribute predictive variation to specific variables. These results underscore the limitations of the linear framework in settings characterised by high collinearity.

Overall, the Durbin results do not undermine our substantive interpretation or the policy implications derived from the EN analysis. Rather, they reinforce two of our central methodological claims: first, that spatial interdependencies among regions are a fundamental feature of the data that any credible empirical framework must account for; and second, that combining ML with theory-grounded econometric specifications can foster a more complete understanding of the complex socio-economic dynamics associated with the sustainability transition.

1.A.4 Descriptive statistics of target variables under scenarios and sensitivity analysis

Table 1.A.3 reports the results of a sensitivity analysis in which the original scenario parameters are subjected to variations of ±5% and ±10%. Note that the sensitivity tests confirm the robustness of our results, as all three metrics exhibit variations only in the decimal places for all target variables and under all four scenarios.

Table 1.A.3: Distribution of actual vs. predicted outcomes and robustness checks.

	Unemployment			**GDP_rel**			**Disp_rel**			**Egap**		
Row	SD	Mean	Gini	SD	Mean	Gini	SD	Mean	Gini	SD	Mean	Gini
Actual	4.75598	6.59495	0.36870	0.01495	0.03210	0.25218	0.00366	0.01719	0.12174	5.74846	11.14394	0.28164
Baseline predicted	4.75367	7.33976	0.32795	0.01114	0.03172	0.20107	0.00326	0.01672	0.10956	7.81840	13.12435	0.30782
scenario 1 — Predicted	4.73306	7.44429	0.32228	0.01082	0.03131	0.19762	0.00339	0.01692	0.11291	7.82303	12.82930	0.31290
scenario 1 — Robustness + 5%	4.73313	7.44390	0.32230	0.01082	0.03131	0.19763	0.00339	0.01691	0.11290	7.82297	12.83035	0.31288
scenario 1 — Robustness - 5%	4.73299	7.44469	0.32226	0.01082	0.03131	0.19761	0.00339	0.01692	0.11292	7.82308	12.82825	0.31292
scenario 1 — Robustness + 10%	4.73320	7.44351	0.32232	0.01082	0.03131	0.19764	0.00339	0.01691	0.11288	7.82292	12.83140	0.31286
scenario 1 — Robustness - 10%	4.73292	7.44508	0.32224	0.01082	0.03131	0.19761	0.00339	0.01692	0.11292	7.82314	12.82720	0.31294
scenario 2 — Predicted	4.52188	6.90786	0.33080	0.01079	0.03345	0.18494	0.00326	0.01675	0.10928	5.13508	7.68602	0.32009
scenario 2 — Robustness + 5%	4.52182	6.90776	0.33080	0.01079	0.03345	0.18494	0.00326	0.01675	0.10928	5.13460	7.68503	0.32009
scenario 2 — Robustness - 5%	4.52193	6.90796	0.33080	0.01079	0.03345	0.18495	0.00326	0.01675	0.10928	5.13557	7.68702	0.32008
scenario 2 — Robustness + 10%	4.52177	6.90766	0.33080	0.01079	0.03345	0.18494	0.00326	0.01675	0.10928	5.13412	7.68404	0.32010
scenario 2 — Robustness - 10%	4.52198	6.90806	0.33080	0.01079	0.03345	0.18495	0.00326	0.01675	0.10929	5.13605	7.68801	0.32008
scenario 3 — Predicted	3.42148	2.69753	0.59374	0.04989	0.07219	0.31786	0.01041	0.02493	0.19528	7.81840	13.12435	0.30782
scenario 3 — Robustness + 5%	3.37825	2.61044	0.60248	0.05229	0.07422	0.32303	0.01088	0.02534	0.20001	7.81840	13.12435	0.30782
scenario 3 — Robustness - 5%	3.46547	2.79002	0.58462	0.04750	0.07017	0.31245	0.00994	0.02452	0.19045	7.81840	13.12435	0.30782
scenario 3 — Robustness + 10%	3.33577	2.52829	0.61086	0.05469	0.07625	0.32797	0.01136	0.02576	0.20464	7.81840	13.12435	0.30782
scenario 3 — Robustness - 10%	3.51026	2.88846	0.57516	0.04511	0.06814	0.30681	0.00947	0.02411	0.18552	7.81840	13.12435	0.30782
scenario 4 — Predicted	3.23517	2.56977	0.59151	0.04983	0.07351	0.31069	0.01044	0.02515	0.19512	5.14099	7.50262	0.32553
scenario 4 — Robustness + 5%	3.19426	2.48645	0.60026	0.05223	0.07554	0.31602	0.01091	0.02556	0.19977	5.14047	7.50230	0.32551
scenario 4 — Robustness - 5%	3.27681	2.65828	0.58239	0.04743	0.07148	0.30511	0.00997	0.02474	0.19037	5.14152	7.50294	0.32554
scenario 4 — Robustness + 10%	3.15406	2.40786	0.60869	0.05464	0.07757	0.32111	0.01139	0.02597	0.20433	5.13994	7.50198	0.32549
scenario 4 — Robustness - 10%	3.31921	2.75248	0.57290	0.04503	0.06945	0.29926	0.00950	0.02433	0.18553	5.14204	7.50326	0.32556

Notes: SD = standard deviation across NUTS2 regions; Gini computed on regional levels.

Chapter 2

Predicting Regional Unemployment in the EU

This essay is co-authored with Elena Paglialunga (Roma Tre University) and Giuliano Resce (University of Molise).

Abstract

This paper predicts regional unemployment in the European Union by applying machine learning techniques to a dataset covering 198 NUTS-2 regions, 2000 to 2019. Tree-based models substantially outperform traditional regression approaches for this task, while accommodating reinforcement effects and spatial spillovers as determinants of regional labor market outcomes. Inflation—particularly energy-related—emerges as a critical predictor, highlighting vulnerabilities to energy shocks and green transition policies. Environmental policy stringency and eco-innovation capacity also prove significant. Our findings demonstrate the potential of machine learning to support proactive, place-sensitive interventions that aim to predict and mitigate the uneven social and economic impacts of structural change across regions.

Keywords: Regional unemployment; Inflation; Environmental policy; Spatial spillovers; Machine learning.

JEL Classification: E24; J64; Q52; R23.

2.1 Introduction

Unemployment is a central concern for European policymakers, particularly at the regional level, where economic shocks and structural imbalances produce persistent spatial disparities (Rodríguez-Pose, 2013). These differences not only reflect unequal access to employment opportunities but also challenge the effectiveness of cohesion policy, social protection systems, and the political legitimacy of European integration (Dijkstra et al., 2015). Anticipating regional unemployment trends is therefore essential for the timely design and targeting of active labor market policies (Boeri and Terrell, 2002), especially in light of recent crises that have revealed the vulnerability of specific regions to economic downturns (Bagliano and Morana, 2025; Iammarino et al., 2019). These vulnerabilities are further exposed in the context of the green transition, where ambitious climate and energy policies, while crucial for decarbonisation, may introduce challenges linked to uneven sectoral exposures, regional labour market resilience, and the distributive impacts of higher energy and production costs (Vona, 2023).

This paper develops a predictive framework to anticipate short-term unemployment rates across NUTS2 regions in Europe. Framing the task as a prediction policy problem (Kleinberg et al., 2015), it becomes evident the practical value of forecasting for policy. In particular, timely and reliable predictions can help identify regions most at risk of rising unemployment, thereby supporting the targeting of place-based interventions such as cohesion policies, where the effectiveness of action depends critically on anticipating rather than merely reacting to shocks. For tasks of this kind, where the goal is accurate out-of-sample prediction rather than unbiased estimation of the effects of a policy on the outcome variable, machine learning (ML) methods are a powerful ally to economic analysis (Kleinberg et al., 2015; Mullainathan and Spiess, 2017).

Traditional econometric models used in labor economics, such as fixed-effects panel regressions or structural macro models, often prioritize inference over prediction and tend to perform poorly in out-of-sample forecasting tasks (Stock and Watson, 2002). In contrast, ML techniques are explicitly designed to maximize predictive accuracy and can flexibly capture complex, non-linear relationships and high-dimensional interactions among variables (Athey and Imbens, 2019). Recent research has demonstrated the value of ML methods across a variety of policy domains, ranging from criminal justice to social welfare. In particular, these techniques are increasingly being applied to address challenges such as poverty targeting (Jean et al., 2016), evaluating the effectiveness of public programs and spending (Andini et al., 2018; Christensen et al., 2024; Di Stefano and Resce, 2025), analyzing education and training systems (Sansone, 2019; Sansone and Zhu, 2023), and identifying instances of financial distress, corruption, and political connections (Antulov-Fantulin et al., 2021; D'Ecclesiis et al., 2025; Titl et al., 2024). They have also proven useful in the analysis of environmental and health-related issues (Caravaggio et al., 2025; Carrieri et al., 2021; Fabra et al., 2022; Girma and Paton, 2024). In the context of labor market analysis, recent studies have applied ML to forecast aggregate unemployment in the euro area (Gogas et al., 2022) and in countries like Germany (Katris, 2020). However, these approaches often overlook subnational disparities—an aspect we aim to address more directly.

We apply a set of ML models to forecast unemployment rates at the NUTS-2 level, using a rich panel dataset that collects a variety of regional socioeconomic indicators. In line with recent literature, we also incorporate environmental variables to capture potential trade-offs between green transition policies and labor market outcomes. We include spatial lags to account for geographic spillovers, following studies on the spatial dimension of unemployment. Additionally, we include inflation-related indicators to reflect the well-established unemployment-inflation relationship. Our level of spatial granularity is particularly valuable, as NUTS-2 regions are a key target of European cohesion policy and the allocation of related funding instruments. Our models has the potential to support proactive policy design by identifying which regions are likely to experience rising unemployment in the near future—providing early-warning signals that can inform the deployment of European Social Fund resources, the design of retraining programs, and the planning of counter-cyclical investment initiatives.

Our findings show that ML models—particularly Random Forest and Gradient Boosting—effectively capture the complex, non-linear dynamics behind regional unemployment in the EU. Key predictors include lagged unemployment and its spatial lag (*i.e.*, the unemployment rates of neighboring regions), highlighting the relevance of geographical concentration effects in shaping unemployment patterns. Environmental policy and environmental performance also emerge as important predictors. Inflation-related variables retain notable influence, while challenging the validity of the Phillips relationship in local contexts.

The remainder of the paper is structured as follows: Section 2.2 reviews the relevant literature, Section 2.3 presents our methodology and data, Section 2.4 discusses the empirical results, and Section 2.5 concludes with policy implications and directions for future research.

2.2 Literature Review

2.2.1 The unemployment-inflation relationship in the green transition era

Unemployment dynamics are strongly shaped by macroeconomic factors such as economic growth and price fluctuations, with uneven consequences that amplify regional disparities. Shocks to real disposable income, whether from energy and food price increases, or broader economic crises, directly affect consumption capacity and firms' labour demand, thereby reinforcing spatial disparities in unemployment outcomes (Iammarino et al., 2019; Rodríguez-Pose, 2013). These vulnerabilities have been repeatedly exposed during recent crises, from the economic downturn to the pandemic and the current energy price shock, where weak income growth combined with inflationary pressures exacerbate regional labour market fragilities.

Within this broader macroeconomic context, the Phillips Curve provides a classical framework for examining the inflation–unemployment nexus, originally formulated as an inverse relationship between change of monetary wages and unemployment (Phillips, 1958) and later revised in the more familiar form of an inverse relationship between the rate of change of prices and those of unemployment (Samuelson and Solow, 1960). .

Over time, empirical studies have provided mixed evidence regarding the stability of this relationship. Some research suggests that the Phillips Curve has flattened in recent decades due to more stable inflation expectations (Stock and Watson, 2008). This finding is in line with expectations-augmented version of the relationship, as introduced by Friedman (1968) and Phelps (1967), who argued that the curve turns vertical in the long run, stabilizing around a "natural rate" of unemployment.

In the short-run, however, the problem persists (Gordon, 2018) and decades of research highlighted many open questions in the short-run Phillips relationship. Some argue that regional variations and structural breaks must be considered to better understand its empirical validity (Babb and Detmeister, 2017; Hooper et al., 2020; Smith, 2024). Despite these issues, renewed interest has emerged, especially in light of contemporary economic conditions and recent inflationary episodes observed in the last years (Hazell et al., 2022). A substantial body of literature is concerned with empirically assessing the occurrence of the curve. One line of research focuses on the re-emergence of empirical support for the Phillips hypothesis, suggesting that inflation-unemployment trade-offs remain relevant (Hazell et al., 2022). Other studies emphasize the role of stable inflation expectations in reducing the sensitivity of inflation to unemployment (Stock and Watson, 2008). Further research underscores the heterogeneity in Phillips Curve coefficients across regions (Babb and Detmeister, 2017; Hooper et al., 2020), while additional contributions examine structural breaks that arise due to economic shifts (Smith, 2024). Despite these findings, key unresolved issues persist, including the potential non-linearity of the Phillips relationship (Babb and Detmeister, 2017), the challenge of disentangling the direction of causality between inflation and employment (Hazell et al., 2022), and the endogeneity concerns in estimating the relationship (Smith, 2024).

Recent global events have underscored two critical considerations: the imperative to accelerate energy transition for achieving strategic autonomy, and the adverse macroeconomic consequences associated with abrupt shifts in energy mix. Indeed, climate and energy policies designed to reduce carbon dependency and accelerate the shift to renewable sources may entail short-term increases in energy prices, whether through carbon taxation, the phasing out of subsidies, or the costs of renewable integration. While essential for long-term sustainability, these measures may introduce new inflationary pressures that interact with labour markets in non-trivial ways. On the one hand, higher energy and input costs raise firms' production costs, potentially leading to lower output and weaker labour demand. On the other hand, rising consumer prices erode households' real disposable income, reducing consumption and further depressing economic activity. Together, these dynamics undermine the classical Phillips trade-off and can give rise to episodes in which inflation and unemployment may rise simultaneously, underscoring the importance of incorporating both price and income dynamics into the analysis of unemployment. These developments have sparked interest in examining the nexus between climate performance and inflationary pressures (Boneva and Ferrucci, 2022; Moessner, 2022).

In this light, the inflation–unemployment relationship cannot be interpreted solely as a matter of monetary expectations, as it increasingly reflects also how regions absorb the distributive consequences of the energy transition, with uneven effects across space and sectors. This perspective underscores the importance of accounting for climate and energy policy shocks in empirical analysis, and highlights the need for regionally tailored policies that can mitigate labour market vulnerabilities while sustaining progress towards decarbonisation.

2.2.2 Environmental policies and unemployment

Even though the green transition may not have sizeable effects on aggregate employment in the EU, it certainly introduces complexities to local labor markets, producing both job creation and destruction, and causing labor reallocation (Vandeplas et al., 2022). While some studies highlight the potential for green jobs to drive employment growth (Marin and Vona, 2019; Vona, 2023), others underscore the risks associated with the decline of traditional energy-intensive industries (OECD, 2012). Displaced workers often encounter significant barriers to reemployment, as green jobs (Consoli et al., 2016) may demand distinct skill sets, potentially leading to structural unemployment in affected regions. Moreover, the development of new and sustainable technological trajectories can have unintended consequences, particularly for low-skilled workers, who may face adverse employment effects as industries adapt to environmental regulations (Consoli et al., 2016).

Rodríguez-Pose and Bartalucci (2024) identify unemployment as an indicator of territorial discontent toward the green transition process. The effects of environmental policies on labor markets depend on local factors. Stringent regulations can alter sectoral compositions, leading to job creation in some areas while causing job losses elsewhere, particularly in regions dependent on carbon-intensive industries (Vona, 2023). The local labor market response is shaped by reconversion and migration costs, wealth distribution, and broader economic conditions. Evidence suggests that the impact of green transition policies varies based on labor mobility, the pace of green innovation and technological adoption (Costantini et al., 2018), existing income inequalities (Napolitano et al., 2022), market concentration and firm size distribution (Zaman and Borsky, 2021), and the evolving demand for green skills (Marin and Vona, 2019). These factors collectively determine the extent to which a shift to more sustainable productive systems can lead to an increase in unemployment.

2.2.3 The spatial dimension of unemployment

Since our analysis investigates labour market dynamics at the subnational level, the spatial dimension of unemployment becomes particularly salient. Indeed, economic activity is highly spatially concentrated in most countries, and this structure drives significant disparities in labour market outcomes. Regional unemployment is thus strongly shaped by local structural conditions, including sectoral specialization, innovation capacity, demographic composition, and institutional quality (Iammarino et al., 2019; Kline and Moretti, 2013). Yet, these factors rarely operate in isolation. Core regions often benefit from agglomeration economies, such as better access to capital, denser labor markets, and knowledge spillovers, while peripheral or less connected areas may face persistent disadvantages due to structural weaknesses, low innovation uptake, or poor infrastructure (Basile, 2009; Moretti, 2012). As a consequence, unemployment rates vary markedly across cities and regions. These spatial disparities are not merely reflections of individual characteristics but arise from geographically bounded dynamics. As documented in the literature on local labor markets, agglomeration effects, labor mobility, economic specialization, and the spatial sorting of workers and firms tend to reinforce regional divergence in wages and employment outcomes (Greenstone et al., 2010; Kline and Moretti, 2014; Moretti, 2011a,b, 2012). High-productivity regions attract capital and skilled labor, setting in motion virtuous cycles of employment growth. In contrast, regions with lower innovation capacity or overexposure to declining industries, including brown sectors, face the dual challenge of economic stagnation and outmigration (Iammarino et al., 2019). Importantly, economic shocks, whether local (*e.g.*, plant closures, automation) or global (*e.g.*, energy price volatility, climate policy reforms), rarely remain confined within administrative borders. They propagate across space via commuting flows, supply chains, and migration patterns, giving rise to significant spillover effects (Moretti, 2011a). These interdependencies challenge the assumption of independent regional labor markets and highlight the importance of modeling spatial interactions explicitly (Niebuhr, 2003). Evidence from the Italian context (Cracolici et al., 2007) shows that both high- and low-unemployment areas exhibit spatial clustering, confirming that labor market outcomes are subject to spatial autocorrelation. In particular, regions in Southern Italy have persistently high unemployment, often linked to structural disadvantages and weak integration with innovation-intensive northern areas. This pattern is mirrored across many EU countries and is closely tied to the legacy of path-dependent regional development trajectories, lock-in effects, further compounded by differences in innovation adoption, regional variety and firm dynamism (Corradini and Vanino, 2022; Overman and Puga, 2002). Such disparities have prompted policymakers to adopt place-based development strategies aimed at reducing territorial inequality. In the EU, Cohesion Policy explicitly targets lagging regions, recognizing the need for tailored interventions that address spatially concentrated disadvantages to foster regional resilience. Empirical evidence supports this approach: Caro and Fratesi (2023) found that Cohesion Policy bolsters labour market resilience during economic downturns. These concerns are particularly relevant in the context of the green transition since climate and environmental policies, such as carbon taxation or the phasing out of polluting industries, are expected to have asymmetric effects across space (Coenen et al., 2012). Regions

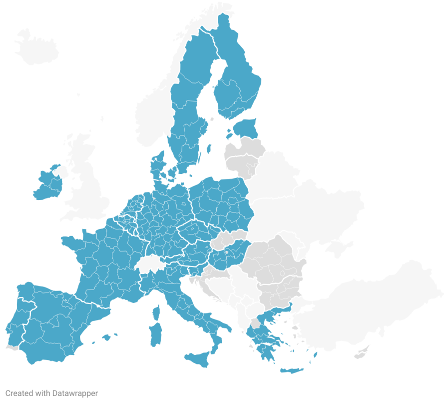

Figure 2.3.1: Map of NUTS2 regions included in the panel

with high concentrations of brown-sector employment may face significant labor displacement, while innovation hubs and urban centers stand to benefit from green job creation and investment flows (Rodríguez-Pose and Bartalucci, 2024; Vona, 2023).

This emerging research trajectory establishes important foundations for both theoretical frameworks and empirical investigations exploring the interconnections among labor market dynamics, sustainable economic transitions, and price stability.

2.3 Data and methods

2.3.1 Data

Our dataset incorporates regional data for EU NUTS 2 regions, supplemented with national data on policy settings, over a 20-year time frame (2000 to 2019). The dataset includes 38 variables (see Table 2.3.1 for a detailed description of variables' sources and codification in our dataset). To preserve as much information as possible, breaks in time series are interpolated. When observations for region i still present missing values, the region is dropped to ensure a balanced dataset. This operation leaves us with a panel of 198 regions out of the original 256: Figure 2.3.1 shows the regions included in the dataset.

The regional unemployment rate, expressed as a percentage of the labour force has been widely used to measure the societal impacts of socio-technical transitions (Castellanos and Heutel, 2019; Vona, 2021). The features considered for the prediction task are inspired by the literature review in Section 2.2. To assess regional economic disparities and potential development constraints related to climate policy implementation, we employ gross domestic product at the NUTS 2 level, expressed in million euros, as our primary indicator of overall economic development (Diemer et al., 2022; Iammarino et al., 2019). For examining household-level impacts of climate policies, we utilize disposable income data, given its direct connection to consumption patterns and behaviors (Rodríguez-Pose and Bartalucci, 2023). As an indicator of systemic inclusivity and intersectionality within socio-economic frameworks, we incorporate the gender employment gap across NUTS 2 regions, expressed as a percentage. This approach aligns with established research utilizing gender metrics as proxies for broader systemic inclusivity (Johnson et al., 2020). The importance of gender representation in climate governance for achieving equitable transitions is well-documented (Kronsell, 2013), particularly given evidence of continued male dominance in green employment sectors (Maldonado et al., 2024). Considering gender employment balance is even more relevant to our research objectives, as Borsekova and Korony (2023) see gender balance as a key dimension of regional labour market efficiency.

