A Counterfactual Approach to Energy Poverty Mitigation: A Case Study for Australia (Preliminary Report)

Diogo Nuno Freitas[1,2,3,*], Eduardo Fermé[1,2] and Santiago Budría[4,5]

[1] Universidade da Madeira, Campus Universitário da Penteada, Portugal
[2] NOVA Laboratory for Computer Science and Informatics (NOVA LINCS), Portugal
[3] Interactive Technologies Institute (ITI/LARSyS and ARDITI), Portugal
[4] Universidad Antonio de Nebrija, Spain
[5] Centro de Estudos de Economia Aplicada do Atlântico (CEEAplA), Portugal

Abstract

Energy poverty is a persistent global issue where households lack access to adequate energy services. Research often focuses on contemporaneous factors, overlooking the predictive power of long-term socioeconomic trajectories. This study addresses two questions: How do past socioeconomic conditions affect future energy poverty? And what are the minimal interventions that could prevent a household from becoming energy-poor? We aim to shift the focus from reactive mitigation to proactive prevention by developing a framework that forecasts risk and identifies actionable pathways to avoid it. Using seven years of Australian longitudinal data, we train a machine learning classifier to predict energy poverty in the following year. We then apply counterfactual analysis to identify minimal, interpretable changes that alter the predicted outcome. The model successfully predicts future energy poverty with a ROC AUC of 70.01%. The counterfactual analysis consistently reveals that modest increases in household income, often less than 5%, are the most effective single intervention. Other factors, such as decreases in energy prices and reductions in unemployment, also contribute to preventing energy poverty, often in combination with income gains. "What-if" scenarios suggest that external shocks, like a sudden rise in energy prices or job loss, can be offset by small, timely income adjustments. The effectiveness of these changes, whether recent or past, highlights the importance of long-term financial stability. Also, it enables proactive policy by identifying at-risk households a year in advance and suggesting targeted interventions as a more efficient alternative to broad measures.

Keywords

energy poverty, predictive models, machine learning, counterfactuals, public policy, decision support systems

1. Introduction

Energy poverty remains a pressing global challenge, affecting the quality of life for millions of people. It refers to a situation where households cannot access or afford adequate, reliable, and clean energy services for their daily needs. Despite international efforts to address this issue, approximately 750 million people still do not have access to electricity, and over 2 billion lack access to clean cooking fuels [1].

The consequences of energy poverty extend far beyond basic comfort. Limited access to energy reduces educational opportunities, restricts access to jobs [2], and negatively affects well-being [3] and health [4]. For instance, Nawaz [5] showed that households facing energy poverty are 9 to 13% more likely to experience poor health compared to those who are not energy-poor. Poor indoor conditions, such as a lack of adequate heating, have a considerable impact on both physical and mental health. Long periods of low indoor temperatures have been linked to higher rates of illness and death [6, 7], as well as worsening mental health issues [8]. Vulnerable groups, including infants, are especially at risk for respiratory diseases [9, 10]. There is also evidence that energy poverty can increase the risk of developing metabolic disorders like diabetes [11].

Workshop on the Foundations and Future of Change in Artificial Intelligence (FCAI 2025) at the 28th European Conference on Artificial Intelligence (ECAI 2025, October 25–30), Bologna, Italy

*Corresponding author.
diogo.freitas@staff.uma.pt (D. N. Freitas); eduardo.ferme@staff.uma.pt (E. Fermé); sbudria@nebrija.es (S. Budría)
0000-0002-2351-8676 (D. N. Freitas); 0000-0002-9618-2421 (E. Fermé); 0000-0001-5537-4059 (S. Budría)

Many factors contribute to energy poverty. These include rising energy prices, economic instability, and the presence of inefficient or poor-quality housing [12]. Research has extensively examined the socioeconomic aspects of energy poverty, with the aim of identifying households at risk and designing targeted policy solutions [13, 14, 15]. However, most studies focus mainly on contemporaneous socioeconomic conditions and assume that only current variables are needed to explain energy poverty. As a result, there is limited understanding of the long-term impact that past socioeconomic factors have on future energy poverty.

Besides identifying households at risk, it is equally important to consider what actions can help prevent energy poverty. In other words, if a household is predicted to become energy-poor, what actions could be taken, or could have been taken, to prevent this outcome? Ideally, these actions should require minimal intervention but provide the greatest benefit in terms of prevention.

Our work addresses these two gaps in the existing research: First (Contribution 1), we explore how past factors, such as income, energy prices, regional circumstances, and other socioeconomic variables, affect *future* energy poverty outcomes. We apply a machine learning algorithm that uses household-level data from seven previous years to predict energy poverty status in the following year. Second (Contribution 2), we use counterfactual analysis on the model developed in Contribution 1 to identify specific, minimal, and human-interpretable changes that could help a household shift from being energy-poor to non-energy-poor. These changes might result from personal actions or targeted policies. We refer to these as *chains of minimal changes*. We also present "what-if" scenarios that identify shocks which increase poverty risk and the minimal counterfactual response required to mitigate that risk.

Contribution 1 and Contribution 2 are complementary, and allow us to move from reacting to energy poverty after it occurs to taking proactive measures to prevent it. Our approach identifies households at risk one year in advance and provides practical suggestions for actions that require the least effort but have the greatest impact on preventing energy poverty. In this way, we bridge the gap between prediction and real-life intervention, providing practical guidance for policymakers, community organizations, and individuals seeking to reduce the long-term effects of energy poverty.

To conduct our analysis, we use the Household, Income and Labour Dynamics in Australia (HILDA) Survey[1], which is a nationally representative, long-term study that tracks the same individuals and households in Australia from 2007 to 2021. The survey collects data on income, jobs, education, health, and family life, providing a strong foundation to study how energy poverty changes over time. For measuring "energy-poverty", we employ the Multidimensional Energy Poverty Index (MEPI) [16], which combines both objective (expenditure-based) and subjective (self-assessed) indicators of energy poverty into a single, comprehensive metric.

2. Literature Review

Energy poverty can be defined as a household's inability to afford or access energy services needed to support adequate living conditions and human development. While conceptual definitions of energy poverty have been the subject of extensive discussion in the literature (for an overview see Sy and Mokaddem [17]), the focus has generally been on the inability of households to afford and have access to adequate energy services.

Based on the literature, energy poverty is a complex phenomenon stemming from a wide range of factors. These include macroeconomic conditions like GDP, governance, and a country's energy mix [18, 19, 20, 21, 22, 23], as well as household-level characteristics such as income, dwelling type, and size [24, 14, 25]. Individual attributes like educational attainment, health status, age, and employment also play a significant role [26, 4, 27, 2, 28, 29, 13]. Furthermore, factors like spatial disparities, cultural behaviors, and energy subsidies add to the complexity [30, 31, 32].

