A Methodology to Address Perceived Hazards from the Twin Transition Among Disadvantaged Groups: A Nine-Step Ordinal-Modelling Approach with Survey Evidence from Six European Countries

Special Session S69 — Embedding Eco-Social Justice in Europe's Twin Transition: Multi-Method Approaches for Policy Innovation

Mehdi Zarehparast Malekzadeh[1][0009-0009-7652-1439], Aditya Kapoor[1][0009-0007-6878-0313], Gemma Dolores Molero[1][0000-0001-7296-2282], Bojana Bjegojevic[2][0000-0002-8421-5572], Ilgiz Nurlanov[3][0009-0000-1509-7333], Oriana Salomón[4][0009-0005-2580-1626], Maria Sangiuliano[4][0000-0003-0468-107X] and Francisco Enrique Santarremigia[1*][0000-0001-5466-9269]

[1]AITEC Asesores Internacionales, SRL (AITEC), Parque Tecnológico. C/Charles Robert Darwin, 20, 46980 Paterna-Valencia, Spain. https://www.aitec-intl.com/

[2] Human Factors in Safety and Sustainability Research Centre, Technological University Dublin. https://www.tudublin.ie/explore/faculties-and-schools/sciences-health/food-science--environmental-health-grangegorman/

[3] SFC Nonprofit Ltd. 1191 Budapest, Fő u. 11. 7/20., Hungary, https://szekely.family/

[4] Research and European Projects Department, Smart Venice, Fondamenta di Cannaregio 969, Venice, Italy, https://www.smartvenice.org/

Abstract. This paper presents and empirically applies a methodology for predicting and prioritising the hazards that European citizens perceive in the green and digital ("twin") transition, with explicit attention to disadvantaged groups. Developed within the Horizon Europe project FITTER-EU, the methodology is a nine-step, fully reproducible analytical pipeline that converts harmonised, EU-SILC-aligned perception surveys into interpretable, profile-specific risk intelligence. After data cleaning and encoding, each hazard is modelled independently: candidate predictors are shortlisted with LASSO, an ordinal proportional-odds model is fitted, and predictors are confirmed through likelihood-ratio testing with Benjamini–Hochberg false-discovery-rate correction, an events-per-predictor cap, and four resampling robustness tests. Confirmed effects are quantified as odds ratios and marginal effects, the most affected groups are identified by ANOVA, and results are compared across countries and hazards. The pipeline is applied identically to all three FITTER-EU sectors—Energy (n=407), Housing (n=263) and Transport (n=358), 1,028 quality-controlled respondents across six countries. Three findings hold across sectors: financial vulnerability is the strongest and most consistent predictor of high concern—in Energy, respondents with repeated utility-bill arrears have nearly six times the odds of high concern (OR 5.76); the most widely shared hazards (heating and cooling costs, electricity bills, fuel and repair costs) are the hardest to attribute to any group; and the country a person lives in explains only a small share (around 1–4%) of how concerned they are. The framework closes the loop with a participatory phase in which mitigation measures are co-created with stakeholders—26 national-lab meetings, 237 participants, 82 measures—and qualitatively assessed through an LLM-based agent ("Dr. Transition"), advancing anticipatory, evidence-based and participatory governance for an inclusive twin transition.

Keywords: Twin transition, Hazard perception, Disadvantaged groups, EU-SILC, LASSO, Proportional-odds regression, Odds ratios, False-discovery-rate control, Co-creation, Large Language Models, Eco-social justice.

1. Introduction

Consider a household in a mid-century apartment block being told that, within a few years, its gas boiler must be replaced,

[1*]Corresponding author. Mobile phone: +34 647 97 81 12.
E-mail address: fsantarremigia@aitec-intl.com

its building retrofitted, and its petrol car retired in favour of an electric one. For a comfortable, well-informed owner-occupier, these are manageable upgrades. For a tenant on a tight budget, with arrears already accumulating and limited say over the building they live in, the same policies can read as a cascade of unaffordable obligations arriving faster than the means to meet them. The twin transition is experienced very differently depending on where a person stands—and it is precisely this difference that decides whether the transition is felt as an opportunity or as a threat.