For capturing regional innovative capacity, we utilize patent data as a standard indicator (Barbieri et al., 2023). Specifically, we employ green patent counts to measure eco-innovative potential, which research consistently links to successful environmental transitions (Costantini et al., 2018; Napolitano et al., 2022). Our patent analysis draws from the OECD REGPAT Database, which compiles patent application data from both the European Patent Office (EPO) and Patent Cooperation Treaty (PCT) frameworks, though we focus exclusively on EPO applications. Using the Spring 2023 dataset version, we extract technological classification information through cooperative patent classification (CPC) codes. Patent applications are allocated to NUTS 2 regions based on inventor location, with our count including applications with priority years spanning 2000-2019. To identify eco-innovations specifically, we filter applications containing CPC code Y02, which designates inventions focused on climate change mitigation and adaptation.

Recognizing that patent metrics alone may incompletely represent regional innovation ecosystems, we incorporate additional knowledge capital indicators from the Regional Innovation Scoreboard (RIS), particularly given research emphasizing R&D support's role in climate preparedness (Cappellano et al., 2024). From the available RIS metrics, we select three key variables: employment within technology and knowledge-intensive industries, educational participation rates, and tertiary education completion rates for the 25-64 age cohort. While the latter two metrics are already expressed as percentages, we convert the employment data to relative terms by calculating per-capita values using regional population figures.

We also account for regional vulnerability through employment concentration in transition-sensitive sectors. Research consistently identifies dependence on specific productive sectors as a critical factor of vulnerability to the distributional impacts of climate policies (Rodríguez-Pose and Bartalucci, 2023; Vandeplas et al., 2022; Vona, 2021). To account for this strand of research, our analysis includes employment data for NACE sectors A through F at the NUTS 2 level, measured in thousands of workers.[1]

To control for demographic factors, we include regional median age as an additional variable capturing population structure dynamics. For environmental performance measurement, we utilize greenhouse gas emission indicators, which, while not comprehensively capturing all environmental impacts of human activities, remain widely accepted metrics for evaluating green transition progress and sustainable development achievements (Bianchini et al., 2023). The EDGAR database provides harmonized emission estimates across multiple pollutants by integrating data from various sources. At the NUTS 2 level, we consider carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), and fluorinated gases (F-gases), all expressed in kilotons of CO_2 equivalent. F-gases are potent greenhouse gases primarily emitted from highly polluting industrial activities such as refrigeration, electronics and semiconductor manufacturing, and some specialized chemical processes, including

[1] In Rodríguez-Pose and Bartalucci (2023), the so-called "brown" sectors encompass agriculture, forestry, and fishing (A); mining operations (B); petroleum refining (C19); chemical manufacturing (C20); non-metallic mineral production (C23); basic metals (C24); utilities including electricity and waste management (D, E); and transportation (H). Due to data constraints, we retain the NACE sector C entirely. Conversely, because Eurostat incorporates transportation in a broader category including sectors G and I, we exclude it from our analysis. However, we retain construction (NACE F) based on evidence suggesting significant future transformation potential through green job demand (Maldonado et al., 2024).

the production of magnesium and aluminum. (Crippa et al., 2023).

Following recent geopolitical developments that have highlighted both the urgency of energy transition for strategic independence and the potential macroeconomic disruptions from rapid energy mix changes, the relationship between climate performance and inflation has gained renewed attention (Boneva and Ferrucci, 2022; Moessner, 2022). Inflation dynamics are examined by incorporating the Harmonized Index of Consumer Prices (HICP) and the HICP for energy products to account for the differential impact of energy inflation compared to global inflation (Moessner, 2022). To capture regional income effects and their relationship with price changes, we include a new variable *rate_ disp*, which represents the rate of change in disposable income at the regional level.

We incorporate OECD Environmental Policy Stringency (EPS) indicators to capture national commitment regarding environmental protection. The EPS composite index includes 17 indices; for a detailed descriptio of the index structure, see Figure 2.3.2. Higher stringency scores, reflecting supportive institutional and political frameworks, strongly predict both environmental performance outcomes (Frohm et al., 2013) and eco-innovation development (Ghisetti and Pontoni, 2015; Xie et al., 2023). We apply national EPS scores to all regions within each respective country.

2.3.2 Methods

We formulate the prediction task as follows. For each region i in year t, based on a set of lagged features $\text{Features}_{i,t-1}$, we aim to determine the function $f(.)$ (a ML model) that best predicts the regional unemployment rate $\text{Unemp}_{i,t}$:

$$\text{Features}_{i,t-1} \xrightarrow{f(.)} \text{Unemp}_{i,t}. \tag{2.3.1}$$

Following established approaches in the ML literature, we randomly partition the dataset into 80% for training and 20% for out-of-sample testing. Given the limited number of regional-level observations, different data partitions may yield varying results. To address this limitation, we employ a nested cross-validation strategy, performing 10 different random splits and averaging model performance

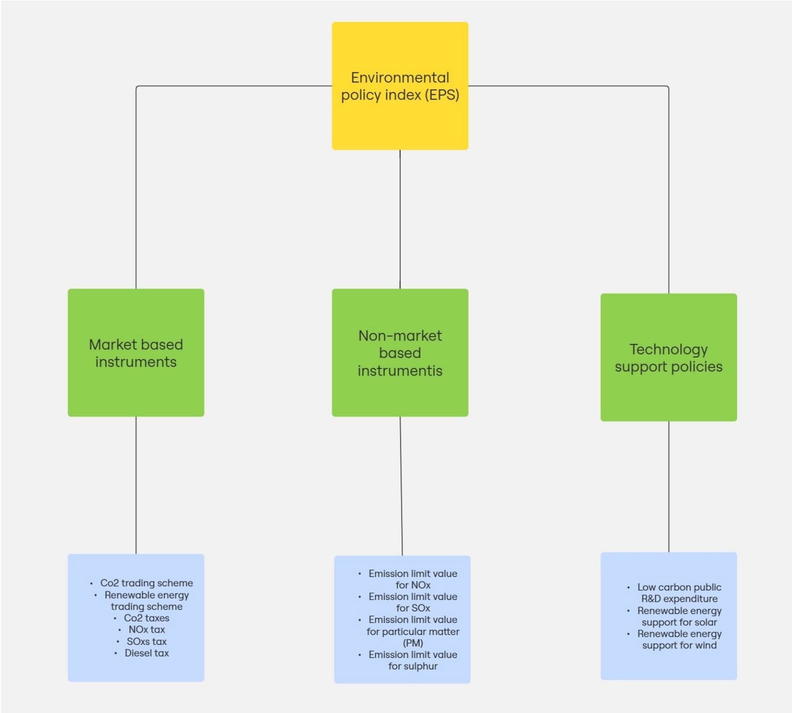

Figure 2.3.2: Structure of the OECD Environmental Policy Stringency index

Table 2.3.1: List of Variables

Source	Name	Description	Unit of Measure
Cohesion Open Data Platform	coh	Cohesion fund recipient	1= yes, 0= no
EDGAR – Emissions Database for Global Atmospheric Research	ch4	Methane emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	co2	Carbon dioxide emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	fgas	Fluorinated gases emissions	Kt CO2eq
EDGAR – Emissions Database for Global Atmospheric Research	n2o	Nitrous oxide emissions	Kt CO2eq
Environmental Policy Stringency (EPS) Index Dataset	c_co2ets	CO2 Trading Scheme	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_co2tax	Carbon dioxide (CO2) Tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_diesel	Diesel tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_enepol	Adoption of solar and wind energy support measures	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_lowrd	Low-carbon R&D expenditures	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_market	Market-based instruments	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_nonmark	Non-market based instruments	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_noxlim	Emission limit value NOx	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_noxtax	NOx taxation	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_pmlim	Emission limit value PM	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_renets	Renewable Energy Trading Scheme	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_solar	Solar Energy support (Auctions & FITs)	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sox	Emission limit value SOx	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sulflim	Emission limit value sulphur	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_sulftax	Sulphur Oxides (SOx) Tax	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_tech	Technology support policies	Integer
Environmental Policy Stringency (EPS) Index Dataset	c_wind	Wind Energy support (Auctions & FITs)	Integer
Eurostat	disp	Disposable income	Million purchasing power standards
Eurostat	egap	Gender employment gap	Percentage
Eurostat	emp	Employment in NUTS2 regions	Thousands of persons
Eurostat	gdp	Gross domestic product (GDP)	Million euros
Eurostat	htec	Employment in technology and knowledge-intensive sectors	Thousands of persons
Eurostat	lllearn	Participation rate in education and training	Percentage
Eurostat	median	Median age of the population	Years
Eurostat	nacea	Employment in Agriculture, forestry and fishing	Thousands of persons
Eurostat	nacebe	Employment in Industry (except construction)	Thousands of persons
Eurostat	nacef	Employment in Construction	Thousands of persons
Eurostat	pop	Population by NUTS 2 regions	Thousands of persons
Eurostat	ted	Tertiary educational attainment, age group 25-64	Percentage
Eurostat	unemp	Unemployment rate	Percentage
REGPAT (2024)	ecopat_count	Number of eco-innovative patent applications (EPO)	Integer
REGPAT (2024)	pat_count	Number of patent applications (EPO)	Integer
Own elaboration on Eurostat data	vuln	Employment in vulnerable sectors (NACE A to F)	Thousands of persons

Note: All variables, except those from the EPS dataset (which are at the country level), are at the NUTS 2 level. Employment data by economic activity are classified according to NACE Rev. 2. For consistency, in the EN model we convert all absolute values to relative measures using two approaches: (1) dividing the variable by regional population, or (2) dividing by the total number of employed persons in the region, in the case of sectoral employment variables. Variables that undergo this transformation will be marked with the suffix "_*rel*".

across repetitions (Resce and Vaquero-Piñeiro, 2022). Hyperparameter tuning is conducted using a repeated (10 times) five-fold cross-validation procedure.

To prevent data leakage during the training process, we consider only the lagged value (at $t-1$) for all features, including the dependent variable, as suggested by Cerqua et al. (2024). Finally, we incorporate a spatial lag for all features to account for potential regional interdependencies and agglomeration effects, consistent with the literature reviewed in Section 2.2. The spatial lag is computed following the standard approach in spatial econometrics (LeSage and Pace, 2009), which calculates the average value of each feature across neighboring regions. This operation doubles the number of features in our model.

To carry out our analysis, we employ five ML algorithms along with a classical multivariate regression model:

- **Elastic Net (EN):** A regression technique that balances L1 (LASSO) and L2 (Ridge) penalties to improve prediction accuracy and interpretability (Tibshirani, 1996; Zou and Hastie, 2005).

- **Random Forest (RF):** A tree-based ensemble learning method that randomly selects feature subsets to enhance predictive performance (Breiman, 2001).

- **Gradient Boosting Machines (GBM):** An iterative method where each learner corrects the pseudo-residuals of its predecessor (Friedman, 2001).

- **Support Vector Regression (SVR):** A method that finds a hyperplane in a high-dimensional space to minimize the error between predicted and actual values (Drucker et al., 1996).

- **Neural Networks (NNs):** A method where data are processed through hidden layers using a series of weighted connections and activation functions, whose parameters are determined iteratively, to produce a prediction (Bishop, 1995).

- **Multivariate Regression (MR):** An Ordinary Least Squares (OLS) regression approach to use as benchmark.

Performance measures are computed on the test set, with Mean Absolute Percentage Error (MAPE) as the primary evaluation metric due to its unit-independent property (Moreno et al., 2013). Our analysis, whose results are presented in Section 2.4, is carried out using the Caret package in R (Kuhn, 2008).

2.4 Results and discussion

2.4.1 Prediction performances

The comparison of model performances reveals substantial differences in predictive power. As shown in Table 2.4.1, Gradient Boosting Machine emerges as the superior algorithm, demonstrating a robust capacity to capture complex patterns in regional unemployment data. Gradient Boosting follows as the second-best performer, while more traditional models like GLM and Elastic Net show considerably weaker performance. Neural Networks, despite their theoretical capacity to model complex relationships, underperform significantly in this context, with the highest RMSE and MAPE.

Table 2.4.1: Models' performances (median across runs)

	MR	EN	**GBM**	RF	SVR	NNs
MAE	0.8301	0.8298	**0.6926**	0.7187	0.7005	7.1939
MAPE	0.1175	0.1171	0.1001	0.1011	**0.0985**	0.8337
RMSE	1.1875	1.1888	**1.0048**	1.0623	1.0523	8.8031
nRMSE	0.2386	0.2391	**0.2017**	0.2118	0.2105	1.7516
R^2	0.9430	0.9428	**0.9593**	0.9551	0.9556	−2.0730

Note: The table displays the median MAE, MAPE, RMSE, nRMSE, and R2 across the ten cross-validation runs. The negative R2 for NNs indicates that, on the test set, the neural network predictions are less accurate than a naive baseline that simply predicts the mean unemployment rate.

With a MAPE lower than 10%, Gradient Boosting achieves highly accurate unemployment predictions, demonstrating the practical value of ML approaches for this policy-relevant forecasting task (Lewis, 1982). This level of accuracy is particularly significant given the complexity of regional unemployment dynamics and the heterogeneous nature of European labor markets. The superior performance of tree-based methods (GBM and Random Forest) and of SVR over linear approaches confirms the importance of capturing non-linear relationships and variable interactions in unemployment forecasting (Athey and Imbens, 2019)—a finding that aligns with our theoretical expectation that regional unemployment patterns result from complex interdependencies among economic, demographic, sectoral, and environmental factors. The strong predictive performance of these models validates our methodological approach and suggests that ML techniques can effectively support the proactive policy design emphasized in recent literature on prediction policy problems (Kleinberg et al., 2015). The ability to forecast regional unemployment with such accuracy provides European policymakers with valuable early-warning capabilities. The excellent predictive performance of tree-based models is confirmed in settings where the lagged unemployment rate is not included among the predictors (see Appendix).

2.4.2 Feature importance

Once the good performance of our models has been established, we turn to the analysis of feature importance in the predictions. Feature importance analysis quantifies how much each predictor contributes to the model's predictive accuracy by measuring its average contribution to reducing prediction error during the training process. Although feature importance reflects predictive relevance rather than causal impact, identifying which inputs carry the most information for forecasting unemployment remains a useful exercise. The plot in Figure 2.4.1 ranks predictors according to their average contribution to error reduction across boosting iterations in the fitting procedure. Feature importance is reported for the best-performing model, which in our case is the GBM (see Table 2.4.1). The results (Figure 2.4.1) offer valuable insights into the key determinants of regional unemployment.

In line with our theoretical expectations, the lagged unemployment rate (*lag_ unemp*) constitutes the most influential predictor in the GBM model, capturing the majority of explained variance and confirming strong autocorrelation effects that reflect persistent regional unemployment performances.

The strong predictive role of the spatial lag of unemployment (*lag_ slag_ unemp*) further indicates that regional labor markets are deeply interdependent and unem-

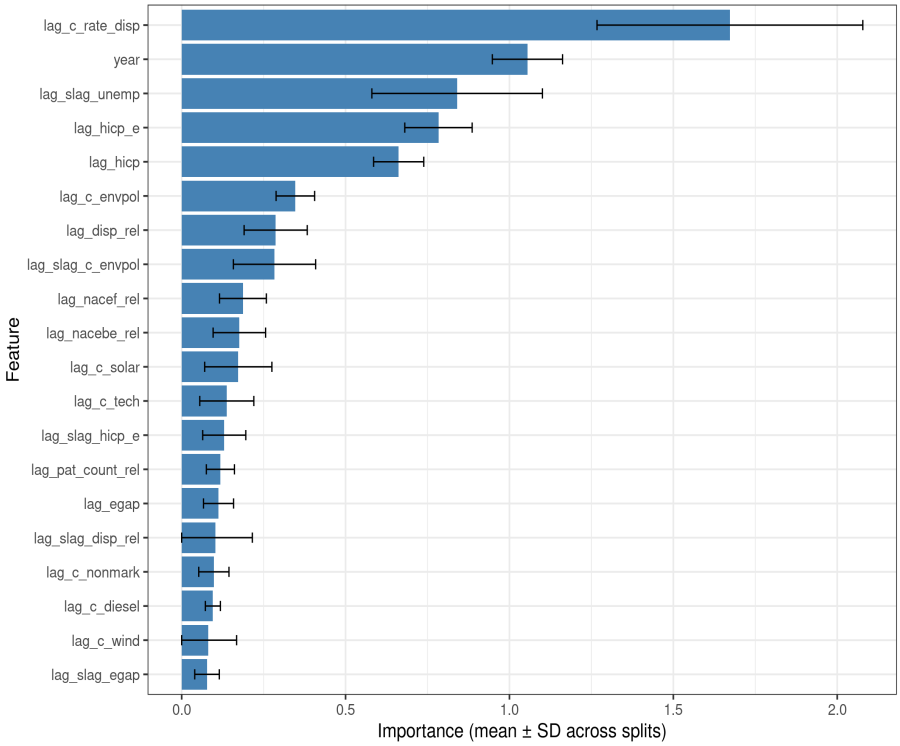

Figure 2.4.1: Feature importance plot - first 20 features by importance
Note: Feature importance is computed on both train and test set and averaged across splits. Bars indicate +− a standard deviations. The lagged unemployment rate (lag_unemp) constitutes the most influential predictor. Even so, to enhance the interpretability of other substantive predictors, we exclude this dominant variable from the visualization. See Fig. 2.A.1 in Appendix for the case where the lagged unemployed rate is not included among the predictors.

ployment outcomes are shaped not only by local structural conditions but also by spillovers from neighboring regions. This result confirms the presence of spatial clustering and agglomeration dynamics in regional unemployment. Note that spatial interdependence appears significant, with the spatial lag of disposable income (*lag_slag_disp_rel*) and of the harmonized index of consumer prices for energy products (*lag_slag_hicp_e*) and the EPS indicator (*lag_slag_envpol*) also featuring among the most important predictors. This validates our methodological choice to incorporate spatial dynamics and corroborates the literature emphasizing the geographic components of unemployment patterns (Iammarino et al., 2019).

The temporal trend (*year*) emerges as the second most important predictor, followed by the harmonized index of consumer prices for energy products (*lag_hicp_e*) and the national HICP (*lag_hicp*). Income-related measures, including disposable income (*lag_disp_rel*), its spatial lag (*lag_slag_disp_rel*), and the rate of change in disposable income (*lag_c_rate_disp*) emerge as relevant predictors, too. The importance of the time (*year*) component can be partially explained by technological progress and structural change in production systems (Acemoglu and Restrepo, 2018; Autor et al., 2003) and partially by the significant institutional and regulatory reforms affecting the European labour market over the last decades (Bentolila and Ichino, 2008; Boeri and Garibaldi, 2019). The other factors underscore the central role of income dynamics and inflationary pressures in shaping unemployment outcomes, highlighting the interconnected relevance of economic growth and price stability (two core macroeconomic objectives alongside unemployment) in shaping regional labor market dynamics. Moreover, it supports the claims in Moessner (2022) that energy-specific inflation dynamics and their role in shaping socioeconomic outcomes deserve more in-depth investigation.

Environmental policy stringency (EPS) variables constitute the next relevant block of predictors, with both the national indicator (*lag_c_envpol*) and its spatial lag (*lag_slag_c_envpol*) ranking among the top features. This result highlights that environmental policies, although decided at the national or even supranational level, yield significant effects at the local scale. Notably, the presence of the spatial lag (*i.e.*, the environmental policy stringency of neighboring regions) among the top predictors suggests the presence of relevant country effects in regional economic and labor market outcomes. Beyond the overall EPS index, policy instruments specifically related to the energy technology mix also appear to play a key role. In particular, *lag_c_tech* captures the EPS associated with technology support policies, including public R&D spending on low-carbon technologies relative to GDP and price support mechanisms for renewable energy (*e.g.*, feed-in tariffs for solar and wind). Among these, support to solar energy (*lag_c_solar*) appears particularly relevant, featuring among the top predictors, while wind-related support (*lag_c_wind*) also ranks among the top 20. Complementing these, *lag_c_diesel* and *lag_c_nonmark* — capturing EPS stringency related to diesel taxation and non-market instruments respectively — also feature among the top predictors. Together, these results highlight the centrality of regional green technology and innovation capacity, and demonstrate that both market-based regulation and technology support are crucial in shaping regional employment, in line with the findings of Costantini et al. (2018); Marin and Vona (2019); Napolitano et al. (2022).

The number of workers in the construction sector (relative to regional population) (*lag_nacef_rel*), as well as workers in the manufacturing sector (*lag_nacebe_rel*), also feature among the 20 most important predictors. This suggests that the economic structure of local economic systems, together with their vulnerabilities to the green transition, constitutes a key determinant of regional unemployment dynamics. These findings align with Iammarino et al. (2019); Rodríguez-Pose and Bartalucci (2024); Vandeplas et al. (2022), who identify reliance on labor in "brown" sectors as a key factor of vulnerability to distributive effects of climate policies. In addition, the gender employment gap (*lag_egap*) and its spatial lag also rank among the top 20 features, indicating that in regions (and even more broadly at the macro-regional level) where disparities between male and female employment are wider, labor markets tend to be structurally weaker, reflecting economic vulnerability but also persistent social inequalities.

2.4.3 Partial dependence plots

Upon identifying the most important features in predicting regional unemployment, we also investigate the functional impact of each feature on our target variable, relying on partial dependence plots (PDP). This exercise is useful to identify non-linear relationships and to empirically test whether the assumptions that are common to the literature hold when observing real-world data. Before presenting these results, an important interpretive caveat is in order. PDPs trace the marginal relationship between the outcome and each predictor by averaging over the joint distribution of all other features — an approach that, by construction, abstracts from interaction effects and correlation structure among covariates. We acknowledge this simplification, and accordingly treat the plots as exploratory rather than confirmatory evidence. Within these limits, however, PDPs provide a valuable lens through which to assess whether relationships commonly assumed in the literature find empirical support in real-world regional data. Indeed, PDPs presented in Figure 2.4.2:2.4.4 reveal complex non-linear relationships between key features and unemployment rates.