Given the multi-layered nature of these determinants, a recent body of literature has introduced machine learning techniques to predict energy poverty outcomes. Evidence based on a Extreme Gradient

[1] https://melbourneinstitute.unimelb.edu.au/hilda

Boosting (XGBoost) framework to predict the risk of experiencing energy poverty in the Netherlands identifies income, house value, and house ownership as the main drivers of energy poverty [33]. In a similar setting, and based on 11 European countries, income, household size, and floor area were consistent predictors [34]. Evidence based on an Random Forest (RF) classifier across the European Union uncovers household- and country-level predictors like dwelling conditions, energy efficiency, and gas supplier switching rates [35].

While the previous studies are based on a single energy poverty indicator, other studies define a multidimensional energy poverty index similar to ours [16]. These studies have shown that in Asian and African countries, wealth, marital status, and residence attributes are significant predictors of poverty [29]. Recent research has further advanced these methodologies by employing ensemble models, such as XGBoost, combined with RF and Artificial Neural Networks (ANN), revealing the critical importance of education and food security indicators in determining energy poverty [36].

3. Materials and Methods

We use the HILDA Survey, a comprehensive, nationally representative longitudinal study that examines the economic, social, and demographic dynamics of Australian households. Initiated in 2001 and conducted annually, it tracks individuals and households over time, providing important information on income, labor market activities, health, education, and family relationships, among other factors. The original 2001 sample included approximately 7,600 households and 13,000 individuals, with periodic updates to account for attrition.

The variables used to model energy poverty are described in Table 3 (Section C), and include labor market indicators (such as part-time employment rate, unemployment rate, and labor force participation), economic measures (GSP per capita and energy prices), and household characteristics (income, household size, and region). We also include individual information such as age, years of education, marital status, employment, health, and the presence of children in the household.

In this study, we approach energy poverty as a classification problem. Let $\mathcal{X} \subseteq \mathbb{R}^d$ be the feature space, and denote an element (feature vector) by $\mathbf{x} \in \mathcal{X}$. We employ a machine learning model $f : \mathcal{X} \to [0, 1]$, where $f(\mathbf{x})$ predicts the likelihood that a household, characterized by the features $\mathbf{x}$ from the past seven years, will experience energy poverty in the upcoming eighth year. This prediction relies on the MEPI indicator, which serves as our dependent variable.

Outcome—The MEPI Indicator: The MEPI variable captures both expenditure-based and subjective dimensions. The expenditure-based measures include the 2M, TPR, and LIHC indicators. The subjective dimension is reflected by two self-assessed indicators: the household's inability to pay for heating due to a lack of money (Heat), and the inability to pay electricity, gas, or telephone bills on time (Arrears).

The MEPI index is calculated as follows: Let $J = \{1, \dots, m\}$ be the set of $m = 5$ poverty indicators. Let $\mathcal{I}$ be a set of individuals, with element i, $i \in \mathcal{I}$, and $\mathcal{T}$ be a set of time periods, $t \in \mathcal{T}$, representing a specific moment when the survey was conducted. Let EP_{ijt} denote the status of the i-th individual in the j-th indicator during period t. If an individual i is poor under indicator j in the period t, then EP_{ijt} takes the value of one, and zero otherwise. Following the family of indexes typically described in the literature on material deprivation [37], individual i's weighted poverty score is given by:

$$\mathrm{MEPI}_{it} = \sum_{j \in J} w_j \mathrm{EP}_{ijt}, \quad \forall\, i \in \mathcal{I},\ t \in T_i,\ T_i \subseteq \mathcal{T}, \tag{1}$$

where w_j denotes the weight assigned to the poverty indicator j, with $\sum_{j \in J} w_j = 1$. Hence, the MEPI_{it} indicator ranges from 0 to 1 and captures the percentage of dimensions in which the individual is deprived.

Although it is common to give the same importance (w) to each indicator, we focus more on the indicators where deprivation is rare. This method is known as the Frequency-Based Weighting Approach [38]. The weight given to an indicator is proportional to the percentage of individuals *not* classified as poor under that specific indicator within a particular state. In other words,

$$w_j = \frac{(1 - n_j)}{\sum_{j \in J}(1 - n_j)}, \tag{2}$$

where n_j is the proportion of poor individuals in dimension j. This choice is based on the belief that lacking access to everyday items should be considered a more significant indicator of deprivation than lacking access to less common items. The weights are calculated separately for each wave.

In the context of this work, an individual i is regarded as energy poor if $\text{MEPI}_{it} > 0$ [16]. This means that a household is considered "energy-poor" if it is deprived in any one of the five dimensions. Therefore, the variable of interest in this study, which indicates whether a person is energy-poor, is binary: it is 1 if the person is energy-poor, and 0 if not.

Variables and Data Preparation: We model energy poverty at time $T = 8$ as a function of socioeconomic and demographic characteristics observed in periods $T - 1, \ldots, T - 7$. We exclude the period T to ensure that our forecasting relies entirely on historical data. Our focus is on the role of past factors. Including contemporaneous variables could potentially mask the effects of lagged factors, especially if there is autocorrelation in the data. More importantly, the inclusion of contemporaneous variables may introduce reverse causality between energy poverty and socio-demographic characteristics such as health and education [39, 4]. By considering only past variables, we eliminate the risk that current energy poverty influences these characteristics.

To capture the temporal dynamics of the variables, we created lagged features, which serve as the input to the predictive models. Generically, for each original feature, we obtained new features representing its values from each of the previous years.

We then split the dataset into training, validation, and test subsets to facilitate model development and evaluation. Out of the 7,977 participants in our dataset, 6,382 (80%) were randomly selected for training and validating the predictive models, while the remaining 1,595 participants (20%) were included in the test set. The test set was held out and used exclusively to evaluate the final performance of the models, providing a fair estimate of their forecasting accuracy. Moreover, the dataset is split by individual to prevent data leakage.

Before training the model, we standardized the data to ensure consistency and reliability in our modeling process. The standardization parameters were estimated solely from the training set to avoid information leakage. Specifically, we removed the median and scaled the data using the interquartile range, as described in [40]. These parameters were then applied to transform the training, validation, and test data.

Model Development: We treat the energy poverty forecasting task as a classification problem. Specifically, households are classified as energy-poor depending on whether their MEPI is greater than 0 (cut-off point). To model the relationship between the socioeconomic and demographic factors and the MEPI indicator, we used a balanced bagging classifier.

A balanced bagging classifier [41] is an ensemble technique that combines the predictions of multiple base models, in our case, decision trees, to improve the robustness and accuracy of the outcomes. In order to further refine the modeling approach, we implemented the classifier in an One-vs-the-Rest (OvR) binary classification framework [42].

We optimized the hyperparameters of our classifier using a grid search. For details on the specific hyperparameters and grid configurations, see Section A. We employed 5-fold cross-validation on the training dataset to ensure the robustness of the hyperparameters across different data splits, selecting the best set based on the highest Receiver Operating Characteristic - Area Under Curve (ROC AUC) score.