The European Union has committed to a double structural transformation: a green transition towards climate neutrality and a digital transition towards technological leadership, jointly framed as the "twin transition". While the aggregate welfare case is strong, its distributional consequences are not neutral. Fast-moving interventions—the roll-out of photovoltaic (PV) systems, the deep retrofitting of the housing stock, and the electrification of road transport—reallocate costs, infrastructure access and regulatory exposure in ways that can disadvantage specific population groups. Low-income households, the elderly, people with disabilities, tenants, migrants and rural populations may face higher up-front and recurring costs, thinner infrastructure, split-incentive problems and unequal access to the benefits of decarbonisation. If these effects are not anticipated and mitigated, the twin transition risks reproducing—or even widening—existing inequalities, eroding the social legitimacy on which the transition itself depends.

Anticipating these effects requires knowing two things that conventional statistics do not readily supply: not just who is objectively exposed, but who feels at risk, and why. Perceived risk matters in its own right—it shapes whether people support, resist, or disengage from transition policies—and it is unevenly distributed in ways that official indicators, built from registers and averages, tend to miss. The challenge this paper addresses is therefore both measurement and method: how to capture citizens' own perceptions of transition hazards, and how to turn those perceptions into reliable, profile-specific evidence that policymakers can act on without over-claiming.

The Horizon Europe project FITTER-EU ("Fair and Inclusive Twin Transitions for a stronger social Europe", Grant Agreement No. 101132546) responds to this challenge by building an anticipatory-governance ecosystem. At its centre is an interactive Digital Platform powered by a Predictive Decision Support System that enables policymakers to anticipate which social groups perceive themselves at risk under different twin-transition policy scenarios, to simulate the implementation of those policies, and to retrieve proposed mitigation measures. The project couples quantitative risk assessment with participatory policy design, so that the perspective of disadvantaged groups is embedded in the evidence base that informs policy.

This paper operationalises the predictive and participatory components of that ecosystem. Its contribution is threefold. First, it presents a rigorous and transferable nine-step analytical pipeline that converts perceptual hazard data into profile-specific risk intelligence, combining modern variable selection (LASSO), an ordinal model appropriate to Likert-type concern data (proportional-odds regression), and a multi-layered confirmation regime (likelihood-ratio testing, Benjamini–Hochberg false-discovery-rate control, an events-per-predictor cap, and four resampling robustness tests). Second, it reports the first full empirical application of this pipeline across all three FITTER-EU sectors, drawing on 1,028 quality-controlled survey responses collected across six European countries. Third, it closes the analytical loop with a participatory phase in which the quantitative evidence feeds the co-creation of mitigation measures with stakeholders, qualitatively assessed by a Large Language Model (LLM)-based agent.

The distinctive feature of the approach is that it is explainable end-to-end and disciplined against over-claiming: every confirmed predictor must survive selection, estimation, multiple-testing correction and resampling before it is reported, and findings are separated into those defensible at the individual-predictor level and those defensible only at the construct level. The methodology builds on the FITTER-EU analytical framework first introduced for the transport sector by Molero et al. [1] and the evolutionary-tree framing of Kapoor et al. [2] and complements the AI-guided policy-making process described by Kapoor et al. [3].

The remainder of the paper is organised as follows. Section 2 describes the methodology—survey design, sampling, and the nine analytical steps with their formulas. Section 3 reports the empirical results across the three sectors and the cross-sector conclusions. Section 4 presents the participatory co-creation phase and the operational platform. Section 5 discusses the findings against the just-transition literature, and Section 6 concludes.

1.1 Related work and contribution

The paper sits at the intersection of three literatures. The first is the measurement of transition-related vulnerability and energy/transport poverty, where composite indicators have advanced the mapping of who is exposed but typically rely on macro- or register-level proxies rather than on the perceptions of citizens themselves [13]. The second is the application of disciplined statistical-learning methods to survey data: LASSO for high-dimensional variable selection [6], ordinal proportional-odds regression for Likert-type outcomes [7], and false-discovery-rate control for honest multiple testing [8]. The third is participatory and co-creative policy design for just transitions, in which the perspective of disadvantaged groups is built into the policy response rather than inferred from above.