Figure 2.4.2 displays the relationships for income- and price-related variables. In the top panel, the rate of change in disposable income (*rate_ disp_ rel*) shows a positive relationship with unemployment when the rate of change is low, while unemployment remains relatively stable at higher growth rates. Regional disposable income levels (*disp_ rel*) exhibit a similar negative relationship with predicted unemployment, with the steepest marginal effect concentrated in the intermediate income range. This evidence suggests that lower unemployment rates are more likely to be observed in wealthy regions, suggesting that more stable income dynamics contribute to more resilient labor market outcomes. At the same time, a positive correlation emerges between inflation-related variables and unemployment growth. This result, confirmed for both the national HICP (*hicp*) and the HICP for energy products (*hicp_ e*) (bottom left and bottom right panel), indicates that inflationary pressures linked to general consumption prices and, more markedly, to energy costs, are associated with deteriorating labor market conditions. Rather than indicating a trade-off between inflation and unemployment, the evidence points to a scenario where rising prices erode the purchasing power of households and increase production costs for firms simultaneously, thereby depressing demand and supply. Such dynamics, while challenging the Phillips (1958) hypothesis, highlight the potential for the green transition and energy price shocks to generate stagflation-like outcomes, reinforcing the importance of policies that cushion vulnerable households and firms from the short-term inflationary costs of decarbonisation.

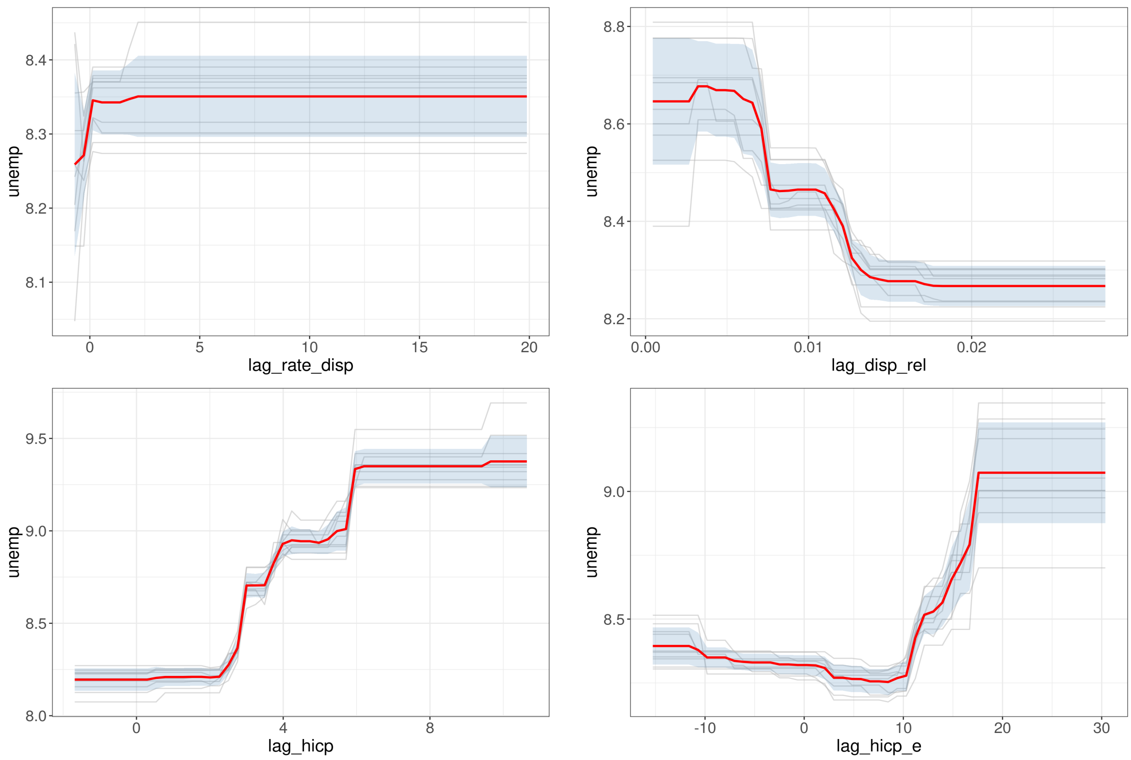

Figure 2.4.2: Partial dependence plot on variables of income and prices. *Notes:* Top Left: Rate of change in national income. Top Right: Growth rate of regional income. Bottom Left: HICP. Bottom Right: HICP for energy products. The red lines represent the averages of the grey lines, which correspond to the PDPs calculated for each algorithm. The shaded grey band spans one standard deviation above and below the mean.

Figure 2.4.3 presents partial dependence plots for the most relevant indicators of environmental policy, which reveal important threshold effects and a heterogeneous pattern across instrument types. Starting from the top left (and proceeding clockwise), the overall EPS index (*c_envpol*) shows a U-shaped relationship with unemployment, suggesting that at lower levels greater environmental policy stringency is associated with a lower unemployment rate, whereas at higher levels further increases in stringency correspond to higher unemployment. The technology support policies variable (*c_tech*) displays a pronounced step pattern, with unemployment rates rising sharply at low levels of stringency and stabilizing quickly at higher levels. Price support for solar energy (*c_ solar*) also shows an overall positive association with unemployment, with relatively lower rates at low levels of stringency and higher rates at higher levels. On the other hand, the stringency of renewable energies markets (*c_ renets*), a market-based measure, exhibits a pattern in which unemployment declines at relatively low levels of stringency and then declines. Overall, this suggests that while technology support policies are essential for fostering innovation, they do not necessarily translate into broad employment gains, as innovation activities often rely on high-skilled labor concentrated in high-tech sectors. By contrast, stronger market-based measures (such as renewable energy trading schemes) appear to generate greater benefits for the local labor market, as reflected in lower unemployment rates for moderate levels of stringency. This may be due to their fostering effect on the commercialization, adoption, and diffusion of renewable energy technologies and on the related products and services that enable their large-scale economic deployment.

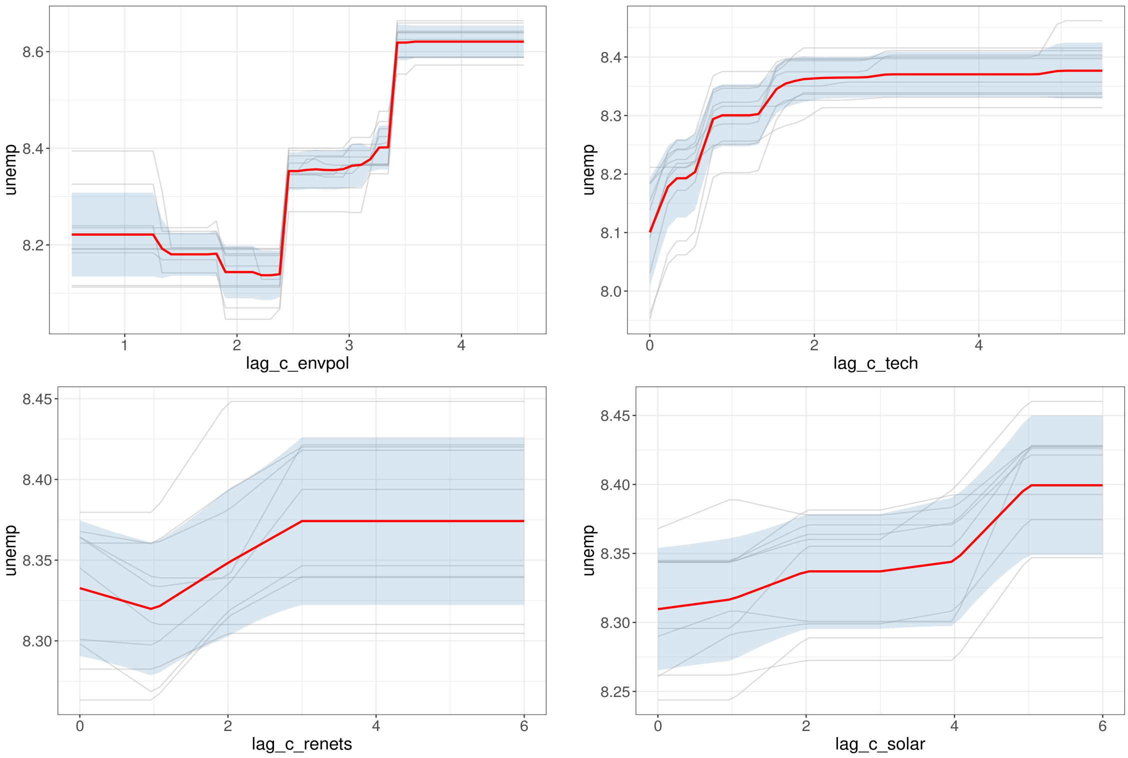

Figure 2.4.3: Partial dependence plot on selected EPS indicators. *Note:* Top Left: EPS composite index. Top Right: Technology support policies. Bottom left: renewable energy trading scheme. Bottom Right: Solar energy support. The red lines represent the averages of the grey lines, which correspond to the PDPs calculated for each algorithm. The shaded grey band spans one standard deviation above and below the mean.

Figure 2.4.4 presents partial dependence plots for CO2 emissions per capita, CO2 emission trading schemes, and sectoral employment, revealing important insights into the environment-employment nexus. Starting with CO2 emissions (Top Left), the PDP exhibits a positive relationship with unemployment that increases at low emission levels before lowering and stabilizing at high unemployment values. This trajectory suggests that regions characterised by carbon-intensive production structures, often relying on older, less energy-efficient technologies, may still be attractive in terms of labor market outcomes. Turning to the stringency of the (market-based) CO2 trading schemes (*c_co2_ets*, Top Right), the results show a non linear relationship between ETS stringency and predicted unemployment, with more stringent carbon pricing generating higher labor market pressures. Although our analysis does not allow us to establish a causal nexus, this relationship may be interpreted in light of two mechanisms that are known to the environmental policy debate. First, higher carbon prices increase production costs for regulated firms, especially in energy-intensive sectors; in the absence of rapid technological upgrading or compensatory policies, this may reduce output, compress margins, and weaken labor demand. Second, carbon pricing may accelerate the reallocation of resources from high-emission to low-emission sectors. While this process may enhance efficiency and stimulate green innovation in line with the Porter hypothesis (Porter and Linde, 1995), it may also generate short-run adjustment frictions, due job destructions and skill mismatches between displaced workers and emerging labor demand, thereby translating into higher unemployment. Together, these patterns point to a tension inherent in the green transition: carbon-intensive regions already face higher predicted unemployment, and the policy instruments designed to accelerate decarbonisation are also associated with additional labor market pressure.

The last two PDPs provide further evidence on the structural channels through which transition pressures may operate. The relative number of workers employed in the industrial sector (Bottom Left) and construction sector (Bottom Right) also exhibit distinctive patterns that reveal important structural differences across economic activities. While higher unemployment rates are associated with greater employment in construction activities,[2] employment in the industrial sector shows a complex non-monotonic relationship with unemployment. Regions with very low industrial employment tend to exhibit relatively higher unemployment rates. As industrial employment increases, unemployment initially drops sharply, likely reflecting the uptake of labor-intensive sectors, but later rises, possibly due to sectoral reallocation and adjustment frictions. At higher levels of industrial employment, however, unemployment declines and stabilises, suggesting that more industrialised labor markets are better able to absorb changes, as job creation increasingly offsets job destruction and skills gradually adjust to industrial demands. This evidence suggests that the industrial sector may operate as a resilient employment anchor with greater opportunities for integrating environmental constraints with job stability, especially if in presence of strong innovative capacity. By contrast, regions relatively more specialized in construction activities, in addition to being particularly exposed to the impacts of climate change, risk facing additional adjustment costs due to the green transition. The fact that the building sector is expected to be directly included in the EU ETS2 carbon market (alongside road transport and other activities) further highlights the need for targeted policy interventions.

Finally, further nuances appear in the PDPs of variables related to the regional innovation system (Figure 2.4.5). Patent applications (*pat_count_rel*) show a generally negative relationship with unemployment rates, though with diminishing returns at higher patent counts. The same applies to the spatial lag of patent counts (*slag_pat_count_rel*), indicating potential positive externalities coming from highly innovative neighbour regions. Employment in high-tech sectors (*htec*) also shows a negative association with unemployment, indicating that regions with a higher share of workers in high-tech sectors and knowledge-intensive activities experience lower unemployment and greater labor-market resilience, consistent with the evidence by Marin and Vona (2019). Finally, tertiary education (*ted*) displays a non-linear and largely inconclusive relationship with unemployment rates. This result is somewhat unexpected given the predictions of human capital theory, which posits that higher educational attainment enhances labour-market resilience — particularly during pe-

[2]A similar pattern is observed for employment in the agricultural sector.

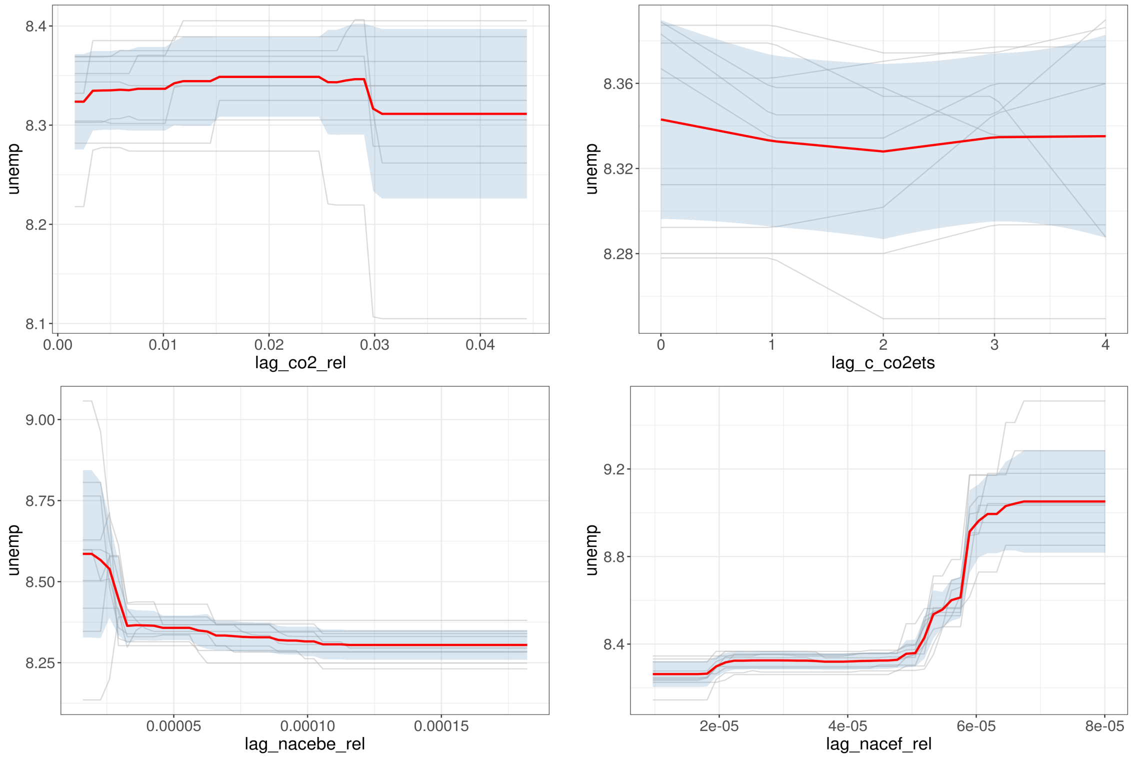

Figure 2.4.4: Partial dependence plot on variables of emissions and sectoral composition. *Notes:* Top-left: CO2 emissions. Top-right: F-gas emissions. Bottom-left: share of employment in industry relative to the number of employed persons in the region. Bottom-right: share of employment in construction relative to the number of employed persons in the region. The red lines represent the averages of the grey lines, which correspond to the PDPs calculated for each algorithm. The shaded grey band spans one standard deviation above and below the mean.

riods of economic and structural transition — and with Consoli et al. (2016); Vandeplas et al. (2022), who emphasise the role of skills in differentiating green from non-green employment. A plausible interpretation, however, is that the relevance of this indicator is partially absorbed by other regional characteristics, such as the availability of high-tech employment and advanced knowledge infrastructure, which appear comparatively more decisive in shaping unemployment outcomes.

Taken together, the patterns emerging from innovation capacity and technology-related policy variables suggest the possibility of heterogeneous development trajectories – regions endowed with stronger innovation capacity, higher shares of high-tech and knowledge-intensive employment, and a more educated workforce appear to be better equipped to benefit from environmental and technological transitions, as reflected by lower unemployment rates. In contrast, regions with weaker human capital endowments and less advanced sectoral structures are less able to capture the employment gains associated with innovation and may thus face greater adjustment costs.

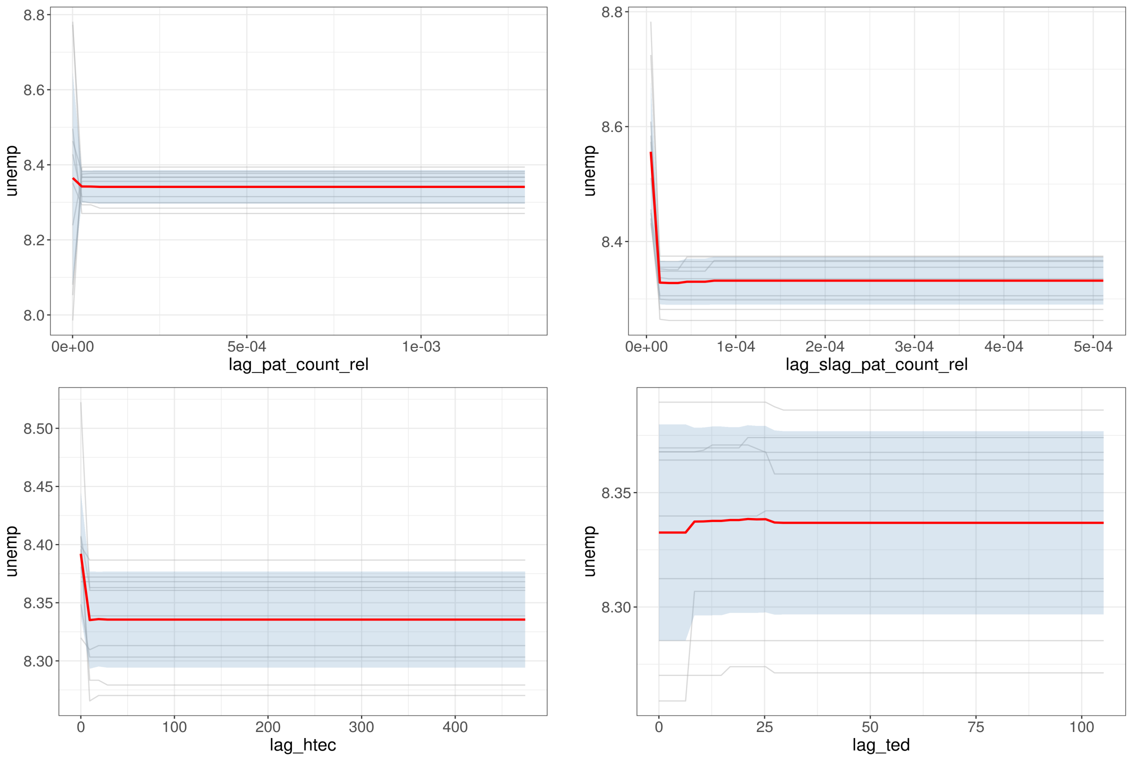

Figure 2.4.5: Partial dependence plots on variables of innovation. *Note:* Top Left: number of patent applications. Top Right: spatial lag of number of patent applications. Bottom Left: percentage workforce employed in technology and knowledge intensive sectors. Bottom Right: tertiary education attainment.The red lines represent the averages of the grey lines, which correspond to the PDPs calculated for each algorithm. The shaded grey band spans one standard deviation above and below the mean.

2.5 Conclusions and policy implications

This paper set out to predict regional unemployment patterns in the EU by leveraging ML methods and a high-dimensional panel of socioeconomic, environmental, and policy-related variables. Our results demonstrate that tree-based models, particularly Gradient Boosting and Random Forest, substantially outperform traditional linear approaches, achieving remarkably accurate forecasts of unemployment dynamics at the NUTS-2 level. This confirms the usefulness of ML techniques for capturing the non-linearities, spatial dependencies, and complex interactions that shape labour market outcomes. The superior performance of tree-based methods over linear approaches confirms the importance of capturing non-linear relationships and variable interactions in unemployment forecasting, aligning with recent applications of ML in economic policy analysis.

By providing highly accurate and spatially granular forecasts, our model offers policymakers an effective early-warning tool to anticipate unemployment risks across regions. Such predictive capacity can support the timely design of place-based labour market interventions, guide the allocation of cohesion funds, and help mitigate the uneven social costs of the green transition through proactive rather than reactive policy responses. Building on the excellent predictive performance of tree-based models, we take a further step by leveraging the model to generate feature importance analysis and partial dependence plots. These analyses, while they should not be interpreted in a causal sense, offer a window into the mechanics of the algorithm and can inspire substantive economic insight. Which variables drive regional unemployment outcomes, and how do they relate to the target variable? We find that the patterns uncovered are broadly consistent with theoretical expectations and established empirical findings, lending credibility to the model's internal logic. At the same time, these results surface associations that warrant more in-depth study, and we hope they serve as a useful starting point for future research seeking to establish rigorous causal relationships.

As a matter of fact, our results contribute to several streams of literature by providing data-driven insights on key determinants of regional unemployment. First, we identify a strong link between inflation and unemployment, consistent with decades of economic research. Energy-related inflation appears particularly relevant, underscoring the urgency of considering the direct labor market effects of energy supply shocks. This finding carries important implications as concerns about energy security and affordability emerge in relation to the transition away from fossil fuels. Second, the high predictive importance of spatial lag variables confirms the relevance of agglomeration effects and regional spillovers emphasized in the literature. This spatial interdependence has important implications for European cohesion policy, suggesting that regional interventions must account for cross-border spillover effects and the interconnected nature of regional labor markets. Third, the high predictive importance of environmental policy variables validates concerns raised in the green transition literature about the heterogeneous labor market impacts of climate policies.