The final model was trained on the complete training set using the identified optimal hyperparameters and subsequently evaluated on a held-out test set of 1,595 participants.

In addition to the balanced bagging classifier, we tested two other class-imbalance ensembles. Namely, we benchmarked the random under-sampling boosting [43] and the easy ensemble [44]. All models used the same feature set, preprocessing pipeline, and 5-fold cross-validation on the training partition, with hyperparameters tuned via grid search to maximize ROC AUC. However, the balanced bagging classifier was the model that achieved the highest ROC AUC. Therefore, we chose the balanced bagging classifier as our final model.

Counterfactual Explanations: Counterfactuals are a post-hoc means to understand and explain the model predictions [45]. In the context of this work, counterfactuals are used to generate alternative household profiles, and in combination with the machine model f, help determine if one change (e.g., an increase in household income) or a set of changes in the user profile can increase or decrease the likelihood of energy poverty. Additionally, they are also used to identifying the minimum response needed to mitigate the risk of becoming energy poor after a shock that increases poverty.

Our focus is on *minimal* changes in order to ensure that the counterfactual recommendations remain practical and actionable, allowing individuals to make small adjustments that could have a meaningful impact on their future energy poverty risk. The minimal changes align with the concept of *proximity*, where the suggested profile modifications are as close as possible to the individual's original state. Besides proximity, another relevant aspect is *diversity*, meaning that multiple plausible pathways out of energy poverty profiles should be considered.

In this work, counterfactual explanations are formulated as a constrained optimization problem, as proposed by Mothilal et al. [46]. Given an individual feature vector $\mathbf{x}$ with $f(\mathbf{x}) = 1$ (classified as energy poor at $T = 8$), we search for a new feature vector $\mathbf{x}'$ such that $f(\mathbf{x}') = 0$. The goal is to find $\mathbf{x}'$ that is as close as possible to $\mathbf{x}$, so only minimal and realistic changes are required.

Formally, if k is the total number of counterfactual examples to be generated, then the set of counterfactuals $\{\mathbf{x}'_i\}_{i=1}^k$ is obtained by solving the following optimization problem:

$$\begin{aligned}\{\mathbf{x}'_i\}_{i=1}^k = \arg\min_{\mathbf{x}'_i} \Bigg\{ & \frac{\lambda}{k} \sum_{i=1}^{k} \sum_{j \in \mathcal{F}} m_j w_j \cdot d_j(x_j, x'_{i,j}) \\ & + \frac{1}{k} \sum_{i=1}^{k} \ell(f(\mathbf{x}'_i), 0) - \gamma D(\{\mathbf{x}'_i\}_{i=1}^k) \Bigg\},\end{aligned} \tag{3}$$

where the first term measures the weighted distance between the original instance $\mathbf{x}$ and each counterfactual $\mathbf{x}'_i$. The per-feature distance $d_j(x_{i,j}, x'_{i,j})$ measures the change for feature $j \in \mathcal{F}$. The term m_j is a binary mask, where $m_j = 1$ if feature j is mutable and $m_j = 0$ otherwise. The weight w_j reflects the difficulty of changing feature j and is defined as the inverse of the median absolute deviation (MAD).

The second term is a hinge-style loss on the logit of the predicted probability:

$$\ell\big(f(\mathbf{x}'_i), 0\big) = \max\!\big(0,\ 1 + \mathrm{logit}\big(f(\mathbf{x}'_i)\big)\big), \quad \text{where } \mathrm{logit}(p) = \ln \frac{p}{1-p}. \tag{4}$$

This penalizes any counterfactual $\mathbf{x}'_i$ whose predicted probability $f(\mathbf{x}'_i)$ is not confidently below the decision boundary for class 0.

The third term promotes diversity among the generated counterfactuals by maximizing their pairwise distances. To avoid overloading notation, we distinguish between a per-feature distance $d_j(\cdot,\cdot)$ and a vector-level distance $\Delta(\cdot,\cdot)$:

$$D\big(\{\mathbf{x}'_i\}_{i=1}^k\big) \ = \ \sum_{i=1}^{k} \sum_{a=i+1}^{k} \Delta(\mathbf{x}'_i, \mathbf{x}'_a), \tag{5}$$

with

$$\Delta(\mathbf{x}, \mathbf{x}') = \sum_{j \in \mathcal{F}} d_j(x_j, x'_j), \tag{6}$$

and, following Mothilal et al. [46], the per-feature distance is

$$d_j(u,v) = \begin{cases} |u-v|, & j \in \mathcal{F}_{\text{cont.}}, \\ \mathbb{1}[u \neq v], & j \in \mathcal{F}_{\text{cat.}} \end{cases}. \tag{7}$$

That is, d_j corresponds to the L^1 change for continuous features, while categorical feature modifications are penalized uniformly to discourage unnecessary changes [46].

The parameters λ and γ control the trade-off between proximity, classification change, and diversity. In our analysis, we set $\lambda = 5$ and $\gamma = 2.5$. We used a genetic algorithm to solve Equation (3) and generated $k = 50$ counterfactuals for each of the prototype profiles described below.

In our analysis, we set $\lambda = 5$ and $\gamma = 2.5$, based on preliminary tests with alternative parameter values. These tests indicated that a lower λ resulted in counterfactuals needing larger changes in features, making them unrealistic for households. On the other hand, higher λ values limited proximity too much. Similarly, varying γ showed that values that were too small limited diversity, and values that were too large reduced interpretability. The selected values offered the best trade-off between minimal, realistic changes and sufficient diversity.

To encode immutability explicitly, let $\mathcal{M} \subseteq \mathcal{F}$ denote the set of mutable features and $\mathcal{I} = \mathcal{F} \setminus \mathcal{M}$ the set of immutable features. We impose the following constraints, for all $i = 1, \ldots, k$:

$$x'_{i,j} = x_j, \qquad \forall\, j \in \mathcal{I}. \tag{8}$$

To ensure realism and actionability, we also restrict changes in continuous mutable features to lie within a $\pm 5\%$ band around their original values. Let $\rho = 0.05$. Then, for all $i = 1, \ldots, k$ and all $j \in \mathcal{F}_{\text{cont.}} \cap \mathcal{M}$,

$$(1-\rho)\, x_j \;\leq\; x'_{i,j} \;\leq\; (1+\rho)\, x_j. \tag{9}$$

We note here that counterfactual explanations are closely related to, and often considered a form of, contrastive explanations [47]. The literature, however, draws a subtle distinction between the two. Contrastive explanations primarily identify the features responsible for a classification (e.g., "The household is energy-poor *because of* its low income") [48], whereas counterfactuals identify the minimal changes that would alter the outcome (e.g., "The household would *not be* energy-poor *if* its income were increased by X") [45]. As our work focuses on providing actionable recourse by showing *what must change* to mitigate energy poverty, we have adopted the term "counterfactual explanation" as it most precisely describes our objective.