Against this backdrop, the contribution of the present work is to integrate all three—perception measurement, explainable and conservative statistical inference, and participatory co-creation—within a single, end-to-end and reproducible methodology, and to add a conversational-AI layer that makes the resulting evidence and measures navigable for non-technical policy users. Two design choices distinguish the pipeline from simpler approaches. First, it never reports a predictor on the strength of a single test: every confirmed predictor must survive LASSO selection, ordinal estimation, FDR-corrected likelihood-ratio testing, an events-per-predictor cap, and four resampling tests, and findings are explicitly partitioned into predictor-level claims (for doubly-robust predictors) and construct-level claims (for the rest). Second, country is represented through national macro indicators rather than country dummies, so that concern is attributed to transferable circumstances (utility-arrears rates, electricity-price bands, degree-days) rather than to nationality, making the findings portable to other countries with similar conditions.

2. Methodology

The methodology comprises a data-collection design, a sampling strategy and a nine-step analytical pipeline. It follows the FITTER-EU consortium protocol and is reported in full in project Deliverables D4.1 (survey design and data collection) and D4.4 (computational analysis). The pipeline is sector-agnostic: the identical sequence of steps and the same statistical formalism are applied to Energy, Housing and Transport. This section summarises the protocol and the empirical instrument.

2.1 Survey design and data collection

Three sector-specific questionnaires were developed under a harmonised protocol informed by WP2 case studies and aligned with the European Union Statistics on Income and Living Conditions (EU-SILC) variable framework. Each questionnaire targets one prioritised policy intervention and a sector-specific catalogue of hazards: Energy assesses PV-installation hazards (7 hazards); Housing assesses retrofitting hazards (13 hazards); and Transport assesses electric-vehicle (EV) acquisition hazards (6 hazards). For every hazard, respondents reported, on two 5-point Likert scales, the perceived likelihood of the hazard and the perceived severity (impact) of its consequences. Part B of each questionnaire recorded socio-demographic and EU-SILC-aligned variables (gender, age, income, tenancy, disability, urbanity, migrant background, education, employment, ethnicity and LGBTIQA+ status), and country-level macro indicators (e.g. utility-arrears rate, cooling/heating degree-days, household electricity-price band) were appended to each respondent.

2.2 Sampling strategy

The target population is all residents aged 18 and over in the FITTER use-case countries potentially affected by the policy interventions. Participants were recruited through the academic panel platform Prolific, integrated with LimeSurvey for real-time quota management and quality control [4]. Recruitment used non-probability quota sampling with EU-SILC-aligned filters (country, income band, sex, and a sector-specific variable—tenancy for Energy/Housing or disability for Transport) so that the achieved sample approximates Eurostat/EU-SILC population distributions. The consortium adopted a design target of 456 respondents per sector, stratified by EU-SILC proportions; this value satisfies the rule of thumb of 15–20 observations per predictor and general survey-research thresholds [5]. The limitations of non-probability sampling—self-selection and unknown inclusion probabilities—are acknowledged; quota controls and EU-SILC stratification keep overall bias comparatively low. After fieldwork, a quality-control protocol removed respondents who completed the survey in under four minutes, fixed data-entry errors, merged categories with fewer than five respondents, and deduplicated to one row per participant. The achieved analytical samples are 407 (Energy), 263 (Housing) and 358 (Transport), totalling 1,028 verified respondents across Germany, Italy, Spain, Portugal, Hungary, and Ireland.

2.3 The nine analytical steps

At its heart, the pipeline answers three plain questions, in order. First, for each transition hazard, who is most worried about it, and can we say so reliably? Second, how much does each characteristic actually matter—by how much does being, say, in arrears on utility bills raise the odds of being highly concerned? Third, do these patterns hold across countries and across different hazards, or are they one-offs? The nine steps are simply the disciplined route to trustworthy answers: they take raw survey responses, screen out the noise, fit a model suited to the data, and—crucially—refuse to report any finding that cannot survive repeated, stringent testing.

The pipeline groups these nine steps into three phases (Figure 1). Phase 1 (Data Preparation, Steps 1–2) cleans and encodes the data once for the full dataset. Phase 2 (Modelling and Confirmation, Steps 3–5c) builds a per-hazard ordinal model, shortlists candidate predictors with LASSO, and confirms which predictors are real through likelihood-ratio testing, false-discovery-rate correction, an events-per-predictor cap, and resampling. Phase 3 (Interpretation, Steps 6–9) measures effect sizes, identifies the most affected groups, profiles country-anchored macro predictors, and compares concern across countries and hazards. Two analyses run in parallel throughout: a per-hazard analysis (A, primary), in which every hazard is modelled independently with its own predictors; and a composite analysis (B, reference baseline), in which all hazards in a sector are averaged into one overall concern outcome. A guiding principle runs through every step: it is better to confirm fewer findings with confidence than to publish many that might not survive a second look.