These results may stand in sharp tension with one another in practice. Yet, our findings serve as a warning on two fronts: first, we reiterate that uniform policies generate heterogeneous outcomes across regions with different productive structures and labour market conditions; and second, that effective policy design should be explicitly place-sensitive — not merely acknowledging spatial transmission channels but actively exploiting them, for instance through interregional coordination mechanisms that internalise cross-border spillovers.

Our partial dependence analysis reveals several patterns that merit deeper investigation. Most notably, the relationship between inflation metrics and unemployment contrasts with the traditional Phillips curve hypothesis. The positive correlation between inflation and unemployment growth challenges the conventional inverse relationship, echoing recent critiques regarding the stability and regional applicability of the Phillips curve. This finding is particularly relevant given renewed interest in inflation-unemployment trade-offs following recent inflationary episodes and energy price volatility. These findings carry particular relevance in the context of the EU's energy and climate policies. They underscore that energy price stability is a deeply social concern and that an increase in price pressures, such as the one that could result from a shift to expensive renewables, risks exacerbating unemployment. At the same time, recent experience — most notably the energy crisis triggered by the war in Ukraine — has demonstrated that dependence on fossil fuels offers no guarantee of price stability either, and may in fact amplify exposure to shocks. This being the case, a more diversified energy mix could be the opportunity to prioritise affordability alongside decarbonisation objectives. The relationship with emission variables also warrants further investigation. The pattern observed for CO2 emissions suggests that regions characterized by more carbon-intensive structures tend to exhibit lower unemployment levels. On the other hand, a stricter CO2 ETS regulation yields the opposite result, suggesting that a greater exposure to transition pressures may lead to labour disruptions in areas characterized by of carbon-intensive activities. Environmental policy stringency indicators reveal that different policies produce varying employment effects. Technological support measures, in particular, may not generate the employment gains associated with other market-based instruments. Overall, our results confirm that regions with stronger human capital endowments and environmental protection frameworks demonstrate greater resilience to unemployment shocks. This could raise concerns of equity as the green transition unfolds.

This study exemplifies how ML approaches can advance policy-relevant economic research while providing an ideal testing ground for economic theories. The policy implications are substantial: ML methods can predict unemployment rates with high accuracy, providing policymakers with powerful tools for targeted interventions. As European regions navigate the challenges of green transition and macroeconomic instability, such approaches can provide the granular, predictive insights necessary for effective and spatially-targeted policy interventions, enabling proactive rather than reactive policy responses in regions at risk of rising unemployment.

Bibliography

D. Acemoglu and P. Restrepo. Artificial intelligence, automation, and work. In *The economics of artificial intelligence: An agenda*, pages 197–236. University of Chicago Press, 2018.

M. Andini, E. Ciani, G. de Blasio, A. D'Ignazio, and V. Salvestrini. Targeting with machine learning: An application to a tax rebate program in italy. *Journal of Economic Behavior & Organization*, 156:86–102, 2018.

N. Antulov-Fantulin, R. Lagravinese, and G. Resce. Predicting bankruptcy of local government: A machine learning approach. *Journal of Economic Behavior & Organization*, 183:681–699, 2021.

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

D. H. Autor, F. Levy, and R. J. Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4): 1279–1333, 2003.

N. R. Babb and A. K. Detmeister. Has the us phillips curve flattened? a semi-structural exploration. *Finance and Economics Discussion Series 2017-065*, 2017.

F. C. Bagliano and C. Morana. Eurozone economic integration: Historical developments and new challenges ahead. *European Economic Review*, page 105023, 2025.

N. Barbieri, D. Consoli, L. Napolitano, F. Perruchas, E. Pugliese, and A. Sbardella. Regional technological capabilities and green opportunities in europe. *Journal of Technology Transfer*, 48:749–778, 4 2023. ISSN 15737047.

R. Basile. Productivity polarization across regions in europe: The role of nonlinearities and spatial dependence. *International Regional Science Review*, 32(1): 92–115, 2009.

S. Bentolila and A. Ichino. Unemployment and consumption near and far away from the mediterranean. *Journal of Population Economics*, 21(2):255–280, 2008.

S. Bianchini, G. Damioli, and C. Ghisetti. The environmental effects of the "twin" green and digital transition in european regions. *Environmental and Resource Economics*, 84:877–918, 4 2023. ISSN 15731502.

C. M. Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995.

T. Boeri and P. Garibaldi. A tale of comprehensive labor market reforms: evidence from the italian jobs act. *Labour Economics*, 59:33–48, 2019.

T. Boeri and K. Terrell. Institutional determinants of labor reallocation in transition. *Journal of economic perspectives*, 16(1):51–76, 2002.

L. Boneva and G. Ferrucci. Inflation and climate change: the role of climate variables in inflation forecasting and macro modelling. 4 2022.

K. Borsekova and S. Korony. Resilience and vulnerability of regional labour markets: principal component analysis of labour market efficiency in the eu. *Regional Studies*, 57(12):2373–2390, 2023.

L. Breiman. Random forests. *Machine Learning*, 45:5–32, 10 2001. ISSN 08856125.

F. Cappellano, A. M. Santos, and N. F. Dotti. Regional ri ventures to tackle climate change: A new geography of challenge-oriented innovation landscape. *Papers in Regional Science*, 103:100052, 10 2024. ISSN 1056-8190.

N. Caravaggio, G. Resce, and C. Vaquero-Piñeiro. Predicting policy funding allocation with machine learning. *Socio-Economic Planning Sciences*, 98:102175, 2025.

P. D. Caro and U. Fratesi. The role of cohesion policy for sustaining the resilience of european regional labour markets during different crises. *Regional Studies*, 57 (12):2426–2442, 2023.

V. Carrieri, R. Lagravinese, and G. Resce. Predicting vaccine hesitancy from area-level indicators: A machine learning approach. *Health Economics*, 30(12):3248–3256, 2021.

K. A. Castellanos and G. Heutel. Unemployment, labor mobility, and climate policy. Technical report, National Bureau of Economic Research, 2019.

A. Cerqua, M. Letta, and G. Pinto. On the (mis) use of machine learning with panel data. *arXiv preprint arXiv:2411.09218*, 2024.

G. Cerulli. Fundamentals of supervised machine learning. *Cham: Springer*, 2023.

P. Christensen, P. Francisco, E. Myers, H. Shao, and M. Souza. Energy efficiency can deliver for climate policy: Evidence from machine learning-based targeting. *Journal of Public Economics*, 234:105098, 2024.

L. Coenen, P. Benneworth, and B. Truffer. Toward a spatial perspective on sustainability transitions. *Research policy*, 41(6):968–979, 2012.

Cohesion Open Data Platform. Cohesion open data platform. URL `https://cohesiondata.ec.europa.eu/`. Accessed: September 2025.

D. Consoli, G. Marin, A. Marzucchi, and F. Vona. Do green jobs differ from non-green jobs in terms of skills and human capital? *Research Policy*, 45(5):1046–1060, 2016.

C. Corradini and E. Vanino. Path dependency, regional variety and the dynamics of new firm creation in rooted and pioneering industries. *Journal of Economic Geography*, 22(3):631–651, 2022.

V. Costantini, F. Crespi, and E. Paglialunga. The employment impact of private and public actions for energy efficiency: Evidence from european industries. *Energy Policy*, 119:250–267, 2018.

M. F. Cracolici, M. Cuffaro, and P. Nijkamp. Geographical distribution of unemployment: An analysis of provincial differences in italy. *Growth and Change*, 38 (4):649–670, 2007.

M. Crippa, D. Guizzardi, F. Pagani, M. Banja, M. Muntean, E. Schaaf, W. Becker, F. Monforti-Ferrario, R. Quadrelli, A. Risquez Martin, et al. Ghg emissions of all world countries. *Publications Office of the European Union, Luxembourg*, 10: 953322, 2023.

R. Di Stefano and G. Resce. The determinants of missed funding: Predicting the paradox of increased need and reduced allocation. *Journal of Economic Behavior & Organization*, 231:106910, 2025.

A. Diemer, S. Iammarino, A. Rodríguez-Pose, and M. Storper. The regional development trap in europe. *Economic Geography*, 98:487–509, 10 2022. ISSN 19448287.

L. Dijkstra, E. Garcilazo, and P. McCann. The effects of the global financial crisis on european regions and cities. *Journal of Economic Geography*, 15(5):935–949, 2015.

H. Drucker, C. J. Burges, L. Kaufman, A. Smola, and V. Vapnik. Support vector regression machines. *Advances in Neural Information Processing Systems*, 9, 1996.

E. A. D'Ecclesiis, E. Levi, and F. Patriarca. Exploring the multifaceted relationship between environmental attitudes and political voting. *Ecological Economics*, 230: 108518, 2025.

EDGAR – Emissions Database for Global Atmospheric Research. Edgar – emissions database for global atmospheric research. URL `https://edgar.jrc.ec.europa.eu/`. Accessed: January 2024.

Environmental Policy Stringency (EPS) Index Dataset. Environmental policy stringency (eps) index dataset. URL `https://www.oecd.org/environment/indicators-modelling-outlooks/environmental-policy-stringency.htm`. Accessed: January 2024.

Eurostat. Eurostat regional statistics database. URL `https://ec.europa.eu/eurostat/data/database`. Accessed: January 2024.

N. Fabra, A. Lacuesta, and M. Souza. The implicit cost of carbon abatement during the covid-19 pandemic. *European Economic Review*, 147:104165, 2022.

J. H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001. ISSN 00905364, 21688966. URL `http://www.jstor.org/stable/2699986`.

M. Friedman. The role of monetary policy. *The American economic review*, 58(1): 1–17, 1968.

E. Frohm, F. M. D'arcangelo, T. Kruse, and M. Pisu. Environmental policy stringency and co2 emissions: Evidence from cross-country sector-level data. 2013.

C. Ghisetti and F. Pontoni. Investigating policy and rd effects on environmental innovation: A meta-analysis. *Ecological Economics*, 118:57–66, 10 2015. ISSN 0921-8009.

S. Girma and D. Paton. Using double-debiased machine learning to estimate the impact of covid-19 vaccination on mortality and staff absences in elderly care homes. *European Economic Review*, 170:104882, 2024.

P. Gogas, T. Papadimitriou, and E. Sofianos. Forecasting unemployment in the euro area with machine learning. *Journal of Forecasting*, 41(3):551–566, 2022.

R. J. Gordon. Friedman and phelps on the phillips curve viewed from a half century's perspective. *Review of Keynesian Economics*, 6(4):425–436, 2018.

M. Greenstone, R. Hornbeck, and E. Moretti. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of political economy*, 118(3):536–598, 2010.

J. Hazell, J. Herreño, E. Nakamura, and J. Steinsson. The slope of the phillips curve: evidence from us states. *The Quarterly Journal of Economics*, 137(3):1299–1344, 2022.

P. Hooper, F. S. Mishkin, and A. Sufi. Prospects for inflation in a high pressure economy: Is the phillips curve dead or is it just hibernating? *Research in Economics*, 74(1):26–62, 2020.

S. Iammarino, A. Rodriguez-Pose, and M. Storper. Regional inequality in europe: evidence, theory and policy implications. *Journal of Economic Geography*, 19: 273–298, 3 2019. ISSN 1468-2702.

N. Jean, M. Burke, M. Xie, W. M. Alampay Davis, D. B. Lobell, and S. Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.

O. W. Johnson, J. Y. C. Han, A. L. Knight, S. Mortensen, M. T. Aung, M. Boyland, and B. P. Resurrección. Intersectionality and energy transitions: A review of gender, social equity and low-carbon energy. *Energy Research Social Science*, 70: 101774, 12 2020. ISSN 2214-6296.

C. Katris. Prediction of unemployment rates with time series and machine learning techniques. *Computational Economics*, 55(2):673–706, 2020.

J. Kleinberg, J. Ludwig, S. Mullainathan, and Z. Obermeyer. Prediction policy problems. *American Economic Review*, 105(5):491–495, 2015.

P. Kline and E. Moretti. Place based policies with unemployment. *American Economic Review*, 103(3):238–243, 2013.

P. Kline and E. Moretti. Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly journal of economics*, 129(1):275–331, 2014.

A. Kronsell. Gender and transition in climate governance. *Environmental Innovation and Societal Transitions*, 7:1–15, 6 2013. ISSN 2210-4224.

M. Kuhn. Building predictive models in R using the caret package. *Journal of Statistical Software*, 28(5):1–26, 2008. doi: 10.18637/jss.v028.i05.

J. LeSage and R. K. Pace. *Introduction to spatial econometrics*. Chapman and Hall/CRC, 2009.

C. D. Lewis. International and business forecasting methods, 1982.

J. E. Maldonado, A. Vandeplas, I. V. Vigani, and A. Turrini. Assessing green job dynamics in the eu: A comparison of alternative methodologies, 7 2024.

G. Marin and F. Vona. Climate policies and skill-biased employment dynamics: Evidence from eu countries. *Journal of Environmental Economics and Management*, 98:102253, 2019.

R. Moessner. Evidence on climate policy, carbon dioxide emissions and inflation. *International Journal of Global Warming*, 28:136–151, 2022. ISSN 17582091.

J. J. M. Moreno, A. P. Pol, A. S. Abad, and B. C. Blasco. Using the r-mape index as a resistant measure of forecast accuracy. *Psicothema*, 25(4):500–506, 2013.

E. Moretti. Local labor markets. *Handbook of Labor Economics*, 4B:1237–1313, 2011a.

E. Moretti. Local multipliers. *American Economic Review: Papers and Proceedings*, 100(2):1–7, 2011b.

E. Moretti. *The new geography of jobs*. Houghton Mifflin Harcourt, 2012.

S. Mullainathan and J. Spiess. Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106, 2017.

L. Napolitano, A. Sbardella, D. Consoli, N. Barbieri, and F. Perruchas. Green innovation and income inequality: A complex system analysis. *Structural Change and Economic Dynamics*, 63:224–240, 2022.

A. Niebuhr. Spatial interaction and regional unemployment in europe. *European Journal of Spatial Development*, 1(5):26–26, 2003.

OECD. The jobs potential of a shift towards a low-carbon economy: Final report for the eu commission, dg employment, 2012.

H. G. Overman and D. Puga. Unemployment clusters across europe's regions and countries. *Economic policy*, 17(34):115–148, 2002.

E. S. Phelps. Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, pages 254–281, 1967.

A. W. Phillips. The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957. *Economica*, 25(100):283–299, 1958.

M. E. Porter and C. v. d. Linde. Toward a new conception of the environment-competitiveness relationship. *Journal of economic perspectives*, 9(4):97–118, 1995.

REGPAT. Regpat database – oecd regional patent database, 2024. URL `https://www.oecd.org/sti/inno/regpat-database.htm`. Accessed: January 2024.

G. Resce and C. Vaquero-Piñeiro. Predicting agri-food quality across space: A machine learning model for the acknowledgment of geographical indications. *Food Policy*, 112:102345, 2022.

A. Rodríguez-Pose. Do institutions matter for regional development? *Regional studies*, 47(7):1034–1047, 2013.

A. Rodríguez-Pose and F. Bartalucci. The green transition and its potential territorial discontents. *Cambridge Journal of Regions, Economy and Society*, 17(2): 339–358, 2024.

A. Rodríguez-Pose and F. Bartalucci. The green transition and its potential territorial discontents. *Cambridge Journal of Regions, Economy and Society*, 17(2): 339–358, 11 2023. ISSN 1752-1378.

P. A. Samuelson and R. M. Solow. Analytical aspects of anti-inflation policy. *The American economic review*, 50(2):177–194, 1960.

D. Sansone. Beyond early warning indicators: high school dropout and machine learning. *Oxford bulletin of economics and statistics*, 81(2):456–485, 2019.

D. Sansone and A. Zhu. Using machine learning to create an early warning system for welfare recipients. *Oxford Bulletin of Economics and Statistics*, 85(5):959–992, 2023.

C. L. Smith. The phillips curve: A relation between real activity and inflation? *Journal of Macroeconomics*, 79:103543, 2024.

J. H. Stock and M. W. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American statistical association*, 97 (460):1167–1179, 2002.

J. H. Stock and M. W. Watson. Phillips curve inflation forecasts. *Journal of Monetary Economics*, 44:99–116, 2008.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58:267–288, 1 1996. ISSN 1369-7412.

V. Titl, D. Mazrekaj, and F. Schiltz. Identifying politically connected firms: A machine learning approach. *Oxford Bulletin of Economics and Statistics*, 86(1): 137–155, 2024.

A. Vandeplas, I. Vanyolos, M. Vigani, L. Vogel, et al. *The possible implications of the green transition for the EU labour market*. Publications Office of the European Union, 2022.

F. Vona. Managing the distributional effects of environmental and climate policies: The narrow path for a triple dividend, 2021.

F. Vona. Managing the distributional effects of climate policies: A narrow path to a just transition. *Ecological Economics*, 205:107689, 2023. ISSN 0921-8009.

P. Xie, Y. Xu, X. Tan, and Q. Tan. How does environmental policy stringency influence green innovation for environmental managements? *Journal of Environmental Management*, 338:117766, 7 2023. ISSN 0301-4797.

B. Zaman and S. Borsky. The impact of supply structure on solar home system installations in rural off-grid areas. *Environmental Innovation and Societal Transitions*, 40:625–644, 2021.

H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67:301–320, 4 2005. ISSN 1369-7412.

2.A Appendix

Features: descriptive statistics

In Table 2.A.1 we present the descriptive statistics for the features included in our analysis. Note that these are computed over the whole panel 2000 to 2019.

Table 2.A.1: Descriptive Statistics

Variable	Mean	Median	SD	Min	Max
coh	0.834	1.000	0.372	0.000	1.000
ch4_rel	0.00144	0.00122	0.00106	0.000110	0.00931
co2_rel	0.00869	0.00721	0.00505	0.00144	0.04440
fgas_rel	0.000224	0.000191	0.000166	0.0000245	0.004683
n2o_rel	0.000609	0.000458	0.000482	0.0000333	0.003496
c_co2ets	1.426	1.000	1.253	0.000	4.000
c_co2tax	0.704	0.000	1.569	0.000	6.000
c_diesel	3.936	4.000	0.813	0.000	6.000
c_envpol	2.799	2.889	0.738	0.528	4.722
c_lowrd	2.319	2.000	1.469	0.000	6.000
c_market	1.434	1.167	0.804	0.333	4.167
c_nonmark	4.744	5.250	1.059	1.250	6.000
c_noxlim	4.377	5.000	1.461	0.000	6.000
c_noxtax	1.278	0.000	1.931	0.000	6.000
c_pmlim	4.376	6.000	1.950	0.000	6.000
c_renets	0.548	0.000	1.402	0.000	6.000
c_solar	2.066	2.000	1.896	0.000	6.000
c_sox	4.673	5.000	0.823	0.000	6.000
c_sulflim	5.549	6.000	0.500	4.000	6.000
c_sulftax	0.709	0.000	1.320	0.000	6.000
c_tech	2.221	2.250	1.210	0.000	6.000
c_wind	2.180	2.000	1.797	0.000	6.000
disp_rel	0.01449	0.01463	0.00424	0.000445	0.03008
egap	13.887	12.700	8.170	0.000	104.200
emp_rel	0.000426	0.000430	0.0000547	0.000155	0.000539
gdp_rel	0.02568	0.02539	0.01238	0.003165	0.10160
htec_rel	0.03660	0.03385	0.01996	0.000	0.13232
lllearn	9.428	7.500	6.597	0.000	36.000
median	41.377	41.200	3.286	28.000	51.700
nacea_rel	0.05368	0.03395	0.05835	0.000883	0.37501
nacebe_rel	0.19333	0.18673	0.07216	0.032598	0.38995
nacef_rel	0.07571	0.07250	0.02219	0.020574	0.32706
pop	1,954,886	1,480,048	1,635,841	118,861	12,252,917
ted	24.214	24.200	11.037	0.000	105.200
unemp	8.874	7.600	5.388	1.200	36.300
ecopat_count_rel	0.0000172	0.00000741	0.0000250	0.000	0.000239
pat_count_rel	0.000132	0.0000726	0.000171	0.000	0.001298
vuln_rel	0.3227	0.3166	0.09411	0.10152	0.58843
Spatial lags					
spatial_lag_unemp	8.879	8.115	3.689	3.653	21.300
spatial_lag_gdp_rel	0.02534	0.02701	0.009344	0.006668	0.04536
spatial_lag_disp_rel	0.01444	0.01515	0.002853	0.006252	0.01960
spatial_lag_egap	13.962	13.183	5.957	3.142	32.563

Notes: Statistics computed over the full panel (2000–2019).

2.A.1 ML algorithms and loss functions

All predictive algorithms considered in this study can be framed as different responses to the same prediction task. Let $\{(x_i, y_i)\}_{i=1}^n$ denote a sample of observations where y_i is the response variable and x_i is a vector of predictors.The goal of a predictive model is to estimate a function $f(x)$ that explains the relationship between the predictors and the outcome:

$$y_i = f(x_i) + \varepsilon_i$$

where ε_i is an error term with mean zero. In very general terms, ML algorithms estimate $f(\cdot)$ by minimizing an empirical loss function:

$$\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} L(y_i, f(x_i)) + \Omega(f)$$

where $L(y_i, f(x_i))$ is a loss function measuring the distance between observed and predicted values, $\Omega(f)$ is a regularization term that penalizes model complexity and avoids overfitting, $\mathcal{F}$ denotes the group of admissible prediction functions. For regression problems, the most commonly used loss function is the squared error loss:

$$L(y_i, f(x_i)) = (y_i - f(x_i))^2$$

Different ML algorithms differ in the functional form of $f(x)$ and the regularization structure imposed through $\Omega(f)$. This section provides an overview of the algortihms used in the analysis. For each ML algorithm, the minimization task is described [3].