Prototype Profiles: To obtain a small set of representative borderline cases, we first define borderline profiles as those for which the model predicts class 1 with low confidence. Specifically, when the predicted probability $f(\mathbf{x})$ lies in the range $[0.5,\, 0.5 + \varepsilon]$, with $\varepsilon = 0.05$. That is, the borderline sample set is defined as

$$B = \{\mathbf{x}_i \mid f(\mathbf{x}_i) \in [0.5,\, 0.55] \text{ and } y_i, \hat{y}_i = 1\}, \tag{10}$$

where $\hat{y}_i = \arg\max_y f(\mathbf{x}_i)$ denotes the predicted class label.

We then compute the empirical mean of the set $B = \{\mathbf{x}_i\}_{i=1}^n$ as

$$\overline{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_i. \tag{11}$$

The first prototype p_1 is chosen as the observation in B closest to this mean, i.e.,

$$p_1 = \arg\min_{\mathbf{x} \in B} \|\mathbf{x} - \overline{\mathbf{x}}\|_2 \,. \tag{12}$$

Subsequent prototypes $\{p_j\}_{j=2}^m$ are selected using the farthest-first traversal rule:

$$p_j = \arg\max_{\mathbf{x} \in B} \left[\min_{1 \leq L < j} \|\mathbf{x} - p_L\|_2 \right], \quad j = 2, \ldots, m, \tag{13}$$

ensuring that each new prototype maximizes its minimum Euclidean distance to the already-selected set. This procedure yields m diverse, actual observations from the borderline region, with p_1 capturing the central tendency and the remaining $\{p_j\}_{j\geq 2}$ spanning the range of variability in feature space.

In our application, we selected $m = 3$ prototypes.

4. Results

Predictive Model (Contribution 1): A grid search was conducted to optimize the balanced bagging classifier's configuration for $T = 8$. The best setup included 100 estimators with bootstrapping of features but not samples. Each estimator sampled 50% of the data, and the sampling strategy ensured an equal representation of energy-poor and non-energy-poor instances. Replacement was used in the resampling process.

Our approach predicts energy poverty status in the eighth year for each household, using information from the previous seven years as input features. The model achieved a ROC AUC of 70.01%, which indicates that the model can discriminate between energy-poor and non-energy-poor households across varying decision thresholds. For class-specific metrics, sensitivity was 73.25%, meaning that the model correctly identified most energy-poor cases. The specificity of 66.77% indicates an acceptable rate of correctly classifying non-energy-poor households.

These results show that the model favors sensitivity. This is useful for policy design, where identifying the most energy-poor households is more important than avoiding the misclassification of some non-energy-poor households. High sensitivity supports early interventions from a prevention perspective, while reasonable specificity keeps misclassifications within an acceptable range.

Overall, our results support the use of balanced ensemble methods with appropriate sampling and parameter selection to model energy poverty using historical data.

Counterfactuals (Contribution 2): To interpret the predictive model and develop actionable insights, we conducted a counterfactual analysis on three representative prototype profiles, as described in Section 3. Each prototype represents a borderline energy-poor household, where the model predicts energy poverty with relatively low confidence. It is important to note that, for interpretability, all features were transformed back to their original scale by applying the inverse of the preprocessing transformation. Moreover, the immutable features include demographic attributes such as age and age group classifications, marital status (including being married, divorced, or widowed), the presence of a disability, the presence of children in the household, and the household size.

For each prototype, we generated 50 diverse counterfactual examples that change the predicted class from energy-poor to non-energy-poor (i.e., from class 1 to class 0), while requiring only minimal and realistic changes to the input features. All generated counterfactual profiles satisfy the classification constraint $f(\mathbf{x}') = 0$.

We examined the full set of generated counterfactuals and analyzed the types of changes required to shift a household's classification. Among these, we identified a subset of counterfactuals that involved only a single actionable modification. In all such one-action counterfactuals, the feature that was changed was household income. That is, all one-action classification changes across the three prototypes were achieved solely through an increase in income. No other individual feature was sufficient to produce a change in predicted energy poverty status when modified in isolation.

The timing of these interventions varied across different counterfactuals. Successful income changes tended to occur in more recent years, most frequently in $T = 6$ and $T = 7$, but also appeared in earlier years, including $T = 1$ and $T = 2$. This indicates that while recent changes in income have strong predictive power, improvements in earlier years can also contribute meaningfully to reducing energy poverty risk. The scale of required income changes was generally modest, but varied slightly depending on the year. In year $T = 1$, the required increase ranged from 3.18% to 3.81%; in $T = 2$, from 2.76% to 4.03%; in $T = 6$, from 1.83% to 4.93%; and in $T = 7$, from 1.30% to 4.82%. These ranges show that effective interventions can occur across multiple time points and typically demand income increases of less than 5%. This suggests that relatively small financial improvements, whether recent or distributed over a longer time horizon, can be sufficient to prevent energy poverty for households at the margin.

Figure 1 shows several connected counterfactual (with more than one action) pathways that lead from an energy-poor to a non-energy-poor household profile at time $T = 8$. Each path represents a sequence of minimal changes to selected features in earlier years. The figure illustrates the range of time steps and features through which this transition can be achieved.

Across all paths, an increase in household income appears as a consistent component. Income changes are present at multiple time points, including $T = 1$, $T = 2$, $T = 6$, and $T = 7$. The required increases range from 2.26% to 4.36%, confirming earlier results that show small changes in income are effective in changing the classification of households that are near the energy poverty threshold.

In addition to income, other features involved in the transitions include energy price, unemployment rate, and part-time employment rate. Decreases in energy prices at $T = 1$ and $T = 3$, and reductions in unemployment at $T = 4$ and $T = 7$, contribute to multiple successful pathways. Several paths combine income changes with labor market improvements, such as a decrease in the part-time employment rate at $T = 6$ or $T = 7$.

The timing of the changes is also significant. While most changes occur in the final time step ($T = 7$), effective changes in earlier years, especially $T \leq 4$, also appear. This suggests, similar to the one-action counterfactuals, that both recent and earlier interventions can influence future energy poverty status.