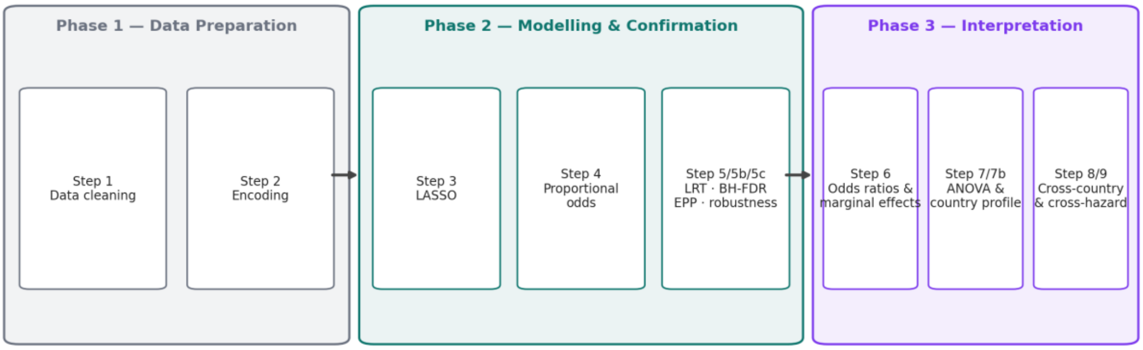

Fig. 1. The nine analytical steps of the methodology, grouped into three phases. Steps 1–2 run once on the full dataset; Steps 3–7b run independently for each hazard; Steps 8–9 operate at the cross-hazard level.

Step 1 — Getting the data ready. After removing system columns, free-text and routing-skipped items, and applying the four-minute completion filter and data-entry checks, each hazard is summarised by a single concern score that combines the two Likert ratings (Formula 1). Multiplying likelihood by impact ensures that a hazard seen as both very likely and very damaging scores far higher than one that is only one or the other.

$$\mathrm{HazardScore}(h) = \mathrm{Likelihood}(h) \times \mathrm{Impact}(h), \quad \in [1, 25] \tag{1}$$

For the composite analysis, the per-respondent score is the average across the J hazards of the sector (Formula 2), and both the per-hazard and composite continuous scores are converted into five ordered concern levels by quintile binning (Formula 3), giving approximately equal-sized groups that the ordinal model can estimate reliably.

$$\mathrm{Composite}(r) = \frac{1}{J}\sum_{h=1}^{J} \mathrm{HazardScore}(r, h) \tag{2}$$

$$\mathrm{ConcernLevel}(r) = Q_5(\mathrm{score}(r)) \in \{1, 2, 3, 4, 5\} \tag{3}$$

Before averaging hazards into a composite, internal consistency is verified with Cronbach's α (Formula 4); all three sectors exceed the conventional thresholds (Energy α=0.82, Housing α=0.84, Transport α=0.79), confirming that the hazards within a sector measure a coherent underlying concept.

$$\alpha = \frac{J}{J-1}\left(1 - \frac{\Sigma_j \mathrm{Var}(\mathrm{item}_j)}{\mathrm{Var}(\mathrm{total})}\right) \tag{4}$$

Step 2 — Preparing the data for the model. Categorical answers are one-hot-encoded, age is binned, and country labels are deliberately excluded from the predictor set—replaced by national macro variables—so that concern is attributed to specific national circumstances (e.g. utility-arrears rate, electricity-price band) rather than to country dummies. This makes findings transferable across countries with similar conditions. Remaining missing cells are filled with iterative k-nearest-neighbour imputation.