Elastic Net The Elastic Net (EN) is a regularized regression technique that combines the penalties of the Least Absolute Shrinkage and Selection Operator (LASSO) and Ridge regression (Tibshirani, 1996; Zou and Hastie, 2005). LASSO applies an L_1 penalty that can shrink some coefficients exactly to zero, thereby performing variable selection. In contrast, Ridge regression uses an L_2 penalty that shrinks coefficients

[3] For an accessible, although rigorous introduction to supervised ML algorithms, we refer to Cerulli (2023).

toward zero and helps stabilize estimates in the presence of multicollinearity. Elastic Net combines these two penalties, making it particularly useful when predictors are highly correlated.

Elastic Net estimates the coefficient vector β by minimizing the following penalized least squares objective function:

$$\min_{\beta} \frac{1}{2n} \sum_{i=1}^{n} (y_i - x_i^\top \beta)^2 + \lambda \left[(1-\alpha) \frac{||\beta||_2^2}{2} + \alpha ||\beta||_1 \right]$$

where $||\beta||_1 = \sum_j |\beta_j|$ denotes the L_1 penalty and $||\beta||_2^2 = \sum_j \beta_j^2$ denotes the L_2 penalty. The hyperparameter λ controls the overall strength of regularization, while $\alpha \in [0, 1]$ determines the balance between the two penalties. When $\alpha = 1$ the model reduces to LASSO, while $\alpha = 0$ corresponds to Ridge regression. The optimal values of λ and α are typically determined through cross-validation.

Random Forest Random Forest (RF) is an ensemble learning method that combines predictions from multiple decision trees constructed in parallel on bootstrap samples of the training data (Breiman, 2001). Each tree is built independently, and at every split only a random subset of predictors is considered, a passage that reduces correlation among trees and improves predictive performance. For regression problems, each decision tree recursively partitions the feature space into independent regions using a procedure called recursive binary splitting. At each node of the tree, the algorithm selects a predictor X_j and a split point s that divide the observations into two subsets. The chosen split is the one that minimizes the mean squared error (MSE) within the resulting nodes:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2.$$

In single-tree models, this process continues until a stopping criterion is reached (e.g., minimum node size or maximum tree depth), in order to prevent overfitting. In Random Forests, instead, trees are typically grown fully, as the risk of overfitting is addressed by the very aggregation of many independent trees. Once all trees are grown, the Random Forest prediction is obtained by averaging the predictions of the B individual trees:

$$\hat{y}(x) = \frac{1}{B}\sum_{b=1}^{B} T_b(x),$$

where $T_b(x)$ denotes the prediction of the b-th tree.

Gradient Boosting Machines Gradient Boosting Machines (GBM) are an ensemble learning method that constructs a predictive model as an additive expansion of weak learners (in the ML literature, simple predictive models that perform slightly better than random guessing), where each iteration aims to correct the errors of the previous model (Friedman, 2001). The model can be expressed as:

$$F(x) = \sum_{b=1}^{B} \beta_b g(x, \phi_b),$$

where $g(x, \phi_b)$ denotes the b-th basis function (in our case, a regression tree) with parameters ϕ_b, and β_b its associated coefficient. The parameters are estimated by minimizing the empirical loss function (in general, the squared error loss). To solve this optimization problem, GBM adopts a sequential procedure that adds new learners to the model. At iteration m, the model is updated as

$$F_m(x) = F_{m-1}(x) + \nu h_m(x),$$

where $h_m(x)$ is a weak learner fitted to the pseudo-residuals (*i.e.*, the negative gradient of the loss function with respect to the current model predictions, which represent the direction in which the model should adjust its predictions to reduce the loss) of the previous iteration and ν is a learning rate controlling the contribution of each new component. The pseudo-residuals correspond to the negative gradient of the loss function:

$$r_{im} = -\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)}.$$

Support Vector Regression Support Vector Regression (SVR) extends the principles of Support Vector Machines to regression problems (Drucker et al., 1996). The method aims to estimate a function that approximates the relationship between predictors and the response while maintaining the model as flat as possible.Unlike classification SVMs, SVR does not aim to separate observations into groups, but instead fits a function such that prediction errors remain within a tolerance margin ϵ whenever possible.

The observations that lie on or outside this margin are called *support vectors*, as they determine the position of the regression function. SVR approximates the target values with an ϵ-insensitive "tube" around the predicted function while keeping model complexity small. The optimization problem can be written as:

$$\min_{w,b,\xi_i,\xi_i^*} \frac{1}{2}||w||^2 + C\sum_{i=1}^{n}(\xi_i + \xi_i^*)$$

Subject to

$$y_i - (w^\top \phi(x_i) + b) \leq \epsilon + \xi_i$$
$$(w^\top \phi(x_i) + b) - y_i \leq \epsilon + \xi_i^*$$
$$\xi_i, \xi_i^* \geq 0$$

Where: $\phi(x)$ maps the input into a higher-dimensional feature space, ϵ defines the width of the tube, ξ_i, ξ_i^* are variables penalizing deviations from the tube, C controls the trade-off between model complexity and prediction error. This formulation corresponds to minimizing an ϵ-insensitive loss function

$$L(y_i, f(x_i)) = \max(0, |y_i - f(x_i)| - \epsilon),$$

which ignores errors smaller than ϵ and penalizes only larger deviations.

Neural Networks (NNs) Neural networks consist of layers of interconnected nodes (neurons) that transform input variables through weighted linear combinations followed by nonlinear activation functions (Bishop, 1995). Data are fed into the input layer, processed through hidden layers using weighted connections and activation functions, and produce predictions at the output layer. The model adjusts these weights iteratively during training through the backpropagation algorithm. In this study, we employ a feedforward neural network with parameters θ, which can be represented as a system of nested transformations. With one hidden layer, the model can be expressed as

$$Z = \sigma(\alpha_0 + \alpha X)$$

$$\hat{y} = \beta_0 + \beta Z$$

where Z represents the vector of hidden-layer activations (neurons), $\sigma(\cdot)$ is a nonlinear activation function, and α and β are weight matrices connecting the layers. The parameters are estimated by minimizing a loss function via gradient-based optimization methods. In regression problems, this is typically the squared error loss $L(y_i, f(x_i)) = (y_i - f(x_i))^2$.

2.A.2 Prediction performance without lagged unemployment rate

To assess the robustness of our predictive framework, we replicate the fitting and prediction procedure described in Section 2.3, this time excluding the lagged unemployment rate from the set of predictors.

Table 2.A.2 reports the prediction performance across models. The results demonstrate considerable robustness to the exclusion of the lagged dependent variable. RF achieves the lowest mean absolute percentage error (MAPE), closely followed by SVR and GBM.

Table 2.A.2: Models' performances for model without lagged unemployment (median across runs)

	MR	EN	GBM	**RF**	SVR	NNs
MAE	1.7850	1.8026	1.2287	**0.8117**	0.9749	7.3359
MAPE	0.2650	0.2669	0.1786	**0.1148**	0.1306	0.8357
RMSE	2.4077	2.4170	1.6480	**1.1933**	1.4779	9.0479
nRMSE	0.4616	0.4637	0.3156	**0.2298**	0.2881	1.7287
R^2	0.7866	0.7846	0.9002	**0.9471**	0.9168	−1.9930

Note: The table displays the median MAE, MAPE, RMSE, nRMSE, and R2 across the ten cross-validation runs when excluding the lagged unemployment rate from the set of predictors.

Figure 2.A.1 displays the feature importance rankings for the Random Forest model estimated without the lagged unemployment rate. The ordering of covariates remains substantively consistent with the ordering reported in Figure 2.4.1, suggesting that the key predictive features—such as regional income dynamics, labor market structure, spatial spillovers, and environmental policy stringency—retain their relative importance. This stability further supports the interpretation that our models identify meaningful economic relationships rather than exploiting autocorrelation patterns in the outcome variable. Notably, inflation dynamics do not rank among the top predictors. Instead, the model features *lag_rate_disp*, *lag_c_rate_disp*, and their spatial lags, which are more directly aligned with the original Phillips curve formulation.

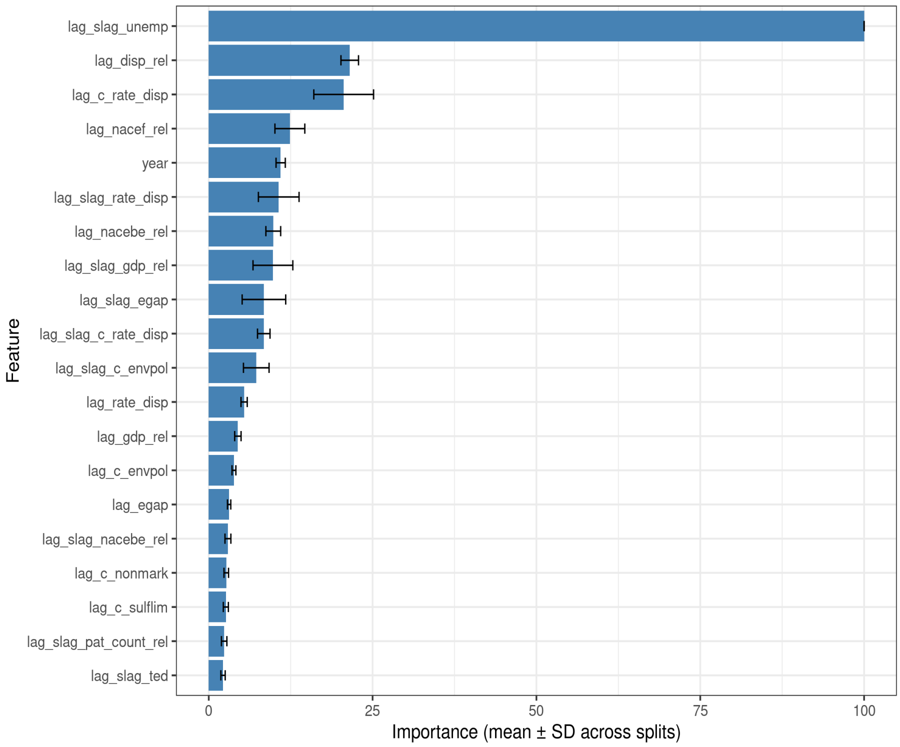

Figure 2.A.1: Feature importance plot - first 20 features by importance. *Note:* Feature importance are computed on both train and test set and averaged across splits.

2.A.3 Temporal holdout validation

To evaluate model performance over a longer forecasting horizon, we implement a temporal holdout validation strategy. Under this specification, models are trained

Table 2.A.3: Model performance with temporal holdout validation

	MR	EN	**GBM**	RF	SVR	NNs
MAE	0.8980	0.8960	**0.7309**	1.1375	1.3080	6.1880
MAPE	0.1628	0.1578	**0.1362**	0.2381	0.2644	0.7946
RMSE	1.2376	1.2200	**1.0132**	1.4906	1.6910	8.0287
nRMSE	0.2418	0.2383	**0.1979**	0.2912	0.3303	1.5683
R^2	0.9415	0.9431	**0.9608**	0.9151	0.8907	−1.4633

Note: Models are trained on data from 2000–2015. Metrics are computed on the test period 2016–2019.

exclusively on data from 2000 to 2015 and evaluated on the out-of-sample period 2016–2019. This design reflects a realistic lag in regional statistical reporting and allows us to assess the practical utility of the proposed framework in a policy-relevant setting, where predictions must be generated several years in advance.

Table 2.A.3 reports the performance metrics for this extended forecasting horizon. The results indicate that the GBM model retains the strongest predictive performance across all evaluation criteria. Although predictive accuracy naturally deteriorates relative to the one-year-ahead forecasts presented in the main analysis, overall performance remains robust.

Figure 2.A.2 displays the yearly MAPE of the best-performing GBM model over the testing period. While predictive accuracy declines as the forecast horizon extends, as we would expect, the model maintains reasonable performance throughout.

2.A.4 NUTS1-stratified split

Table 2.A.4 reports performance metrics under a train/test split stratified by NUTS1 regions. The splitting strategy is designed to ensure that both the training and test sets contain at least one observation for each NUTS1 macro-region in the sample. This guarantees that all regional contexts inform the learning process, addressing the concern that predictive performance could be inflated by the overrepresentation of particular areas in the training data. The results confirm that GBM remains the best-fitting model across all metrics and that predictive accuracy is not materially affected by this alternative splitting strategy, further corroborating the robustness of the main findings.

Figure 2.A.2: Yearly MAPE of out-of-sample prediction under temporal holdout validation

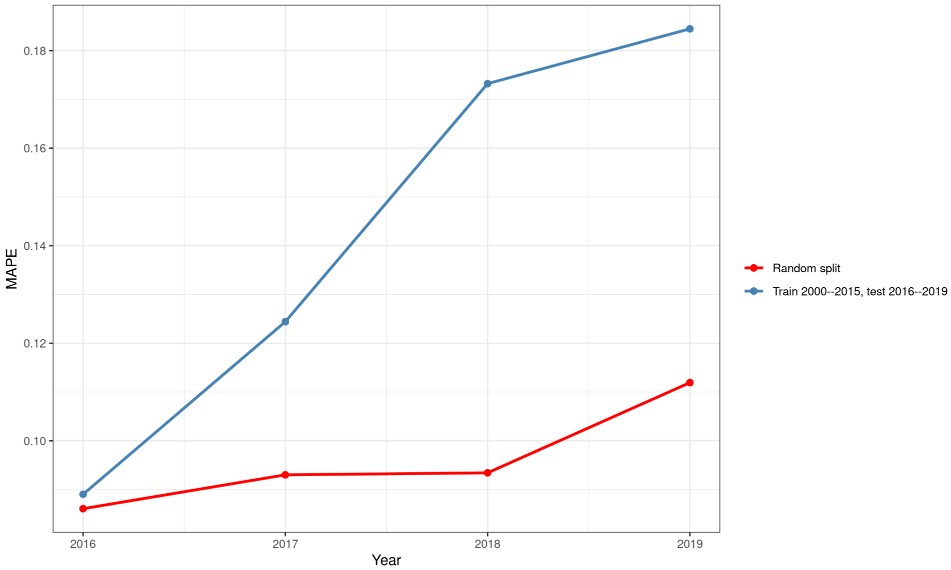

Note: The figure compares yearly (2016–2019) out-of-sample unemployment prediction errors from a GBM model trained on 2000–2015 data (blue line) with the corresponding errors from the original random split design (red line).

	MR	EN	**GBM**	RF	SVR	NNs
MAE	0.8648	0.8613	**0.6977**	0.7471	0.7288	7.3788
MAPE	0.1203	0.1188	**0.0970**	0.0995	0.0983	0.8367
RMSE	1.2416	1.2365	**0.9985**	1.1068	1.1240	8.9678
nRMSE	0.2450	0.2438	**0.1981**	0.2182	0.2154	1.7376
R^2	0.9399	0.9405	**0.9607**	0.9523	0.9535	−2.0239

Table 2.A.4: *Note:* The table displays the median MAE, MAPE, RMSE, nRMSE, and R2 across the ten cross-validation runs on a train-test split that ensures each NUTS1 region is represented in both the train sample and the test sample.

Chapter 3

Inside the Matrix Completion Approach Applications on the Labour Effects of Minimum Wage

This essay is co-authored with Marco Ventura (Sapienza University of Rome).

Abstract

This study advances the understanding of a recent causal machine learning tool through its adaptation to two prominent studies. Specifically, the tool employed is the Matrix Completion method for Causal Panels (MCP), developed by Athey et al. (2021a), and the studies replicated are Card and Krueger (1994) and Callaway and Sant'Anna (2021). The contribution we provide is twofold, both economic and methodological. First, we seek to reconcile divergent evidence on the impacts of minimum wage on (low-skilled) employment and explore potential heterogeneous effects. While most of the literature relies on difference-in-differences (DiD) designs, these approaches rest on the common trends assumption, which is basically non-testable, involving unobservable quantities. Therefore, our second contribution is to explore whether a matrix completion approach relying on an entirely different set of assumptions corroborates previous findings. We find proof of a small, yet significant, negative and amplifying effect of a minimum wage introduction on the employment of targeted groups, in line with the prevailing literature. Second, we delve deep into MCP functioning, highlighting its potential and limitations. Last, but not least, we show that, in principle, MCP offers the possibility to test for untestable hypotheses.

Keywords: matrix completion, minimum wage, difference-in-difference, parallel trends

JEL Classification: C23, C52, J23

3.1 Introduction

Recent political debates have reignited discussions on minimum wage policies, placing the issue back in the spotlight. Mario Draghi's report on EU competitiveness highlights the pressing need to establish fair and adequate remuneration for European workers as a cornerstone for revitalizing the continent's economy Draghi (2024). Even so, a minimum wage policy is regarded as a venture with uncertain political outcomes (Fazio and Reggiani, 2023). Additionally, the idea that minimum wage can have positive macroeconomic effects is still highly debated. The most common argument against makes the case that a rise in labour price would result in a drop of employment and, consequently, to economic backlash. This idea, which has a long history in economic thought (for a review, see Neumark (2019)), has been often challenged by empirical findings.

A landmark example is the work of Card and Krueger (1994) (CK henceforth), who leveraged a natural experiment involving a minimum wage increase in Pennsylvania to examine its impact on firm-level employment. Their contribution was twofold: first, they introduced the Difference-in-Differences (DiD) method as a practical tool for causal inference; second, they challenged the conventional economic theory that minimum wage increases reduce employment. Instead, their findings suggested that the wage increase had a positive effect on employment, particularly in low-wage sectors.

Subsequent studies produced mixed results. Neumark and Wascher (2000) reported negative employment effects for low-wage workers, whereas the same Card and Krueger (2000) found no significant impact when using administrative data, and acknowledged that firm-level heterogeneity might obscure aggregate dynamics. This issue was further explored by Ropponen (2011), who found that large firms tend to experience negative employment effects, while smaller firms may exhibit neutral or even positive responses.

More recently, Callaway and Sant'Anna (2021) (CS henceforth) refined the DiD methodology by incorporating the staggered treatment adoption case, again using minimum wage policy as a case study—this time focusing on teen labor. By developing a statistical framework that accounts for heterogeneity in treatment timing, their findings contrasted with those of Card and Krueger (1994), suggesting that minimum wage increases may negatively impact teen employment, particularly when dynamic effects are considered. Even though the empirical case was not the central contribution of CS's work, their findings are in line with Meer and West (2016), who argue that, while short-term impacts of a minimum wage scheme may be minimal, negative employment adjustments tend to unfold over time.

These results are in contrast with prominent studies indicating that minimum wage policies may reshape the wage distribution rather than eliminating low-wage jobs altogether (Cengiz et al., 2019), adding complexity to impact assessments. More recently, advances in big data and machine learning (ML) have focused on heterogeneity of treatment effects, a factor previously addressed through group-specific or sector-specific analyses (Cengiz et al., 2022). A novel body of research explores the effects of minimum wage policies on reallocation, both geographically and across firms (Dustmann et al., 2022). Finally, the health impacts, particularly on teens, became the subject of study (Matsuzawa et al., 2025).

Despite this persistent interest, the debate over the impact of minimum wage policies on employment remains unresolved. According to Neumark (2019), empirical studies on the topic encounter two crucial challenges. The first is the identification of the control group, given the potential endogeneity or non-randomness of minimum wage increases. The second has to do with the widely imposed parallel-trend restriction (Card and Krueger, 1994, 2000; Cengiz et al., 2019; Matsuzawa et al., 2025; Meer and West, 2016; Neumark and Wascher, 2000).

In this paper, we revisit the empirical cases analyzed in the seminal works of CK and CS using the Matrix Completion approach for Causal Panel (MCP) developed by Athey et al. (2021a). This framework, which emerged contemporaneously with CS's staggered DiD estimator, is particularly well-suited for analyzing complex and rich data structures. MCP explicitly accommodates both similarities and heterogeneities across units, offering a flexible, data-driven alternative to the parallel trends assumption for counterfactual imputation. Ultimately, we seek to uncover whether the introduction of minimum wage yields a significant and long-lasting effect on employment.

Our work relates to two strands of literature. First, we contribute to the literature that analyses the effects of minimum wage on employment by applying a new method to reconcile the divergent findings of two influential studies. Second, we contribute to the recent and growing debate on causal ML by highlighting the strengths and the weaknesses of this approach over simpler frameworks. The remainder of the paper is organized as follows. In Section **??** we briefly introduce the MCP methodology, in Section 3.3 and Section 3.4 we replicate CK and CS, respectively. Section 3.5 provides details on heterogeneous effects. Section 3.6 shows how to test for untestable restrictions and, finally, Section 3.8 offers concluding remarks and recommendations for future research.

3.2 MCP for ATE estimation

The core innovation of Athey et al. (2021a) is viewing the literature on unconfoundness Rosenbaum and Rubin (1983) and on synthetic controls Abadie et al. (2010) as specific cases of the same imputation problem, solvable through a matrix completion machine-learning algorithm. Matrix completion was first theorized in Candes and Recht (2012) and is now an established practice in the field of compressed sensing.

Imagine we observe the outcome Y for N units, over T periods. We have information on the units' treatment status from a matrix W, with dimensions NxT, where $W_{it} = 1$ if unit i is treated at time t and $W_{it} = 0$ otherwise.

$$Y = \begin{pmatrix} Y_{11} & Y_{12} & Y_{13} & \cdots & Y_{1T} \\ Y_{21} & Y_{22} & Y_{23} & \cdots & Y_{2T} \\ Y_{31} & Y_{32} & Y_{33} & \cdots & Y_{3T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ Y_{N1} & Y_{N2} & Y_{N3} & \cdots & Y_{NT} \end{pmatrix} \quad \text{(realized outcome)}$$

$$W = \begin{pmatrix} 1 & 1 & 0 & \cdots & 1 \\ 0 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 1 & \cdots & 0 \end{pmatrix} \quad \text{(binary treatment)}$$

Leading to the following potential outcome matrices for the untreated $Y(0)$ and for the treated $Y(1)$

$$Y(0) = \begin{pmatrix} ? & ? & \checkmark & \cdots & ? \\ \checkmark & \checkmark & ? & \cdots & \checkmark \\ ? & \checkmark & ? & \cdots & \checkmark \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ ? & \checkmark & ? & \cdots & \checkmark \end{pmatrix}$$

$$Y(1) = \begin{pmatrix} \checkmark & \checkmark & ? & \cdots & \checkmark \\ ? & ? & \checkmark & \cdots & ? \\ \checkmark & ? & \checkmark & \cdots & ? \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \checkmark & ? & \checkmark & \cdots & ? \end{pmatrix}$$

$\checkmark$ mark a set of indices i, t for which Y is observed, ? a set for which Y is missing. Hereafter, we will refer to the group of observed entries with O and to the group of missing entries with M.