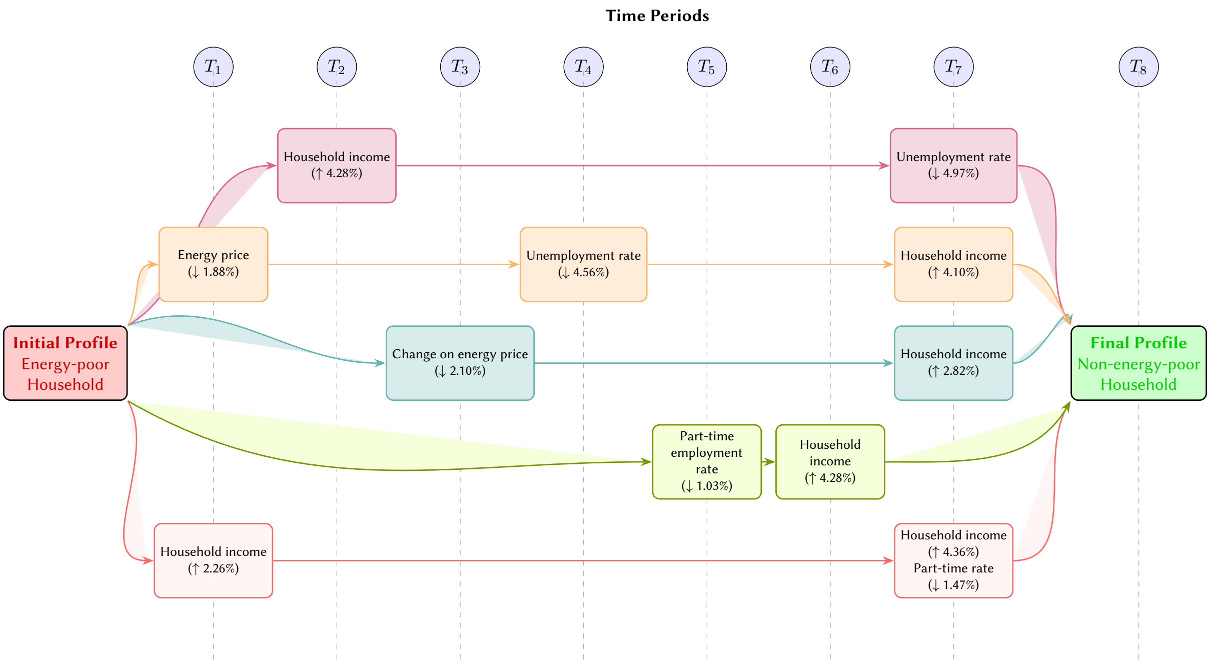

Figure 1: Connected pathways showing updated cases for transitioning from energy poor to non-energy-poor profile, with specific feature changes at each time step.

Shocks (Contribution 2): To further explore how households may avoid energy poverty under different conditions, we present in Table 1 a set of scenarios based on selected multi-feature counterfactuals. These scenarios describe situations in which a change in a macroeconomic or household-level factor (in other words, a shock) increases the risk of energy poverty. For each situation, we identify the minimal change that the model associates with successfully avoiding this outcome.

In one scenario, energy prices rise by 3.03% in year $T = 4$. The model finds that a compensating income increase of 2.92% in year $T = 7$ is required to prevent the household from entering energy poverty. Similarly, when energy prices increase by 3.16% in year $T = 6$, a larger income adjustment of 4.82% in $T = 7$ is needed to offset the rising cost burden.

Changes in labor force participation also play a role. When the total labor force participation rate

Table 1
Counterfactual Scenarios for Avoiding Energy Poverty. The table shows the initial shock that increases poverty risk and the minimal counterfactual response required to mitigate that risk, as identified by the model.

	Shock (Risk Increase)				Required Response (Mitigation)		
Scenario	**Factor**	**Change**	**Time**		**Factor**	**Change**	**Time**
1	Energy Price Increase	3.03%	$T = 4$	→	Income Increase	2.92%	$T = 7$
2		3.16%	$T = 6$			4.82%	$T = 7$
3	Labor Force Drop	−2.15%	$T = 5$	→	Income Increase	3.23%	$T = 7$
4	Energy Price Volatility	3.46%	$T = 2$	→	Income Increase	3.73%	$T = 7$
5	Becomes Unemployed	–	$T = 2$	→	Income Increase	1.83%	$T = 6$

drops by 2.15% in year $T = 5$, the model identifies a 3.23% income increase in $T = 7$ as sufficient to prevent energy poverty. Shifts in energy cost volatility show a similar pattern. When the energy price change rate increases by 3.46% in year $T = 2$, future income must increase by 3.73% in $T = 7$ to avoid a negative classification.

A final scenario involves a complete loss of employment, with the individual becoming fully unemployed in year $T = 2$. In this case, a modest income increase of 1.83% in year $T = 6$ is sufficient to counterbalance the risk introduced by the unemployment shock.

These scenarios show that while the causes may differ, income remains the most frequent feature requiring adjustment. Small income increases, aligned with observed external conditions, are often enough to reduce the likelihood of energy poverty according to the learned model.

5. Discussion

Our study suggests a novel approach to energy poverty prevention by combining a machine learning algorithm with counterfactual analysis. We used a machine learning algorithm to predict energy poverty one year in advance and identify minimal-effort interventions. With this, our work shifts the focus from reactive mitigation to proactive prevention.

The primary contribution of this work is a framework that not only identifies households at risk of future energy poverty but also provides actionable, human-interpretable pathways to avoid it. Our model achieved a ROC AUC of 70.01%, and can therefore distinguish between future energy-poor and non-energy-poor households using only historical data. The model's sensitivity (73.25%) is significant from a policy perspective as it focuses on accurately identifying households at risk, which is essential for successful early intervention. This provides the Contribution 1 of this work.

Our counterfactual analysis highlights the central role of household income. In all the prototype profiles, a modest rise in household income was always the most effective way to change a household's predicted condition from being energy-poor to not being energy-poor. This was true for both single-action and multi-feature counterfactuals. The required income increases were often less than 5%, suggesting that for households on the borderline, relatively small financial improvements can make a significant difference.

Our results also highlight the long-term nature of energy poverty. The use of data from the preceding seven years confirms that energy poverty is not merely a consequence of immediate circumstances but is related to a household's long-term socioeconomic trajectory. This challenges the conventional focus on contemporaneous factors and suggests that policies must consider the cumulative impact of past conditions. Interestingly, our analysis revealed that interventions could be effective at various time points. For example, both recent income gains (in years $T = 6$ and $T = 7$) and earlier improvements (in years $T = 1$ and $T = 2$) were shown to reduce future risk. This implies that both sustained financial health and timely support can be critical for prevention. We highlight here that the contemporaneous relation between income and energy poverty has been highlighted in previous work [33, 34]; however, our results reveal the effect of lagged income and future energy poverty, indicating that a household's long-term financial trajectory, not just its present condition, is an important determinant of its future risk.

Other indicators of energy poverty include the unemployment rate, the rate of part-time employment, and the cost of energy. Concerning energy prices, their connection to energy poverty aligns with recent research [35].

In addition to the counterfactual analysis, we examined whether small year-to-year income increases are associated with improvements in MEPI values in the observed data. The results indicated that households experiencing a modest income increase of less than 5% exhibit a lower average MEPI (0.051) compared to households without such an increase (0.062). Regression results in Table 2 of Section B confirm that these differences are statistically significant. A small income increase is associated with an average reduction in MEPI of 0.011 ($p < 0.001$), even after controlling for year fixed effects and clustering standard errors at the household level. These findings support the plausibility of our counterfactual scenarios, showing that the small changes in income identified through the framework align with realistic transitions found in the HILDA dataset.