Step 3 — Choosing which variables to test (LASSO). For each hazard, LASSO logistic regression with 5-fold cross-validation shortlists the candidate predictors most worth testing, pushing uninformative coefficients to exactly zero (Formula 6) [6]. LASSO is a practical filter—surviving it means a predictor is worth testing, not that it is confirmed.

$$\hat{\beta}^{\mathrm{LASSO}} = \mathrm{argmin}_{\beta}\left[\mathrm{Loss}(\beta) + \lambda\sum_{p}|\beta_p|\right] \tag{6}$$

In plain terms, each survey carries roughly a hundred candidate characteristics per respondent—far too many to test responsibly against a few hundred answers. LASSO acts as a sieve: it keeps the handful of characteristics that carry a genuine signal and discards the rest, so that the formal testing in later steps is not overwhelmed by dozens of spurious comparisons. It casts a deliberately wide net at this stage; the tightening happens next.

Step 4 — The main statistical model. Each hazard's five-level concern outcome is modelled with an ordinal proportional-odds regression (Formula 7) [7]. The proportional-odds assumption is checked with the Brant test (0 violations across all three sectors), and McFadden's R^2 describes how much variation each model explains. These are model-quality gates; they do not confirm individual predictors.

$$\mathrm{logit}(P(Y \le k \mid X)) = \theta_k - \sum_{p}\beta_p X_p, \quad k = 1, \ldots, 4 \tag{7}$$

This model is chosen because concern is naturally ordered— "very high" is more than "high", which is more than “moderate”—but the gaps between levels are not numeric quantities to be added or averaged. A proportional-odds model respects that ordering it estimates, for each characteristic, how much it shifts a respondent up or down the concern ladder, without pretending the rungs are evenly spaced. The Brant test confirms the model is appropriate (no violations in any sector), and McFadden’s R^2 reports how much of the variation in concern the model captures—values of 0.10 or above being a good fit for survey data of this kind.

Step 5 — Confirming which predictors are real. Each LASSO-selected predictor is tested with a likelihood-ratio test (Formula 8), the p-values are corrected for multiple testing using the Benjamini–Hochberg false-discovery-rate procedure at $\alpha=0.05$ (Formula 9) [8], and an events-per-predictor (EPP) cap removes estimates based on subgroups too small for stable standard errors. A predictor is confirmed only when it passes all three filters.

$$\mathrm{LR} = -2(\ell_{\text{reduced}} - \ell_{\text{full}}) \;\sim\; \chi^2_{(df)} \tag{8}$$

$$p^{\text{adj}}_{(i)} = \min_{j\geq i}\, \min\Big(1,\; \frac{m}{j}\, p_{(j)}\Big) \tag{9}$$

Step 5b then validates the whole confirmed set against a null model with a global likelihood-ratio test, and Step 5c subjects the confirmed predictors to four resampling tests: selection stability under bootstrap resampling (Meinshausen–Bühlmann) and under the more conservative complementary-pairs subsampling (Shah–Samworth), a retention test (does FDR-significance survive subsample refitting), and a sign-consistency check. Crossing selection stability with retention places each confirmed predictor into one of four quadrants, of which “doubly robust” (reliably selected and reliably significant) is the strongest. This separation is what allows the paper to distinguish predictor-level claims from construct-level ones.

The reasoning behind this heavy filtering is worth making explicit, because it is what gives the findings their credibility. When dozens of characteristics are tested at once, some will appear “significant” by chance alone; the false-discovery-rate correction holds the expected share of such false alarms below 5%. The events-per-predictor cap then removes any effect resting on too few respondents to be stable. Finally, the resampling tests ask the hardest question of all—would this finding reappear if we had drawn a slightly different sample? A predictor that clears every one of these hurdles is reported as a firm, individual finding; one that is consistently selected but whose significance wobbles under resampling is reported more cautiously, as evidence about a broader construct (for example, “financial pressure”) rather than about one exact variable. The pipeline is, in short, designed to under-claim rather than over-claim.

Step 6 — Measuring the size of each effect. Confirmed coefficients are translated into odds ratios (Formula 10) and into marginal effects—the percentage-point change in the probability of being in the highest concern level when a predictor switches on (Formula 11). All effects are reported with 95% confidence intervals and an EPP-warning flag where the events-per-predictor ratio is low.

$$\mathrm{OR}_p = e^{\beta_p} \tag{10}$$

$$\mathrm{MEM}_p = P(Y=5 \mid X_p=1) - P(Y=5 \mid X_p=0) \tag{11}$$

These two numbers answer complementary questions and are worth reading together. An odds ratio of 4 means a characteristic quadruple the odds of being in a higher concern category—a relative statement. A marginal effect of, say, +25 percentage points means that switching the characteristic on raises the probability of being in the very-highest concern band by 25 points—an absolute statement that is often more intuitive for policy. A large odds ratio attached to a rare characteristic can still move few people, which is why both are reported, each with its 95% confidence interval, and why estimates resting on very small subgroups are flagged rather than highlighted.