The average treatment effect on the treated (ATT) is calculated as:

$$ATT = \frac{\sum_{i,t} W_{it} \left(Y_{it}(1) - Y_{it}(0)\right)}{\sum_{it} W_{it}}$$

Therefore, we need to impute the missing potential outcomes in matrix $Y(0)$ (*i.e.* the entries in Y(0) that correspond to the unit-time couple where $W_{it} = 1$). For this scope, define the two matrices $P_O(Y)$ and $P_O^{\perp}(Y)$:

$$P_O(Y)_{it} = \begin{cases} Y_{it} & \text{if } (i,t) \in O \\ 0 & \text{if } (i,t) \notin O \end{cases}$$

$$P_O^{\perp}(Y)_{it} = \begin{cases} 0 & \text{if } (i,t) \in O \\ Y_{it} & \text{if } (i,t) \notin O \end{cases}$$

In the unconfoundness setup, the ATE is estimated under the assumption that the patterns over time are stable across units. As the core information comes from the time vector, the unconfoundedness approach is suitable when $T \gg N$. On the other hand, in a synthetic control setup, the ATE is estimated assuming patterns across units are stable over time. Intuitively, this approach applies to cases where $T \ll N$.

The intuition behind MCP is that it is possible to exploit the information carried by both patterns across units and patterns over time when imputing Y(0). This being the case, this framework is applicable in settings where both N and T are large. Moreover, this approach is flexible to empirical settings that differ from the traditional block structures (one pre-, one post- treatment period; stable groups of treated and controls), *e.g.* the staggered adoption case. Let us briefly recap the main features of MCP.

Let the matrix of complete outcomes Y be defined as

$$Y = L^{\star} + \epsilon \tag{3.2.1}$$

with ϵ measurement error assumed to be independent of $L^{\star}$, $E[\epsilon|L^{\star}] = 0$. The goal is to estimate $L^{\star}$. $L^{\star}$ is assumed to be a low-rank matrix and that can be considered also as including fixed effects that can be partialed out by rewriting $L^{\star} = L + \Gamma \mathbf{1}_T + \Delta \mathbf{1}_N$, where Γ is a $N \times 1$ matrix of unit fixed effects and Δ a $T \times 1$ matrix of time fixed effects.

An estimate of $L^{\star}$ can be obtained by solving the following regularized minimization problem

$$(\hat{L}, \hat{\Gamma}, \hat{\Delta}) = \arg\min_{L,\Gamma,\Delta} \left\{ \frac{1}{|O|} \| P_O(Y - L - \Gamma^{\top}\mathbf{1}_T - \Delta^{\top}\mathbf{1}_N) \|_F^2 + \lambda \|L\| \right\} \tag{3.2.2}$$

The term $\lambda\|L\|$ indicates a penalty term, and the hyperparameter λ is determined through cross-validation. As the share of observed values is relatively high, the fixed effect can be estimated with great precision. For this reason, only L is regularized, and not also Γ and Δ. At first glance, the problem is very much akin to a LASSO model, but Athey et al. (2021a) argue that a specific machine-learning technique, the nuclear norm minimization (NNM) algorithm, can optimally perform the imputation task. Simply put, a NNM minimizes the sum of the singular values of a matrix, which is used as a convex alternative to the matrix rank. The rank of a matrix corresponds to the number of linearly independent rows or columns; therefore, a low-rank matrix can in principle be reconstructed as a linear combination of such rows and columns. In place of the matrix rank, the literature on matrix completion employs the nuclear norm of the matrix, which is defined as the sum of its singular values. Therefore, the nuclear norm is continuous, which makes it a convenient alternative for rank minimization problems (Candes and Recht, 2012).

In the case of MCP, the NNM is used to minimize the distance between the observed Y and the estimator L, calibrated for λ, based on the observed data. The process is reiterated for different values of λ, such to minimize the mean square error (MSE). NNM has been widely applied to matrix completion problems, particularly through the Soft-Impute algorithm developed by Mazumder et al. (2010). See Figure 3.2.1 for a representation of how MCP algorithm imputes missing values.

MCP relies on minimal assumptions: i) the stable unit treatment value assumption, SUTVA; ii) the once-treated-always-treated assumption; iii) the assumption that the incomplete matrix Y(0) is low-rank. It is under this last strictly mathematical condition that Y(0) can be approximated by its nuclear norm $\hat{L}$. See Sec. 3.2.1 for a more in-depth discussion on the mathematical assumptions behind MCP and 3.2.2 for a grasp on the economic meaning of these conditions.

In Athey et al. (2021a), the superiority of the MCP approach with respect to DiD was presented in a simulation exercise on Abadie et al. (2010)'s California smoking data. In this application, MCP outperforms DiD in terms of MSE in a prediction exercise. However, it is important to highlight that we do not know to what extent this performance of DiD translates into estimation discrepancies when applied to a genuine estimation case of treatment effects. In the subsequent analysis, we

Figure 3.2.1: Representation of the MCP algorithm.

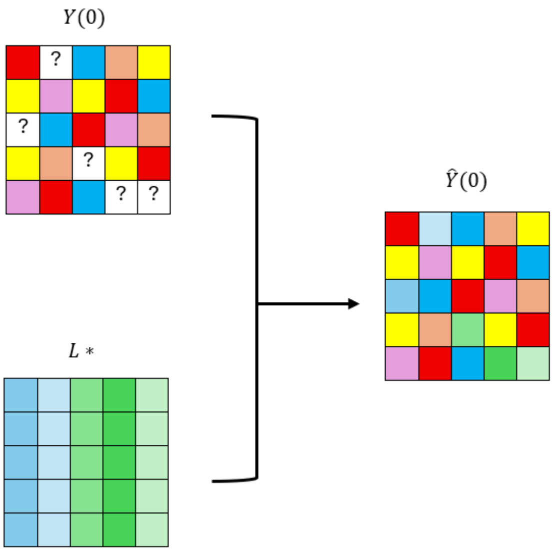

Note: MCP estimates a complete matrix $\hat{Y}(0)$ from the partially observed matrix $Y(0)$ by imputing missing entries (denoted by "?"), exploiting information from the low-rank structure of Y(0), L^*. Notice that the colour of the cells in $\hat{Y}(0)$ corresponding to the "?" in $Y(0)$ is the same as the one in L^*

investigate whether the MCP approach yields significantly different ATE estimates compared to DiD. Observing close results, despite divergent underlying assumptions, would suggest that the parallel trends assumption holds within the data. Consequently, the MCP approach can serve as an indirect validation mechanism for the DiD results.

3.2.1 Mathematical assumptions: low-rank and incoherence

Matrix completion originated as a solution to the so-called "Netflix problem" (Candes and Recht, 2012), which consisted in predicting user preferences for media content when only a small subset of user ratings was observed. To make such predictions feasible, it was necessary to assume that users and titles share common underlying characteristics. Formally, this assumption translates into the hypothesis that the user-title matrix is low-rank, meaning that, despite being very high-dimensional, the data can be well-approximated by the linear combination of a small number of rows and columns. In practical terms, the intuition was that users enjoy content that is in some ways similar, and titles appreciated by a user are likely to appeal to similar users. In practical terms, assuming that Y(0) is low-rank is equivalent to conceding that observations over units and periods can be obtained as a combination of one another. This aspect will be presented in further detail in Section 3.2.2. Note that, in practice, the low-rank condition is assessed on the training matrix P_O. Therefore, it is likely to be violated when there are no common patterns in the data, and/or when the number of entries to be imputed is very limited. A characteristic of MCP that could be a limitation in some practical applications is that it requires a large volume of data, with a disproportionately higher number of values to be imputed compared to those observed.

This assumption adds up to what Candes and Recht (2012) defined as an incoherence condition. The concept of coherence broadly refers to the concentration of the main singular values over rows and columns. The fact that a matrix has low coherence guarantees that its low-rank structure can be faithfully reconstructed starting from a random subsample of its entries. If the matrix were highly coherent, on the other hand, the information about its low-rank structure would be concentrated in a few rows or columns. If those regions were missing, the low-rank structure would become unrecoverable. This mathematical condition is not explicitly mentioned in Athey et al. (2021a), yet it guarantees that any matrix completion algorithm works properly. Candes and Recht (2012) explicitly formalized this concept from a mathematical point of view, while Athey et al. (2021a) do not delve into this criterion. This could be a critical oversight, given the distinctly non-random placement of missing values in panel contexts. Indeed, the fact that missing data appears in "chunks" within columns likely violates the coherence assumption in numerous economic applications—a challenge that warrants careful consideration. We argue that the exploration of coherence properties represents a promising direction for future research.

3.2.2 The economic intuition

In economic applications, assuming that a panel data matrix is low-rank implies the existence of interdependencies across units and over time. This framework contrasts sharply with conventional econometric models that typically assume independent units of observation, such is the case with DiD. Instead, the MCP approach acknowledges and exploits observed patterns to impute missing data. To put this in mathematical terms, minimizing the rank (or, as is our case, the nuclear norm) of the matrix means maximizing the probability of reconstructing missing values from the observed data.

On top of the low-rank assumption, the incoherence condition also requires attention. Even though this intuition concerns the estimator's overall economic sense, the significance of the coherence condition remains unexplored.

3.3 MCP for contemporaneous treatment: the case of Card and Krueger (1994)

The case of CK is one of the most influential empirical studies on the impact of minimum wage on employment. The study analyzed the effect of a minimum wage rise in New Jersey on employment by comparing it to neighboring Pennsylvania, where the minimum wage did not change. In April 1992, New Jersey raised its minimum wage from $4.25 to $5.05 per hour. Pennsylvania also had a minimum wage in place at the time, but the amount did not increase at the time of the rise in New Jersey. Data on employment, wages, and product prices come from a survey carried out in two waves, once in the pre-treatment period and once in the post-treatment period, at fast-food restaurants in the two states. The mean of full-time employment (FTE) in the two states in the two periods is reported in Table 3.3.1. The difference in employment levels before and after the intervention for both the treatment and control groups is attributed to the effect of the intervention.

Table 3.3.1: Mean of FTE, by period

	NJ	PA
FTE_pre	10.21	7.72
FTE_post	7.56	8.45

Table 3.3.2: Difference in FTE between New Jersey and Pennsylvania

variable	NJ-PA
FTE employment before	-2.60
SE employment before	1.30
FTE employment after	0.84
SE employment after	1.19
change in mean FTE	**3.44**
SE of change in mean	1.56

Narrow-sense replications of the study are ubiquitous, although we found no systematic peer-reviewed publications. Thus, using the publicly available data, we have carried out the narrow-sense exercise matching exactly the original result which points to a positive treatment effect, as depicted in Table 3.3.2. As for the replication in a broad-sense, we propose to recast the estimation exercise in terms of the MCP [1], which relies on a completely disjoint set of hypotheses that have nothing to do with parallel trend. This could help to overcome criticism that the original survey data

[1] All the MCP estimations in this work are performed using the MCPanel R package (Athey et al., 2021b).

violated the common trend assumption (Neumark and Wascher, 2000) and reconcile the findings with Card and Krueger (2000).

Using the notation introduced above, CK's problem can be reframed as follows. The outcome variable employment (Y) is collected in 371 fast-food restaurants in New Jersey, and 79 restaurants in Pennsylvania, during two waves ($T = 2$), hence the observed $N \times T$ matrix Y looks like

$$Y = \begin{pmatrix} Y_{NJ_{1,1}} & Y_{NJ_{1,2}} \\ \vdots & \vdots \\ Y_{NJ_{371,1}} & Y_{NJ_{371,2}} \\ Y_{PA_{1,1}} & Y_{PA_{1,2}} \\ \vdots & \vdots \\ Y_{PA_{79,1}} & Y_{PA_{79,2}} \end{pmatrix} \quad \text{(observed outcome)}$$

Restaurants in New Jersey are treated at t = 2, leading to the following incomplete matrix for the untreated $Y(0)$.

$$Y(0) = \begin{pmatrix} Y_{NJ_{1,1}} & ? \\ \vdots & \vdots \\ Y_{NJ_{371,1}} & ? \\ Y_{PA_{1,1}} & Y_{PA_{1,2}} \\ \vdots & \vdots \\ Y_{PA_{79,1}} & Y_{PA_{79,2}} \end{pmatrix}$$

This configuration corresponds to what Athey et al. (2021a) refer to as a "block structure" case, particularly a single-treated-period block structure.

When running the MCP routine on CK empirical case, we found the resulting L matrix to be null. Consequently, the RMSE of the MCP estimation is the same for any value of the hyperparameter λ over the grid.[2] This situation is depicted in Figure 3.3.1, which presents the value of the RMSE as a function of λ, showing a flat line. It follows that, in this case, $L^{\star}$ is determined solely by fixed effects.

We compute the ATT as the average difference between the value of Y(0) for

[2]The cross-validation grid search was conducted over 100 runs, ranging from the minimum number that makes $L = 0$ to $1e^{-3}$ times this minimum number.

Figure 3.3.1: Replication of CK - Plot of RMSE over different values of λ

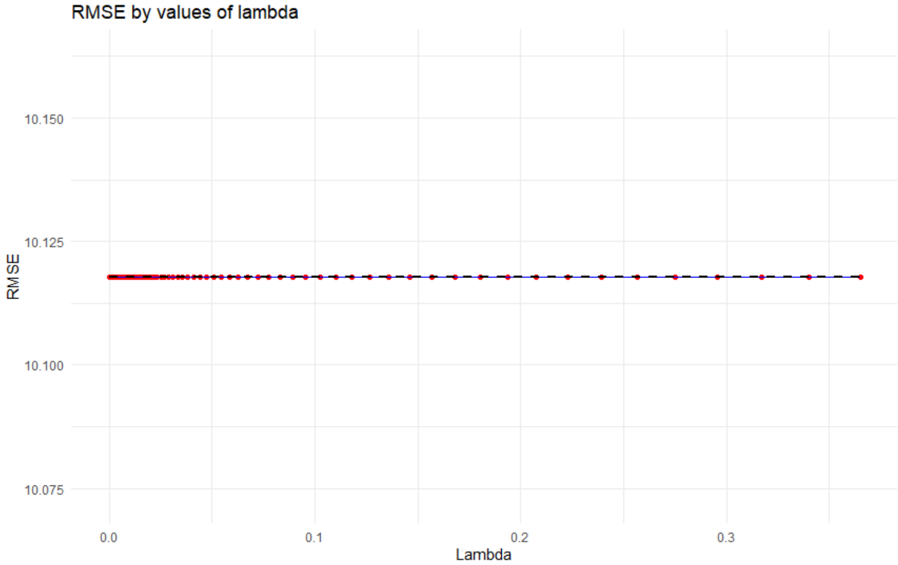

the units where W=1 as imputed by the MCP algorithm and the observed value of Y(1) on the same units. Note that CK results are precisely replicated, resulting in an estimated positive effect of 3.44. The same as in the original paper.

This finding suggests the possibility that the MCP estimator can trivially align with the simplest version of DiD, the so called 2×2 case, in scenarios where the MCP algorithm fails to recover an estimate for L. The most likely explanation is that, if the number of missing entries is relatively high compared to the dimension of the matrix, the algorithm struggles to accurately impute missing values due to the limited information it has access to. Because this happens even though this matrix is extremely low-rank, we hypothesize that the explanation may be linked to a coherence issue, too. Conversely, with a richer observational set and more distributed missing values, the probability of accurately estimating $\hat{L}$ increases. We will discuss this issue in further detail in the next section.

3.4 MCP for staggered adoption: revisiting Callaway and Sant'Anna (2021)

In a similar spirit, we apply the MCP approach to CS data on teen employment, aiming to observe MCP performance in staggered treatment adoption setups.

CS applied the staggered DiD framework to study the impact of state minimum wage increases on county-level teenage employment. Their central innovation is the incorporation of event-study designs to estimate dynamic treatment effects. By comparing treated units to appropriate control groups (either never-treated or the not-yet-treated) at each period, CS construct cohort-time ATTs and propose a battery of partially aggregated measures for treatment effect.

The panel comprises 2,284 observations spanning eight years, from 2001 to 2007. In 2000, the federal minimum wage was increased to \$5.05. Over the subsequent years, many U.S. states independently raised their minimum wages by varying amounts. State-level minimum wage increases occurred in 2004, 2006, and 2007. During this period, 907 counties experienced autonomous minimum wage increases. The remaining counties served as the never-treated control group.

Using the MCP approach, this empirical case can be represented as follows:

$$Y = \left(\begin{array}{ccc} Y_{1,1} & \cdots & Y_{1,8} \\ \vdots & \cdots & \vdots \\ Y_{1377,1} & \cdots & Y_{1377,8} \\ \hline Y_{1378,1} & \cdots & Y_{1378,8} \\ \vdots & \cdots & \vdots \\ Y_{1478,1} & \cdots & Y_{1478,8} \\ \hline Y_{1479,1} & \cdots & Y_{1479,8} \\ \vdots & \cdots & \vdots \\ Y_{1702,1} & \cdots & Y_{1702,8} \\ \hline Y_{1703,1} & \cdots & Y_{1703,8} \\ \vdots & \cdots & \vdots \\ Y_{2284,1} & \cdots & Y_{2284,8} \end{array}\right) \qquad \text{(realized outcome)}$$

- Rows Y_1 to Y_{1377}: the never-treated.
- Rows Y_{1378} to Y_{1478}: counties in states where the minimum wage rise occurred in 2004.
- Rows Y_{1479} to Y_{1702}: counties in states where the minimum wage rise occurred in 2006.

- Rows Y_{1703} to Y_{2284} : counties in states where the minimum wage rise occurred in 2007.

Therefore, the imputation problem takes the form of:

$$
Y(0) = \left(\begin{array}{cccccccc}
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \\
\hline
\checkmark & \checkmark & \checkmark & \checkmark & ? & ? & ? & ? \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\checkmark & \checkmark & \checkmark & \checkmark & ? & ? & ? & ? \\
\hline
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & ? & ? \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & ? & ? \\
\hline
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & ? \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & \checkmark & ?
\end{array}\right)
$$

Notice that in this case we have a much richer dataset and, contrary to the previous exercise, cross-validation on λ is successful and converges to an optimal value $\lambda^{\star} = 0.002$, corresponding to a RMSE of 0.2281, as shown in Figure 3.4.1.

Once obtained an estimate of $L^{\star}$ it is possible to compute the treatment effect of each unit-time observation in the sample, $Y_{it}(1) - Y_{it}(0)$, and aggregate them according to different criteria. In order to compare the quality of the estimate we have obtained with the results of CS we have computed the same treatment effects as in CS and synoptically compared them. Table 3.4.1 presents the result. To ease comparison with the MCP estimator, we present here the case where the control group for the DiD estimation consists of the not-yet-treated units (those for which $W_{it} = 0$). The results with the never-treated control group are presented in the Appendix 3.A.4.

Broadly speaking, MCP greatly aligns with CS's results both in terms of cohort-specific effects and in terms of relative time, *i.e.* event study. The estimates show a negative effect of minimum wage on teen employment, and the effect decreases over

Figure 3.4.1: Replication of CS - Plot of RMSE over different values of λ

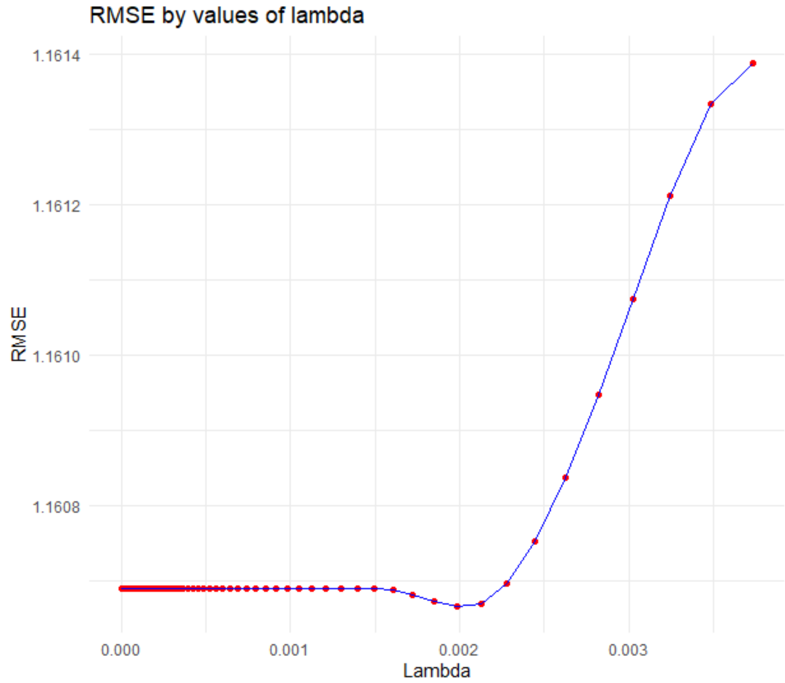

cohorts. As a further interesting result, we observe that the estimates generated by MCP are more precise than those reported by CS and this is an appreciable feature. Figure 3.A.1 provides an immediate view of the proximity of the results obtained with the two methods.

Now we move a step further by extending the replication to the case of conditional parallel trends, as outlined in CS. In the CS's staggered DiD estimator, covariates are kept constant at the last value before treatment. In contrast, the MCP approach allows for the inclusion of time-varying covariates. As part of this extension, we reconstructed the database to ensure a consistent and reliable replication. Our reconstruction adheres as much as possible to the sources indicated by CS. This ensures that our dataset remains methodologically sound and comparable to the original study.