All things considered, the findings of this study offer concrete guidance for policymakers and social support organizations. The predictive model can serve as a tool to identify households that are likely to face energy poverty. This allows for early, targeted support to prevent the problem before it fully develops. The counterfactual analysis complements the model's predictions by generating interpretable scenarios for intervention. It also provides evidence-based strategies with minimal interventions to prevent energy poverty.

For policymakers, this framework supports the design of more efficient and cost-effective programs. Instead of implementing broad, untargeted subsidies, governments could deploy precise interventions, such as small, targeted income supplements, financial counseling, or employment support for households flagged by the model. The scenarios presented in Table 1 illustrate how such support could be tailored to offset specific external shocks, like rising energy prices or a drop in labor force participation. This counterfactual analysis constitutes the Contribution 2 of our work.

For non-governmental and community organizations, our findings reinforce the importance of programs aimed at improving income stability. The evidence that even small income increases can be highly effective provides a strong rationale for prioritizing such initiatives.

However, we acknowledge several limitations in this explanatory study. First, the predictive power of our model (ROC AUC of 70.01%) is good but not perfect. Misclassifications are inevitable, and it is important to consider their real-world consequences. A false positive (wrongly flagging a household as at-risk) may lead to inefficient allocation of resources, while a false negative (failing to identify a household that will become energy-poor) means a missed opportunity for prevention. Second, our counterfactual analysis is based on the patterns learned by the model, but these should not be interpreted as definitive causal links. For instance, while an income increase is associated with avoiding energy poverty, our study does not model the underlying cause of that increase (e.g., a new job, a promotion, or a government benefit), which could have its own complex effects. The prototype approach does not distinguish cases where profiles are close in feature space but differ in which features contribute to that similarity. Another limitation is the uniform treatment of households, which overlooks heterogeneity in responses to energy poverty predictors. In other words, this approach implicitly assumes that all households above the threshold experience energy poverty in a similar way. However, in practice, households may face very different forms of deprivation. Factors such as income, age, education, and personal traits likely influence how individuals experience and respond to energy challenges [49]. Finally, our findings are based on data from the HILDA survey. The specific drivers of energy poverty and the effectiveness of certain interventions may differ in other countries with different climates, energy markets, economic conditions, and social safety nets. Therefore, the generalizability of our results to other contexts should be approached with caution.

Acknowledgments

This work was partially supported by FCT-Fundação para a Ciência e a Tecnologia, Portugal through project PRODY: PTDC/CCI-COM/4464/2020; by NOVA LINCS ref. UIDB/04516/2020 (https://doi.org/10.54499/UIDB/04516/2020) and ref. UIDP/04516/2020 (https://doi.org/10.54499/UIDP/04516/2020) with the financial support of FCT.IP; by LARSyS ref. LA/P/0083/2020 (https://doi.org/10.54499/LA/P/0083/2020) and ref. UIDP/50009/2020 (https://doi.org/10.54499/UIDP/50009/2020) the financial support of FCT.IP; and with the financial support provided by the 2022 R&D&I National Projects and 2021 Strategic Projects Oriented to the Ecological and Digital Transition by the Spanish Ministry of Sciences and Innovation (Refs: PID2022-143254OB-I00 and TED2021-132824B-I00).

Diogo Nuno Freitas was supported by the Portuguese Foundation for Science and Technology with the grant number 2021.07966.BD (https://doi.org/10.54499/2021.07966.BD).

We thank Fábio Mendonça (https://orcid.org/0000-0002-5107-3248) for reviewing the introduction and methodology sections of this work and providing valuable feedback.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-4o and Grammarly in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

[1] International Energy Agency, World Energy Outlook 2024, 2024. URL: https://www.iea.org/reports/world-energy-outlook-2024, (Accessed on 10 November 2024).

[2] M. S. Shabbir, C. W. Cheong, M. Jaradat, R. Lile, G. Niedbała, et al., Energy deprivation to financial prosperity: Unveiling multidimensional energy poverty’s influence, Energy Strategy Reviews 54 (2024) 101473. doi:10.1016/j.esr.2024.101473.

[3] H. T. Nguyen-Phung, H. Le, Elderly well-being amidst energy poverty: Exploring the health, social, and economic impacts in Vietnam, Energy Research & Social Science 118 (2024) 103762. doi:10.1016/j.erss.2024.103762.

[4] T. M. Pondie, F. D. Engwali, B. E. O. Nkoa, E. N. Domguia, Energy poverty and respiratory health in Sub-Saharan Africa: Effects and transmission channels, Energy 297 (2024) 131158. doi:10.1016/j.energy.2024.131158.

[5] S. Nawaz, Energy poverty, climate shocks, and health deprivations, Energy Economics 100 (2021) 105338. doi:10.1016/j.eneco.2021.105338.

[6] K. B. Dear, A. J. McMichael, The health impacts of cold homes and fuel poverty, Bmj 342 (2011) d2807. doi:10.1136/bmj.d2807.

[7] V. Ballesteros-Arjona, L. Oliveras, J. Bolívar Muñoz, A. Olry de Labry Lima, J. Carrere, et al., What are the effects of energy poverty and interventions to ameliorate it on people’s health and well-being? A scoping review with an equity lens, Energy Research & Social Science 87 (2022) 102456. doi:10.1016/j.erss.2021.102456.

[8] R. Bentley, L. Daniel, Y. Li, E. Baker, A. Li, The effect of energy poverty on mental health, cardiovascular disease and respiratory health: a longitudinal analysis, Lancet Reg Health West Pac 35 (2023) 100734. doi:10.1016/j.lanwpc.2023.100734.

[9] J. Sokołowski, J. Frankowski, P. Lewandowski, Energy poverty, housing conditions, and self-assessed health: evidence from Poland, Housing Studies 39 (2024) 2325–2354. doi:10.1080/02673037.2023.2176831.

[10] G. Mohan, Young, poor, and sick: The public health threat of energy poverty for children in ireland, Energy Research & Social Science 71 (2021) 101822. doi:10.1016/j.erss.2020.101822.

[11] S. Shapira, N. Teschner, No heat, no eat:(dis) entangling insecurities and their implications for health and well-being, Social Science & Medicine 336 (2023) 116252. doi:10.1016/j.socscimed.2023.116252.

[12] S. Pye, A. Dobbins, Energy poverty and vulnerable consumers in the energy sector across the EU: Analysis of policies and measures, Technical Report, INSIGHT_E Consortium, 2015.

[13] J. M. Fry, L. Farrell, J. B. Temple, Energy poverty and retirement income sources in Australia, Energy Economics 106 (2022) 105793. doi:10.1016/j.eneco.2021.105793.

[14] A. Awan, F. Bilgili, D. B. Rahut, Energy poverty trends and determinants in Pakistan: Empirical evidence from eight waves of HIES 1998–2019, Renewable and Sustainable Energy Reviews 158 (2022) 112157. doi:10.1016/j.rser.2022.112157.

[15] I. Koomson, S. Awaworyi Churchill, Employment precarity and energy poverty in post-apartheid South Africa: Exploring the racial and ethnic dimensions, Energy Economics 110 (2022) 106026. doi:10.1016/j.eneco.2022.106026.

[16] P. Nussbaumer, M. Bazilian, V. Modi, Measuring energy poverty: Focusing on what matters, Renewable and Sustainable Energy Reviews 16 (2012) 231–243. doi:10.1016/j.rser.2011.07.150.

[17] S. A. Sy, L. Mokaddem, Energy poverty in developing countries: A review of the concept and its measurements, Energy Research & Social Science 89 (2022) 102562. doi:10.1016/j.erss.2022.102562.

[18] R. Banerjee, V. Mishra, A. A. Maruta, Energy poverty, health and education outcomes: Evidence from the developing world, Energy Economics 101 (2021) 105447. doi:10.1016/j.eneco.2021.105447.

[19] Y. Bao, T. Liao, Multidimensional poverty and growth: Evidence from India 1998–2021, Economic Modelling 130 (2024) 106586. doi:10.1016/j.econmod.2023.106586.

[20] C. Boţa-Avram, S. A. Apostu, R. Ivan, M. V. Achim, Exploring the impact of macro-determinant factors on energy resource depletion: Evidence from a worldwide cross-country panel data analysis, Energy Economics 130 (2024) 107341. doi:10.1016/j.eneco.2024.107341.

[21] M. Igawa, S. Managi, Energy poverty and income inequality: An economic analysis of 37 countries, Applied Energy 306, Part B (2022) 118076. doi:10.1016/j.apenergy.2021.118076.

[22] E. Kocak, E. E. Ulug, B. Oralhan, The impact of electricity from renewable and non-renewable sources on energy poverty and greenhouse gas emissions (GHGs): Empirical evidence and policy implications, Energy 272 (2023) 127125. doi:10.1016/j.energy.2023.127125.

[23] Y. Kolati, C. Raghutla, Do renewable energy sources and energy infrastructure contribute to mitigating energy poverty? Exploring uncharted dynamics, Energy Strategy Reviews 58 (2025) 101687. doi:10.1016/j.esr.2025.101687.

[24] L. Karpinska, S. Śmiech, Invisible energy poverty? Analysing housing costs in Central and Eastern Europe, Energy Research & Social Science 70 (2020) 101670. doi:10.1016/j.erss.2020.101670.

[25] B. Manasi, J. P. Mukhopadhyay, Definition, measurement and determinants of energy poverty: Empirical evidence from Indian households, Energy for Sustainable Development 79 (2024) 101383. doi:10.1016/j.esd.2024.101383.

[26] Z. Zhang, H. Shu, H. Yi, X. Wang, Household multidimensional energy poverty and its impacts on physical and mental health, Energy Policy 156 (2021) 112381. doi:10.1016/j.enpol.2021.112381.

[27] Q. Zhang, S. Appau, P. L. Kodom, Energy poverty, children's wellbeing and the mediating role of academic performance: Evidence from China, Energy Economics 97 (2021) 105206. doi:10.1016/j.eneco.2021.105206.

[28] A. O. Crentsil, D. Asuman, A. P. Fenny, Assessing the determinants and drivers of multidimensional energy poverty in Ghana, Energy Policy 133 (2019) 110884. doi:10.1016/j.enpol.2019.110884.

[29] K. Abbas, S. Li, D. Xu, K. Baz, A. Rakhmetova, Do socioeconomic factors determine household multidimensional energy poverty? Empirical evidence from South Asia, Energy Policy 146 (2020)

111754. doi:10.1016/j.enpol.2020.111754.

[30] M. Umair, A. Aizhan, V. Teymurova, L. Chang, et al., Does the disparity between rural and urban incomes affect rural energy poverty?, Energy Strategy Reviews 56 (2024) 101584. doi:10.1016/j.esr.2024.101584.

[31] K. Prakash, S. Awaworyi Churchill, R. Smyth, Are you puffing your children's future away? Energy poverty and childhood exposure to passive smoking, Economic Modelling 114 (2022) 105937. doi:10.1016/j.econmod.2022.105937.

[32] S. Hosan, K. K. Sen, M. M. Rahman, S. C. Karmaker, A. J. Chapman, et al., Evaluating the mediating role of energy subsidies on social well-being and energy poverty alleviation in Bangladesh, Energy Research & Social Science 100 (2023) 103088. doi:10.1016/j.erss.2023.103088.

[33] F. Dalla Longa, B. Sweerts, B. van der Zwaan, Exploring the complex origins of energy poverty in The Netherlands with machine learning, Energy Policy 156 (2021) 112373. doi:10.1016/j.enpol.2021.112373.

[34] W. van Hove, F. Dalla Longa, B. van der Zwaan, Identifying predictors for energy poverty in Europe using machine learning, Energy and Buildings 264 (2022) 112064. doi:10.1016/j.enbuild.2022.112064.

[35] C. Spandagos, M. A. T. Reaños, M. Á. Lynch, Energy poverty prediction and effective targeting for just transitions with machine learning, Energy Economics 128 (2023) 107131. doi:10.1016/j.eneco.2023.107131.

[36] S. Gawusu, S. A. Jamatutu, A. Ahmed, Predictive modeling of energy poverty with machine learning ensembles: Strategic insights from socioeconomic determinants for effective policy implementation, International Journal of Energy Research 2024 (2024) 9411326. doi:10.1155/2024/9411326.

[37] S. Dhongde, P. K. Pattanaik, Y. Xu, Well-being, deprivation, and the Great Recession in the U.S.: A study in a multidimensional framework, Review of Income and Wealth 65 (2019) S281–S306. doi:10.1111/roiw.12411.

[38] K. Decancq, M. A. Lugo, Weights in multidimensional indices of wellbeing: An overview, Econometric Reviews 32 (2013) 7–34. doi:10.1080/07474938.2012.690641.

[39] H. Phoumin, F. Kimura, Cambodia's energy poverty and its effects on social wellbeing: Empirical evidence and policy implications, Energy Policy 132 (2019) 283–289. doi:10.1016/j.enpol.2019.05.032.

[40] J. H. Sullivan, M. Warkentin, L. Wallace, So many ways for assessing outliers: What really works and does it matter?, Journal of Business Research 132 (2021) 530–543. doi:10.1016/j.jbusres.2021.03.066.

[41] M. Galar, A. Fernandez, E. Barrenechea, H. Bustince, F. Herrera, A review on ensembles for the class imbalance problem: Bagging-, boosting-, and hybrid-based approaches, IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews) 42 (2012) 463–484. doi:10.1109/TSMCC.2011.2161285.