CS utilize county-level teen employment data from the Quarterly Workforce Indicators (QWI) along with pre-treatment county characteristics from the 2000 County Data Book. These characteristics include county population in 2000, the proportion of white inhabitants, educational attainment in 1990, median income in 1997, and the fraction of the population below the poverty level in 1997. To construct a panel

Table 3.4.1: Cohort-specific effects and event study, unconditional parallel trends. Control group: not-yet-treated.

	MCP	**Staggered DiD**
Cohort-Specific Effects		
$g = 2004$	-0.081^{***}	-0.093^{***}
	(0.004)	(0.020)
$g = 2006$	-0.014^{***}	-0.042^{***}
	(0.002)	(0.008)
$g = 2007$	-0.045^{***}	-0.028^{***}
	(0.007)	(0.007)
Event Study		
$e = 0$	-0.026^{***}	-0.025^{***}
	(0.001)	(0.006)
$e = 1$	-0.053^{***}	-0.074^{***}
	(0.001)	(0.009)
$e = 2$	-0.106^{***}	-0.118^{***}
	(0.015)	(0.020)
$e = 3$	-0.122^{***}	-0.136^{***}
	(0.018)	(0.022)

Note: The table compares the estimated effects of the minimum wage on teen employment using CS staggered DiD estimator versus MCP. In the latter case, the treatment effect is obtained as the mean difference between the counterfactual outcome $Y(0)$, as imputed by MCP in the absence of covariates, and the observed outcome $Y(1)$ for treated counties. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. See Appendix 3.A.3 for details on the computation of the standard errors.

of time-varying covariates, we leverage data from the county intercensal datasets (county population and racial composition) and the Small Area Income and Poverty Estimates (SAIPE) Program (poverty rate and median income), both provided by the U.S. Census Bureau. For details on variables' names and units of measure, Table 3.4.2. For further information, we refer to the Appendix 3.A.2. We argue that incorporating time-varying covariates challenges the conditional parallel trends argument put forth in CS and introduces additional complexity to the analysis, further justifying the use of the MCP approach.

Unfortunately, annual county-level data regarding educational attainment are not accessible, precluding a precise narrow-sense replication. To achieve the closest possible approximation to the published results, we employed the CS estimator on all covariates, excluding educational attainment.

Figure 3.4.2: Group-time average treatment effects, unconditional parallel trends. Control group: not-yet-treated.

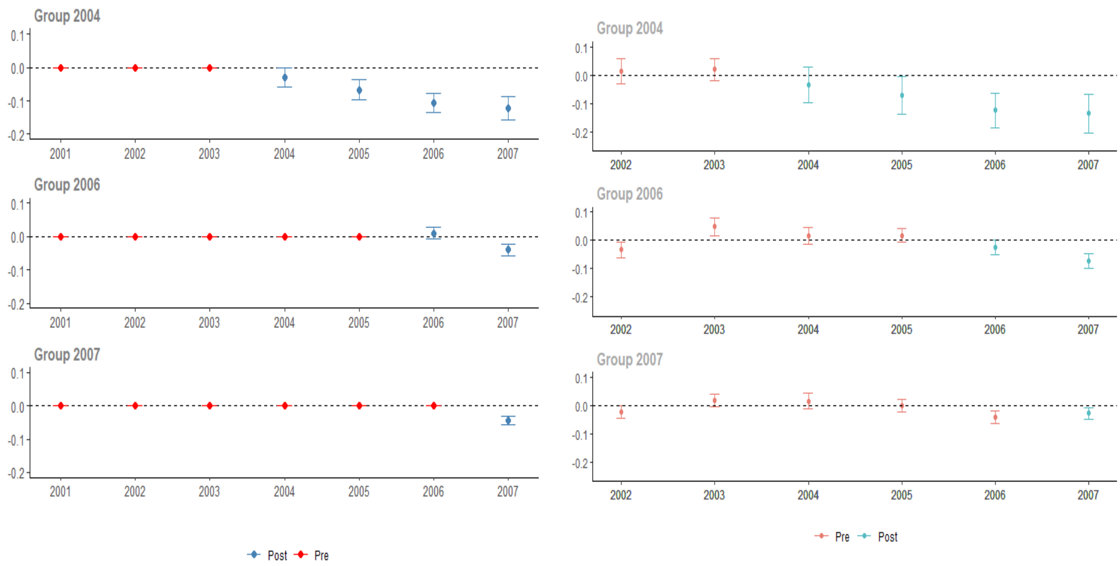

Note: The figure compares group-time treatment effects obtained using MCP (left) with those from CS (right) in the case of unconditional parallel trends. The top row corresponds to states that increased their minimum wage in 2004, the middle row to states that did so in 2006, and the bottom row to states that implemented the increase in 2007. Blue lines represent treatment effect point estimates and the relative 95% confidence bands. Red lines indicate point estimates and 95% confidence bands for pre-treatment periods. For the parallel trends assumption to hold, these estimates should include zero. Note that, in the MCP framework, pre-treatment effects are set to zero by construction.

The resulting estimates, presented in Table 3.4.3 column 2, demonstrate substantial agreement with the published findings and serve as the basis for comparison with the MCP performance, as depicted in Table 3.4.3 column 1. An immediate comparison of the results from the two methods is also offered in Figure 3.A.2. Again, the results are very similar and MCP estimates are more precise. Like for the unconditional case, we present the case where the not-yet-treated constitute the control group, while the comparison with the never-treated can be found in the Appendix 3.A.4.

Table 3.4.2: Covariates and their source

Variable	**Source**
County population (log)	US Census Bureau - County Intercensal Datasets (2024)
Fraction of whites on county population	US Census Bureau - County Intercensal Datasets (2024)
Poverty rate	Small Area Income and Poverty Estimates (SAIPE) Program (2024)
Median income (thousands of $)	Small Area Income and Poverty Estimates (SAIPE) Program (2024)

Table 3.4.3: Cohort-specific effects and event study, conditional parallel trends. Control group: not-yet-treated.

	MCP	**Staggered DiD**
Cohort-Specific Effects		
$g = 2004$	-0.082^{***}	-0.095^{***}
	(0.004)	(0.019)
$g = 2006$	-0.014^{***}	-0.040^{***}
	(0.002)	(0.008)
$g = 2007$	-0.049^{***}	-0.030^{***}
	(0.007)	(0.007)
Event Study		
$e = 0$	-0.023^{***}	-0.026^{***}
	(0.001)	(0.005)
$e = 1$	-0.053^{***}	-0.075^{***}
	(0.001)	(0.009)
$e = 2$	-0.030^{***}	-0.118^{***}
	(0.001)	(0.020)
$e = 3$	-0.067^{***}	-0.141^{***}
	(0.001)	(0.022)

Note: The table compares the estimated effects of the minimum wage on teen employment using CS staggered DiD estimator versus MCP, including the covariates reported in Table 3.4.2. In the latter case, the treatment effect is obtained as the mean difference between the counterfactual outcome $Y(0)$, as imputed by MCP, and the observed outcome $Y(1)$ for treated counties. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. See Appendix 3.A.3 for details on the computation of the standard errors.

Figure 3.4.3: Group-time average treatment effects, conditional parallel trends. Control group: not-yet-treated.

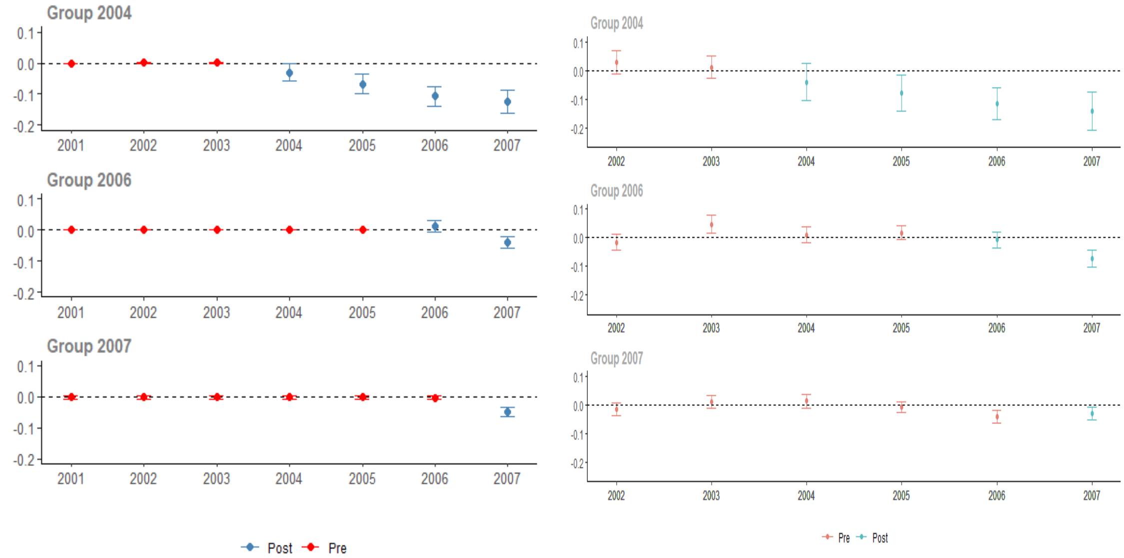

Note: This figure compares group-time treatment effects obtained using MCP (left) with those from CS (right) in the case of conditional parallel trends. The top row corresponds to states that increased their minimum wage in 2004, the middle row to states that did so in 2006, and the bottom row to states that implemented the increase in 2007. Blue lines represent treatment effect point estimates and the relative 95% confidence bands. Red lines indicate point estimates and 95% confidence bands for pre-treatment periods. For the parallel trends assumption to hold, these estimates should include zero. Note that, in the MCP framework, pre-treatment effects are set to zero by construction.

Although seemingly marginal, the positive - yet not significant effect revealed by MCP for the 2006 cohort, observed in both the unconditional and the conditional case, potentially reconciles the two competing perspectives in the literature. Specifically, this finding bridges the consensus regarding modest negative effects following minimum wage introduction (also evident in CS's results) with CK's assertion of positive effects. In this sense, MCP concedes the possibility of at least a neutral result in the first post-treatment period, while showing that a negative outcome emerges in the subsequent periods.

It is interesting to notice that moving from assumptions quite different from DiD, MCP provides close results to DiD, reinforcing the validity of the parallel trends and no anticipatory effect assumptions. In fact, MCP captures similar common patterns in the data.

A primary benefit of the MCP method lies in its capability to assess individual potential outcomes $Y_{i,t}(0)$, rather than depending exclusively on average effects for groups over time. This enables a finer comprehension of how treatment effects vary across different units and periods. The subsequent section will explore this facet in further detail.

3.5 Hidden heterogeneity

Figure 3.5.1 visually represents the estimated treatment effects at the county level, highlighting the heterogeneity across treated units. Specifically, for each treated unit, the figure displays the difference between the counterfactual outcome $Y(0)$, as imputed by MCP, and the observed post-treatment outcome $Y(1)$. Each horizontal line corresponds to a treated unit, color-coded to indicate the intensity and direction of the effect: blue denotes a negative impact of the minimum wage on teen employment at each point in time, while red indicates a positive effect.

This representation underscores the presence of substantial heterogeneity in treatment effects, revealing patterns that remain obscured in aggregated results such as those in Tables 3.4.1 and 3.4.3. By leveraging MCP, we can go beyond average group or event-time effects to examine the full distribution of individual outcomes, and therefore effects. This feature has no equals in the literature on impact assessment and offers a clear advantage over traditional techniques.

Note that estimates are highly sensitive to covariates. As the right panel of Figure 3.5.1 shows, incorporating covariates into the MCP routine both increases the overall magnitude of negative estimates and results in a higher occurrence of units that experience a net negative effect. This finding suggests that controlling for additional covariates, while computationally expensive, greatly enriches estimates.

Figure 3.5.1: Treatment effects - point estimates:

Note: Estimation without covariates (on the left) and with covariates (on the right). Counties are disposed along the vertical axis. Each horizontal line corresponds to a treated unit, blue denotes a negative impact of the minimum wage on teen employment at each point in time, while red indicates a positive effect.

The possibility and value of obtaining point estimates was not explored by Athey et al. (2021a), as the authors focus primarily on the estimation of aggregated treatment effects. We argue, however, that the ability to recover unit-level counterfactual estimates — which is ingrained in the matrix completion procedure itself — represents a relevant methodological contribution in its own right, and one with meaningful implications for policy. Access to point estimates allows researchers and policymakers to evaluate treatment effects case-by-case, and ultimately to design more place-sensitive policies that account for local conditions and experiences. This granularity is especially valuable in settings where unit-level characteristics are likely to interact with the treatment in differentiated ways. The introduction of a minimum wage provides a good illustration. Figure 3.5.1 paints a picture that is consistent with the literature (Dustmann et al., 2022; Ropponen, 2011): heterogeneous employment effects are shaped by the productive and economic structure of each local labour market, underscoring the need for more refined estimates to support place-sensitive policy recommendations.

3.6 Testing the untestable

The output from MCP lends itself to carrying out tests on hypotheses that are untestable by definition. Indeed, having the entire potential outcome matrix under no treatment, $\mathbf{Y}(\mathbf{0})$, one can test the assumptions based on non-observable quantities, such as the parallel trend assumption in DiD contexts.

This possibility depends on the fact that MCP does a good imputation job. If that is the case, we would expect the imputed counterfactuals to respect common trends after treatment if common trends were indeed present in CS data. This is especially true considering how, in our case, the replication of CS using MCP does not yield substantially different results compared to the original paper. Therefore, we expect MCP to broadly support the parallel trends hypothesis in CS data. Nevertheless, MCP does not explicitly require common trends to impute counterfactuals. This feature, which is arguably the greatest advantage of the MCP setup, may result in imputed counterfactuals that deviate from common trends. If this is the case, it may become more difficult to assess whether the imputed values are true to the observational data and consistent with economic understanding.

In this section, we first show a possible way for checking the quality of MCP imputation, and then we will use the $\mathbf{Y}(\mathbf{0})$ matrix to test for parallel trends. In doing so, sections 3.6.1 and 3.6.2, respectively, offer a window into how MCP assumptions are operationalised in practice in our application. However, the tests proposed below should be interpreted as broad diagnostic tools, whose applicability extends beyond the empirical setting presented in this work. As such, we hope they prove useful to practitioners seeking to assess imputation quality and validate counterfactual assumptions in other MCP applications.

3.6.1 A check of the quality of the imputation

Following the idea of placebo tests, we search for treatment effects on $\mathbf{Y(0)}$ in post-treatment years. Our goal is to assess the internal consistency of the assumptions under MCP imputation, rather than to test for the truth.

Under the assumption of no other exogenous shocks that impact the outcome variable, other than treatment, we expect to find no effects. In our specific context and in our particular dataset we re-apply the CS estimator on such an outcome. The results are summarized in Tables 3.6.1 and 3.6.2 for the unconditional and conditional case, respectively. The results provide mixed evidence as three effects out of seven are statistically significant at 5%, and are pretty stable across different conditioning sets and comparison groups.

Table 3.6.1: Cohort-specific effects and event study on the imputed $Y(0)$, unconditional case. Control group: not-yet-treated.

	Never-treated	**Not-yet-treated**
Cohort-Specific Effects		
$g = 2004$	-0.010	-0.012
	(0.0095)	(0.0085)
$g = 2006$	-0.033^{**}	-0.027^{**}
	(0.0073)	(0.0068)
$g = 2007$	0.018^{**}	0.018^{**}
	(0.0056)	(0.0061)
Event Study		
$e = 0$	0.003	0.005
	(0.0037)	(0.0043)
$e = 1$	-0.023^{**}	-0.025^{**}
	(0.0054)	(0.0060)
$e = 2$	-0.019	-0.012
	(0.0105)	(0.0096)
$e = 3$	-0.013	-0.013
	(0.0101)	(0.0102)

Note: The table reports the estimated effects of the minimum wage on teen employment using CS staggered DiD on the counterfactual outcome variable generated by MCP, $Y(0)$. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. See Appendix 3.A.3 for details on the computation of the standard errors.

Table 3.6.2: Cohort-specific effects and event study on the imputed $Y(0)$, conditional case. Control group: not-yet-treated.

	Never-treated	**Not-yet-treated**
Cohort-Specific Effects		
$g = 2004$	-0.010	-0.010
	(0.0089)	(0.0095)
$g = 2006$	-0.032^{**}	-0.033^{**}
	(0.0068)	(0.0073)
$g = 2007$	0.018^{**}	0.018^{**}
	(0.0058)	(0.0056)
Event Study		
$e = 0$	0.003	0.005
	(0.0037)	(0.0043)
$e = 1$	-0.023^{**}	-0.025^{**}
	(0.0057)	(0.0059)
$e = 2$	-0.019	-0.012
	(0.0100)	(0.0105)
$e = 3$	-0.013	-0.013
	(0.0100)	(0.0097)

Note: The table reports the estimated effects of the minimum wage on teen employment using CS staggered DiD on the counterfactual outcome variable generated by MCP, $Y(0)$. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. See Appendix 3.A.3 for details on the computation of the standard errors.

The presence of these significant effects does not completely meet our expectations, as we had hoped to find no significant results. Although these findings may be the result of an incomplete accounting for idiosyncratic shocks in MCP's learning process, a more serious concern regards the quality of the imputation, *i.e.*, the capacity of MCP to generate good approximations of the counterfactual quantities. Indeed, both of these explanations could be true. At any rate, these results emphasize the need to dive deeper into how MCP assumptions compare with previous designs. This conclusion aligns with our understanding that MCP is not a panacea, but rather a tool that, although powerful, trades freedom from untestable assumptions with a degree of mathematical sophistication that may come in the way of clarity, transparency, and interpretability.

3.6.2 Testing the parallel trend assumption

The common trend assumption lies at the heart of the DiD framework, yet it remains fundamentally untestable using traditional methods. MCP, however, offers a promising avenue for directly assessing the plausibility of this assumption—provided that it generates credible counterfactuals.

Let us define G as the time period when a unit first becomes treated, for all units that participate in the treatment, G defines which "treatment group" they belong to. Accordingly, we define G_g to be a binary variable that is equal to one if a unit is first treated in period g (*i.e.*, $G_{i,g} = \mathbf{1}\{G_i = g\}$). Furthermore, we define C to be a binary variable that is equal to one for units that do not participate in the treatment in any time period, and as not-yet-treated a member of group g *i.e.* $G_g = 1$ at time s for $s < \tau_g$ where τ_g represents the first treatment time of group g. In this case $(1 - W_s)(1 - G_g) = 1$

Formally, the unconditional parallel trend assumption taking never-treated as the reference group can be written as:

$$\mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid G_g = 1\right] = \mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid C = 1\right] \tag{3.6.1}$$

and its conditional version is:

$$\mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid X, G_g = 1\right] = \mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid X, C = 1\right] \tag{3.6.2}$$

As for the not-yet-treated, the unconditional and conditional parallel trends assumptions can be written, respectively as:

$$\mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid G_g = 1\right] = \mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid W_s = 0, G_g = 0\right] \tag{3.6.3}$$

$$\mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid X, G_g = 1\right] = \mathbb{E}\left[Y_t(0) - Y_{t-1}(0) \mid X, W_s = 0, G_g = 0\right] \tag{3.6.4}$$

We can easily test these assumptions by performing two mean difference tests between groups, both unconditional and conditional. The comparison groups for these tests are the never-treated group in one instance and the not-yet-treated group in the other. The results are displayed in Table 3.6.3 for the unconditional case, while Table 3.6.4 presents the conditional case. The comparison with not-yet-treated units presents no differences in mean, as expected, for both conditioning sets. Conversely, when comparing with never-treated units, the results show no significant difference for the 2004 treatment group, but statistically significant differences for the 2006 and 2007 cohorts. These results corroborate the general intuition that the not-yet-treated are more similar to the treated than the never-treated group.

Table 3.6.3: Testing the parallel assumption as in eq. (3.6.1) and (3.6.3), unconditional case.

	Never-treated	**Not-yet-treated**
$g = 2004$	-0.003	-0.003
	(0.008)	(0.008)
$g = 2006$	-0.016^{*}	-0.013
	(0.008)	(0.007)
$g = 2007$	0.018^{**}	
	(0.007)	

Note: The table reports the mean difference in the yearly variation of teen employment between imputed untreated outcomes for the treated and control units, in years following treatment. Empirical test of Eq. (3.6.1) (never-treated) and Eq. (3.6.3) (not-yet-treated). Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. See Appendix 3.A.3 for details on the computation of the standard errors.

Table 3.6.4: Testing the parallel assumption as in eq. (3.6.2) and (3.6.4), conditional case.

	Never-treated	**Not-yet-treated**
$g = 2004$	-0.006	-0.005
	(0.008)	(0.008)
$g = 2006$	-0.017^{*}	-0.014
	(0.008)	(0.007)
$g = 2007$	0.016^{*}	
	(0.007)	

Note: The table reports the mean difference in the yearly variation of teen employment between imputed untreated outcomes for the treated and control units, in years following treatment. Empirical test of Eq. (3.6.2) (never-treated) and Eq. (3.6.4) (not-yet-treated). Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Despite the tests in section 3.6.1 not meeting our expectations, the ATE estimation results, notably, remain strikingly similar. In other words, even if moving from entirely different hypotheses to impute counterfactuals, and despite the absence of common trends for the never-treated group in the imputed Y(0) after treatment, MCP and CS's staggered DiD reach the same conclusions not only in terms of economic message, but also in terms of magnitude, and this is somewhat surprising.