[42] K. Murphy, Machine Learning: A Probabilistic Perspective, Adaptive Computation and Machine Learning series, MIT Press, Cambridge, Massuchusetts, 2012.

[43] C. Seiffert, T. M. Khoshgoftaar, J. Van Hulse, A. Napolitano, RUSBoost: A hybrid approach to alleviating class imbalance, IEEE Transactions on Systems, Man, and Cybernetics, Part A (Systems and Humans) 40 (2010) 185–197. doi:10.1109/TSMCA.2009.2029559.

[44] X.-Y. Liu, J. Wu, Z.-H. Zhou, Exploratory undersampling for class-imbalance learning, IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) 39 (2009) 539–550. doi:10.1109/TSMCB.2008.2007853.

[45] S. Verma, V. Boonsanong, M. Hoang, K. Hines, J. Dickerson, et al., Counterfactual explanations and algorithmic recourses for machine learning: A review, ACM Computing Surveys 56 (2024) 312. doi:10.1145/3677119.

[46] R. K. Mothilal, A. Sharma, C. Tan, Explaining machine learning classifiers through diverse counterfactual explanations, in: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (ACM FAccT), Association for Computing Machinery, New York, NY, USA, 2020, pp. 607–617. doi:10.1145/3351095.3372850.

[47] I. Stepin, J. M. Alonso, A. Catala, M. Pereira-Fariña, A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence, IEEE Access 9 (2021) 11974–12001. doi:10.1109/ACCESS.2021.3051315.

[48] A. Dhurandhar, P.-Y. Chen, R. Luss, C.-C. Tu, P. Ting, et al., Explanations based on the missing: Towards contrastive explanations with pertinent negatives, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS), volume 31 of *Advances in Neural Information Processing Systems*, Curran Associates Inc., Montréal, Canada, 2018. doi:10.48550/arXiv.1802.07623.

[49] S. Cong, D. Nock, Y. L. Qiu, B. Xing, Unveiling hidden energy poverty using the energy equity gap, Nature communications 13 (2022) 2456. doi:10.1038/s41467-022-30146-5.

A. Model selection and grid search parameters

To identify the optimal hyperparameter configurations for predicting energy poverty, a grid search approach was implemented. This process systematically tested combinations of model parameters and evaluated their performance using cross-validation.

We used 5-fold cross-validation on the training dataset to ensure that the models were validated on various data splits and that the hyperparameters chosen were robust across different subsets of data. The best set of hyperparameters was then chosen based on the highest average ROC AUC score from the validation folds. The ROC AUC score measures the model's ability to discriminate between energy-poor and non-energy-poor households.

Additionally, the grid search utilized the one-vs-rest framework, which creates a binary classifier for each class.

All models were optimized and trained using Python, and the scikit-learn library for model implementation and evaluation. The model description and the parameters used in the grid search optimization are detailed below.

Balanced bagging

A balanced bagging classifier [41] is an ensemble technique that combines the predictions of multiple base models, e.g., decision trees, in order to improve the robustness and accuracy of the outcomes. This method specifically addresses class imbalance by ensuring that each decision tree in the ensemble is trained on a balanced subset of the dataset. These subsets are created by resampling the original training data, wherein each subset contains a representative distribution of both minority (energy-poor) and majority (not energy-poor) classes.

The parameter grid included the following parameters:

- Number of estimators: 10, 50, 100
- Maximum samples (proportion): 0.5, 1.0
- Maximum features (proportion): 0.5, 1.0
- Bootstrap sampling: True, False
- Bootstrap feature selection: True, False
- Sampling strategy (proportion): Auto, 0.5, 1.0
- Replacement: True, False

B. Robustness Check: Income Increases and MEPI

Table 2
OLS regression of MEPI on small income increase (0–5%), with year fixed effects

Variable	Coefficient	Std. Error	z-stat	p-value
Intercept	0.0550	0.003	20.034	$<$0.001
Small income increase (0–5%)	-0.0112	0.002	-6.282	$<$0.001
Year fixed effects	Yes			
Observations		64,927		

Results are based on OLS regression with robust standard errors clustered by household. Households experiencing a modest year-to-year income increase of less than 5% exhibit a lower average MEPI (0.051) compared to households without such an increase (0.062). Regression results confirm that these differences are statistically significant.

A small income increase is associated with an average reduction in MEPI of 0.011 ($p < 0.001$), even after controlling for year fixed effects and clustering standard errors at the household level.

These findings support the plausibility of our counterfactual scenarios by showing that the small income changes highlighted by the framework correspond to realistic and observed transitions in the HILDA dataset.

C. Variables of the HILDA survey

Table 3: Summary of variables used in the study.

Variable	Description	Type	Observations
Participant ID	Identification of the participant.	––	This variable was not included in the prediction model
Year	Year the participant data was collected.	Discrete	This variable was not included in the prediction model
Part-time employment rate	Percentage of the workforce employed part-time.	Continuous	
Unemployment rate	Total unemployment rate.	Continuous	
Unemployment rate change	Year-over-year change in the total unemployment rate.	Continuous	Reflects shifts in labor market conditions and employment opportunities.
Total labor force participation rate	Labor force participation rate.	Continuous	
Total labor force participation rate change	Year-over-year change in the total labor force participation rate.	Continuous	Indicates changes in the proportion of individuals actively participating in the labor market.
GSP per capita	Gross state product per capita.	Continuous	
GSP per capita growth	Growth rate of gross state product per capita.	Continuous	
GSP per capita growth change	Year-over-year change in the growth rate of gross state product per capita.	Continuous	Captures the economic growth fluctuations at the state level.
Energy price	Energy price.	Continuous	Adjusted for inflation.
Energy price change	Year-over-year change in the energy price.	Continuous	Measures fluctuations in energy costs that may impact household energy affordability.
Years of education	Logarithm of the number of years the participant has been educated.	Continuous	
Age	Age	Discrete	
Married	If married (=1) or not (=0)	Binary	
Divorced	If divorced (=1) or not (=0)	Binary	
Widowed	If widowed (=1) or not (=0)	Binary	
Children presence at household	Number of minors at home.	Discrete	
Unemployment status	If unemployed (=1) or not (=0)	Binary	
Poor health	If the individual perceives their health status as bad or very bad (=1) or not (=0).	Binary	
Employment status	If employed (=1) or not (=0).	Binary	
Household size	Household size.	Discrete	
Household income	Logarithm of the household income.	Continuous	
Household income change	Year-over-year change in the household income.	Continuous	Reflects income volatility and its potential effect on household vulnerability to energy poverty.
Household region 2 to 8	Various one-hot encoded regional and household state indicators.	Binary (one-hot encoded)	
MEPI	Poverty indicator.	Continuous	Includes the 2M, TPR, LIHC indicators, and measures of inability to pay for adequate heating (Heat) or utility bills on time (Arrears).