This puzzling result raises new questions. In principle, concordance between MCP and DiD estimates would suggest that all the assumptions of the DiD framework are compatible with the data structure. This includes the parallel trends assumption, whose empirical foundation is partly supported by the tests in Section 3.6.2. However, the exercise in Section 3.6.1 finds statistically significant effects in the Y(0). This contradiction admits two interpretations: i) parallel trends are somehow imposed into the MCP imputation through fixed effects, even though they would not emerge organically under pure low-rank matrix completion, so that both estimators yield similar but biased estimates, or ii) the parallel trends assumption is valid, but MCP fails to recover reliable counterfactuals, possibly due to information leakage wherein post-treatment observations from untreated units inappropriately influence the imputation.

3.7 When to use MCP? Practical guidelines

This section addresses a practical question that applied researchers are likely to ask: is MCP worth the additional complexity, and if so, under what circumstances?

In our empirical application, DiD and MCP arrive at strickingly similar conclusions, which raises the question of when the greater computational and methodological demands of MCP are justified. Based on our analysis, we identify three conditions under which MCP is preferable to DiD.

First, when there is a good reason to believe the data exhibit a genuine low-rank structure. This is the structural assumption MCP is built around, therefore its presence should be carefully assessed in the data. Finding appropriate diagnostics for low-rank is a task we leave for future research.

Second, MCP is suitable when the data are rich and heterogeneous: a long time dimension T and substantial cross-unit variation are features that favour the algorithm's learning process and improve imputation quality. We expect the advantage over DiD to narrow in short panels or in datasets collecting more homogeneous units.

Third, MCP can be of assistance when unobservables pose serious identification challenges that cannot be resolved through standard parallel trends assumptions. A key advantage of MCP in this regard is that it does not require common trends to hold: this makes it particularly valuable in settings where the parallel trends assumption is difficult to defend — for instance, when treatment is correlated with unit-specific trends or when pre-treatment dynamics differ systematically across groups.

Conversely, when none of these conditions hold — and particularly when panel data are short, units homogeneous, or when the parallel trends assumption is reasonable, like it is probably the case with our CS replication, the additional complexity of MCP may not be worth it. In such cases, DiD remains a more transparent and credible alternative.

3.8 Conclusions and further research

In aggregate terms, MCP's findings are consistent with the original studies. However, the CATE estimation for the CS specification, reveals a positive coefficient for the initial period in the 2006 cohort. This finding may offer a reconciliation between the divergent conclusions of CK and CS, suggesting a temporal dimension to minimum wage effects: such policies may generate neutral or positive short-run impacts on low-wage workers, but these beneficial effects appear to attenuate over subsequent periods. In this sense, our findings align with Meer and West (2016), showing a slightly negative but increasing long-term impact on less-skilled labor. Importantly, these results are produced without imposing common-trend restrictions.

The MCP estimator proposed by Athey et al. (2021a) seems to offer an enticing data-driven alternative for estimating ATE's. In spite of its potential, it is not popular among applied researchers who still seem to prefer more common approaches, such as DiD. One of the aims of our work is to try to understand in more detail the performance of MCP, its limitations, and the actual possibility to use it in applied research. The results are encouraging, albeit mixed.

On the "plus" side, the estimator offers the following advantages. First, the researcher has the possibility to fully control for heterogeneous effects, even at the finest level of detail, *i.e.*, individual effects. This aspect is particularly valuable from a policy perspective. In our specific context, this feature allows for a deeper understanding of how minimum wage policies affect different labor market segments, addressing concerns raised by Neumark (2019). Second, MCP seems to yield more efficient estimates, producing ATE results similar to those of DiD but with substantially lower standard errors. Third, the fact that it provides explicit values to counterfactual quantities allows testing the validity of the identification assumptions for any causal panel estimator, not just DiD.

Despite these strengths, our study highlights that MCP seems to suffer from some drawbacks that limit its widespread adoption. In other words: there is a "minus" side. First, it remains complex to understand for applied researchers unfamiliar with the literature on matrix completion, and the computational burden may be substantial. Second, translating its assumptions into economic terms is challenging, making it harder to interpret compared to more traditional methods like DiD. Third, an open question is raised by our replication of CK estimates: under certain circumstances, MCP is unable to impute the $Y(0)$ values, returning only the fixed effects, which, in turn, coincide with the basic double difference of the DiD. Last, but not least, the application to CS has shown that despite the close replication of the average effects, the imputation of the counterfactual quantities may not be perfectly in line with expectations. This finding casts a shadow on the quality of the MCP imputation and/or on the validity of the parallel trend hypothesis in the original work.

Our analysis suggests that MCP is not a universally superior alternative to traditional methods, but rather a tool whose performance hinges critically on the structure of the data at hand. MCP appears most promising when: (i) both N and T are large; (ii) the number of treated units is relatively small compared to the matrix dimensions. Under these conditions, the algorithm has sufficient "training ground" to recover the underlying low-rank structure with precision. In such settings, MCP appears to deliver more efficient estimates. A third desirable condition is that (iii) the data exhibit a genuine low-rank structure, meaning the researcher has reasons to believe that there exists a structure of interrelations among the units of observation.

Conversely, MCP offers no advantage when the low-rank structure is unrecoverable. Our replication of CK illustrates this limitation: in a low-dimensional setting, MCP was unable to recover the low-rank structure $\mathbf{L}^*$, effectively collapsing to a standard DiD estimator. In such cases, researchers gain no benefit from MCP's additional complexity and should prefer simpler, more transparent methods like difference-in-differences.

An additional concern emerges from the replication of CS: the quality of the imputed counterfactuals may be questionable, as evidenced by our placebo tests in 3.6.1. Therefore, we encourage ensuring that the imputation satisfies internal consistency checks.

In sum, researchers should carefully assess whether their data structure aligns with MCP's requirements and, crucially, should validate imputation quality through placebo tests or similar diagnostic procedures. When data are limited or the imputation fails consistency checks, traditional methods remain the safer choice. Note that these indications are based on empirical observation rather than a systematic theoretical investigation of the algorithm's mathematical conditions. These conditions, and their implications in economic analysis, represent a promising avenue for future research.

Notably, how do these mechanisms compare to the assumptions of more established approaches like DiD? Are they economically meaningful? Under what mathematical conditions is MCP unable to impute the missing values? Addressing these questions will deepen our understanding of MCP's broader applicability in causal analysis. Matrix completion methods can also be leveraged within policy learning frameworks (Athey and Wager, 2021), offering a promising pre-processing step for ML-based approaches. By exploiting latent structures in the data to impute counterfactuals and recover treatment effects, MCP can complement covariate-based strategies and help identify heterogeneity driven by unobservables—such as time trends, reinforcement effects, or path dependence—that are not well captured by the observed covariates.

While ML has been used to identify potential treatment pools in minimum wage studies (Cengiz et al., 2022), it has yet to be applied to optimizing policy design. This could be a promising avenue for research, allowing for a deeper exploration of the interplay between minimum wage policies and local economic conditions, ultimately informing more effective labor market policies. Finally, future research should investigate MCP's performance in scenarios where common trends are openly violated.

Bibliography

A. Abadie, A. Diamond, and J. Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.

S. Athey and S. Wager. Policy learning with observational data. *Econometrica*, 89 (1):133–161, 2021.

S. Athey, M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730, 2021a.

S. Athey, M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi. MCPanel: Matrix completion methods for causal panel data models. `https://github.com/susanathey/MCPanel`, 2021b. R package.

B. Callaway and P. H. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230, 2021.

E. Candes and B. Recht. Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119, 2012.

E. J. Candes and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

D. Card and A. B. Krueger. Minimum wages and employment: A case study of the fast food industry in new jersey and pennsylvania, 1994.

D. Card and A. B. Krueger. Minimum wages and employment: a case study of the fast-food industry in new jersey and pennsylvania: reply. *American Economic Review*, 90(5):1397–1420, 2000.

D. Cengiz, A. Dube, A. Lindner, and B. Zipperer. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454, 2019.

D. Cengiz, A. Dube, A. Lindner, and D. Zentler-Munro. Seeing beyond the trees: Using machine learning to estimate the impact of minimum wages on labor market outcomes. *Journal of Labor Economics*, 40(S1):S203–S247, 2022.

M. Draghi. Eu competitiveness: Looking ahead - european commission. Technical report, 2024.

C. Dustmann, A. Lindner, U. Schönberg, M. Umkehrer, and P. Vom Berge. Reallocation effects of the minimum wage. *The Quarterly Journal of Economics*, 137 (1):267–328, 2022.

A. Fazio and T. Reggiani. Minimum wage and tolerance for high incomes. *European Economic Review*, 155:104445, 2023.

J. Kaddour, A. Lynch, Q. Liu, M. J. Kusner, and R. Silva. Causal machine learning: A survey and open problems. *arXiv preprint arXiv:2206.15475*, 2022.

K. Matsuzawa, D. I. Rees, J. J. Sabia, and R. Margolit. Minimum wages and teenage childbearing in the united states. *Journal of Applied Econometrics*, 2025.

R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, 11:2287–2322, 2010.

J. Meer and J. West. Effects of the minimum wage on employment dynamics. *Journal of Human Resources*, 51(2):500–522, 2016.

D. Neumark. The econometrics and economics of the employment effects of minimum wages: Getting from known unknowns to known knowns. *German Economic Review*, 20(3):293–329, 2019.

D. Neumark and W. Wascher. Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania: Comment. *American Economic Review*, 90(5):1362–1396, 2000.

O. Ropponen. Reconciling the evidence of card and krueger (1994) and neumark and wascher (2000). *Journal of Applied Econometrics*, 26(6):1051–1057, 2011.

P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Small Area Income and Poverty Estimates (SAIPE) Program. SAIPE datasets, 2024. URL `https://www.census.gov/programs-surveys/saipe/data/datasets.html`. Accessed: February 2025.

US Census Bureau - County Intercensal Datasets. Census data explorer, 2024. URL `https://data.census.gov/table`. Accessed: February 2025.

3.A Appendix

3.A.1 Glossary box

Table 3.A.1: Glossary Box

Term	**Definition**
Causal machine learning	An umbrella term for machine learning methods that formalize the data-generation process as a structural causal model Kaddour et al. (2022).
Coherence	A matrix property where the singular vectors are concentrated on a small subset of rows or columns, rather than being spread evenly across all entries. Incoherence, conversely, indicates that the singular vectors are uniformly distributed.
Cross-validation	A resampling method used to evaluate model performance and select hyperparameters in machine learning fitting. The procedure involves partitioning the dataset into K subsets (folds), iteratively training the model on $K-1$ folds and validating on the remaining fold. Performance metrics computed on validation folds provide an estimate of the model's out-of-sample predictive performance.
Machine learning	A bundle of statistical methods that learns patterns from data through iterative optimization processes. Rather than relying on explicit, pre-defined functional relationships, machine learning models adjust their parameters to minimize the prediction error on training data. The learning process typically proceeds through cross-validation.
Matrix completion	An algorithm developed by Candes and Plan (2010) that recovers missing entries of a partially observed matrix by exploiting its low-rank structure.
Mean squared error (MSE)	A widely used loss function and performance metric in machine learning that quantifies the average squared difference between predicted and actual values. For a dataset with n observations, MSE is computed as: $\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$ where y_i denotes the observed value and $\hat{y}_i$ the predicted value for observation i.
Nuclear norm	The sum of all singular values of a matrix. Contrary to matrix rank, the nuclear norm is continuous and, thus, it can be used as a convex relaxation of matrix rank.
Rank	In a matrix, the number of linearly independent rows or columns.
Singular values	The singular value decomposition is a matrix factorization $X = U\Sigma V^\top$ where U and V are orthogonal matrices containing left and right singular vectors, and Σ is a diagonal matrix with non-negative singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r \geq 0$. The singular values are the square roots of the eigenvalues of $X^\top X$.
Sparsity	In a matrix, sparsity indicates that the proportion of non-zero entries is small relative to the matrix dimensions.

3.A.2 Further information on the panel of covariates

In this section, Table 3.A.2 compares the mean of our variables with the ones in CS, taking the first panel year in order to reduce the time distance between the two datasets.

Table 3.A.2: Mean of covariates, comparison with the original CS dataset.

Variable	**CS**	**Our Panel**
Log of population	3.214	3.227
Fraction of white	0.856	0.882
Median income	32.605	36.838
Poverty rate	0.147	0.127

Note: The table presents the mean values of covariates for the first year in our panel (2001) alongside those reported in CS. Note that CS use time-invariant covariates coming from the 2000 County Data Book. These are measured in different years: the county population and the fraction of the population that is white date to the 2000 census, while median income and the poverty rate date to 1997.

Table 3.A.3 presents the yearly descriptive statistics in the reconstructed dataset.

Table 3.A.3: Summary Statistics by Year

	2001	2002	2003	2004	2005	2006	2007
lemp	5.83	5.75	5.71	5.66	5.67	5.69	5.70
medinc_covs	35.7	35.9	36.9	38.2	39.1	40.5	42.6
pop	3.29	3.30	3.31	3.32	3.32	3.33	3.34 pov_covs
0.121	0.122	0.120	0.123	0.134	0.135	0.131	
white_covs	0.889	0.888	0.887	0.885	0.884	0.882	0.880

3.A.3 Replication of CS - Standard errors

The standard error for cohort g in year t is given by

$$se_{g,t} = \frac{\mathrm{sd}_{g,t}}{\sqrt{n_{g,t}}} \tag{3.A.1}$$

where $\mathrm{sd}_{g,t}$ is the standard deviation of individual treatment effects, and $n_{g,t}$ is the number of treated units in cohort g at time t. *Mutatis mutandis*, similar quantities can be obtained for different cohort/time aggregations.

3.A.4 Replication of CS - Comparison with never-treated

We report here the results of the same exercise in Section 3.4, Tables 3.4.1 and 3.4.3, using the never-treated as the control group. In general, we do not observe substantial differences.

Figure 3.A.1: Minimum wage group-time average treatment effects - Unconditional parallel trends. Control group: never-treated

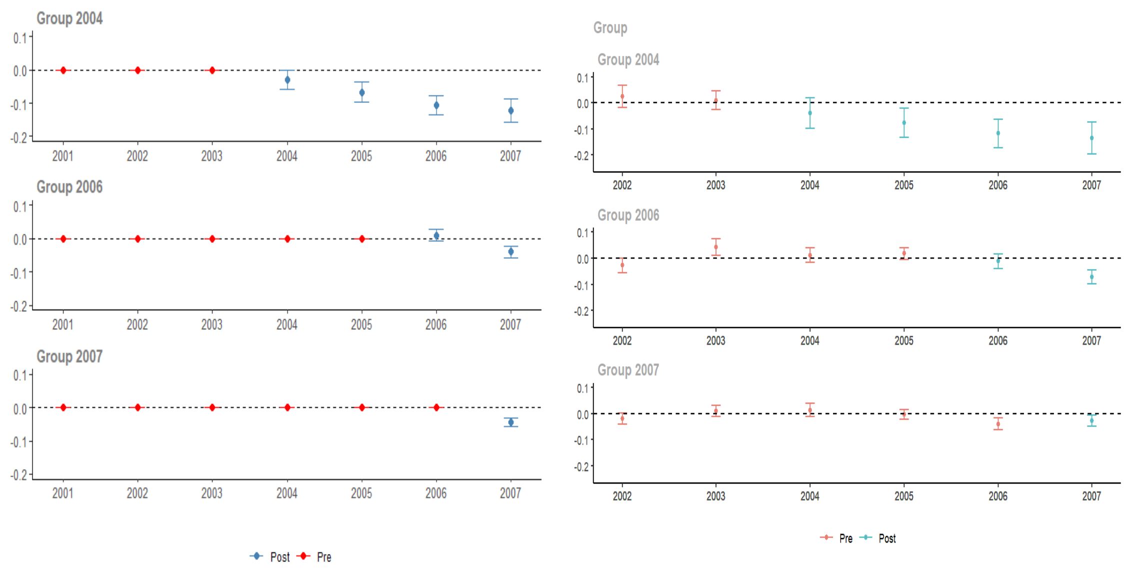

Note: The figure compares group-time treatment effects obtained using MCP (left) with those from CS (right) in the case of unconditional parallel trends. The top row corresponds to states that increased their minimum wage in 2004, the middle row to states that did so in 2006, and the bottom row to states that implemented the increase in 2007. Blue lines represent treatment effect point estimates and the relative 95% confidence bands. Red lines indicate point estimates and 95% confidence bands for pre-treatment periods. For the parallel trends assumption to hold, these estimates should include zero. Note that, in the MCP framework, pre-treatment effects are set to zero by construction.

Table 3.A.4: Cohort-specific effects and event study, unconditional case. Control group: never-treated

	MCP	**Staggered DiD**
Cohort-Specific Effects		
$g = 2004$	-0.081^{***}	-0.091^{***}
	(0.004)	(0.019)
$g = 2006$	-0.014^{***}	-0.047^{***}
	(0.002)	(0.008)
$g = 2007$	-0.045^{***}	-0.028^{***}
	(0.007)	(0.007)
Event Study		
$e = 0$	-0.026^{***}	-0.027^{***}
	(0.001)	(0.006)
$e = 1$	-0.053^{***}	-0.071^{***}
	(0.001)	(0.009)
$e = 2$	-0.106^{***}	-0.125^{***}
	(0.015)	(0.021)
$e = 3$	-0.122^{***}	-0.136^{***}
	(0.018)	(0.023)

Note: The table compares the estimated effects of the minimum wage on teen employment using CS staggered DiD estimator versus MCP. In the latter case, the treatment effect is obtained as the mean difference between the counterfactual outcome $Y(0)$, as imputed by MCP in the absence of covariates, and the observed outcome $Y(1)$ for treated counties. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A.5: Cohort-specific effects and event study, conditional case. Control group: never-treated

	MCP	**Staggered DiD**
Cohort-Specific Effects		
$g = 2004$	-0.082^{***}	-0.095^{***}
	(0.004)	(0.020)
$g = 2006$	-0.014^{***}	-0.047^{***}
	(0.002)	(0.009)
$g = 2007$	-0.049^{***}	-0.030^{***}
	(0.007)	(0.007)
Event Study		
$e = 0$	-0.023^{***}	-0.028^{***}
	(0.001)	(0.005)
$e = 1$	-0.053^{***}	-0.074^{***}
	(0.001)	(0.009)
$e = 2$	-0.030^{***}	-0.129^{***}
	(0.001)	(0.021)
$e = 3$	-0.067^{***}	-0.141^{***}
	(0.001)	(0.023)

Note: the table compares the estimated effects of the minimum wage on teen employment using CS staggered DiD estimator versus MCP, including the covariates reported in Table 3.4.2. In the latter case, the treatment effect is obtained as the mean difference between the counterfactual outcome $Y(0)$, as imputed by MCP in absence of covariates, and the observed outcome $Y(1)$ for treated counties. The section "Cohort-Specific Effects" reports the ATEs by the timing of the minimum wage increase, where g denotes the first year a county is treated. The "Event Study" section presents aggregated treatment effects by the duration of exposure to the minimum wage increase, with e indicating the number of years since treatment began. Standard errors are reported in brackets, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3.A.2: Minimum wage group-time average treatment effects - Conditional parallel trends. Control group: never-treated

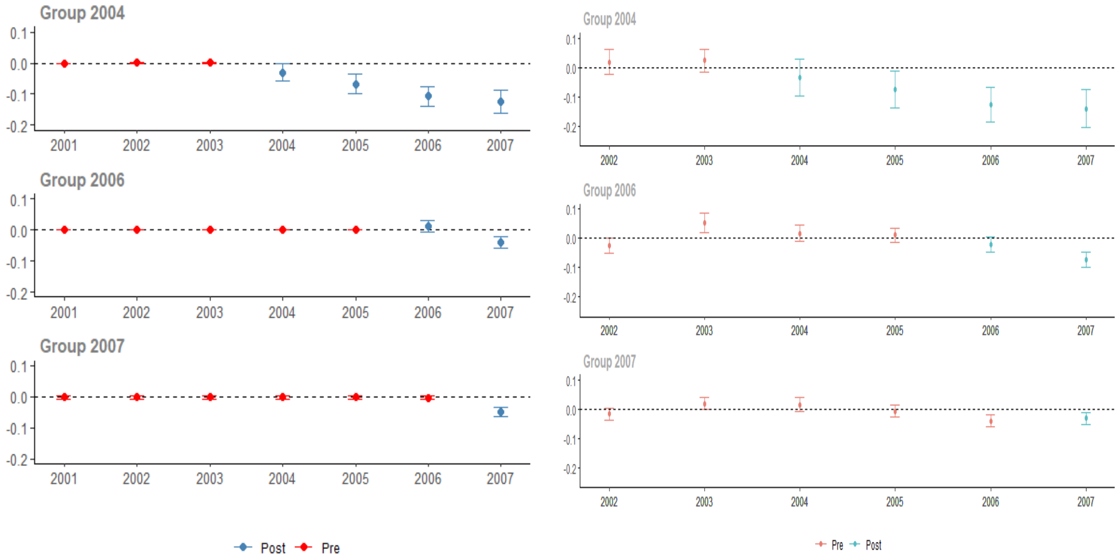

Note: This figure compares group-time treatment effects obtained using MCP (left) with those from CS (right) in the case of conditional parallel trends. The top row corresponds to states that increased their minimum wage in 2004, the middle row to states that did so in 2006, and the bottom row to states that implemented the increase in 2007. Blue lines represent treatment effect point estimates and the relative 95% confidence bands. Red lines indicate point estimates and 95% confidence bands for pre-treatment periods. For the parallel trends assumption to hold, these estimates should include zero. Note that, in the MCP framework, pre-treatment effects are set to zero by construction.

Acknowledgements

During this PhD I had the privilege of working with people that took a real interest in my professional development. To Elena, Valeria, Giuliano, and Marco, for their patience and their guidance, the greatest thanks. A special thanks to the Reviewers who took the time to give their precious feedback to this thesis. Thanks to all my colleagues for contaminating me and my work with countless inspiring ideas.

On a more personal note: Thanks to my truly amazing friends, scattered all around the globe, for always being emotionally within reach. Thank you to my family for always being interested and understanding. Thanks to Mert for being the best partner and my proudest supporter.

Angela Zanoni,

Rome, March 2026