Step 7 — Identifying the most affected groups. For each confirmed categorical predictor, a one-way ANOVA on the raw 1–25 concern score identifies which specific category carries the highest mean concern and how much it differs from the lowest, with post-hoc Tukey HSD or Games–Howell tests. Step 7b profiles the country-anchored macro predictors per country.

Steps 8–9 — Comparing across countries and hazards. Step 8 compares overall concern across countries with a composite-level ANOVA and Tukey HSD. Step 9 synthesises the per-hazard results: it ranks hazards by mean concern, maps which predictors are confirmed across multiple hazards (a predictor confirmed for many hazards is a robust cross-cutting finding), and checks A/B convergence between the per-hazard and composite analyses. Because the family-wise error rate across all hazards in a sector is not controlled, a predictor confirmed for a single hazard is read with caution, while predictors confirmed across multiple hazards carry substantially stronger evidential weight.

2.4 Two parallel analyses and the reproducibility of the pipeline

Throughout Steps 3–9, two analyses run side by side. The per-hazard analysis (A) is the primary track: each hazard is modelled completely independently, with its own LASSO shortlist, its own ordinal model, its own confirmed predictors, its own odds ratios and its own group profiles. This is what makes the cross-hazard synthesis of Step 9 possible and is the level at which the most policy-specific findings emerge—different hazards have genuinely different driver sets, which a single pooled model would obscure. The composite analysis (B) is a reference baseline: all hazards in a sector are averaged into one overall concern outcome (Formula 2) and modelled once, giving a sector-level summary against which the per-hazard findings are cross-checked. Reporting both, and flagging where they agree or diverge, is a built-in internal validation: a predictor confirmed in the composite and in several individual hazards is far stronger evidence than one appearing in a single model.

Two features make the pipeline conservative by construction. First, the events-per-predictor cap removes any effect estimated from a subgroup too small for a stable standard error, so the headline findings are never driven by a handful of respondents. Second, the four resampling tests of Step 5c separate two distinct senses of robustness—whether LASSO would re-select a predictor on a fresh sample (selection stability) and whether its significance survives subsample refitting (retention)—and only predictors strong on both are cited at the individual-predictor level. The entire pipeline is scripted and version-controlled, with fixed random seeds, so that every reported number can be reproduced from the cleaned dataset; the data-cleaning rules (four-minute completion floor, implausible-value checks, small-category merging) are applied identically across the three sectors.

3. Results

This section reports the application of the pipeline to the three sectors (Energy n=407, Housing n=263, Transport n=358). It is organised to move from the big picture to the detail and back again: Section 3.1 gives the cross-sector overview and the one pattern that unites all three sectors, Sections 3.2–3.4 take each sector in turn—always asking the same three questions (which hazards worry people most, who is most concerned, and how reliable is the evidence)—and Section 3.5 draws the conclusions that hold across sectors and matter most for policy. A short reading note before the numbers: throughout, a hazard "score" is the product of perceived likelihood and perceived impact on a 1–25 scale, with 12 the threshold for "high concern"; a "confirmed predictor" is a respondent characteristic that survived the full confirmation regime of Section 2; and a "doubly-robust" predictor is one that additionally reproduced under resampling and can therefore be cited individually rather than only as part of a broader construct.

3.1 Overview across the three sectors

Table 1 summarises the three sectors. All three exceed the Cronbach's α thresholds, the proportional-odds assumption holds everywhere (zero Brant violations), and the number of confirmed predictors and doubly robust predictors decreases with sample size and country homogeneity. Figure 2 shows the confirmed and doubly robust counts.

Table 1. Results overview across the three sectors.

Sector	Sample	Countries	Hazards	Cronbach α	Brant violations
Energy	407	DE, IT, ES, PT, IE	7	0.82 (good)	0 of 153
Housing	263	DE, PT, IE	13	0.84 (good)	0 of 186
Transport	358	IT, ES	6	0.79 (acceptable)	0 of 69

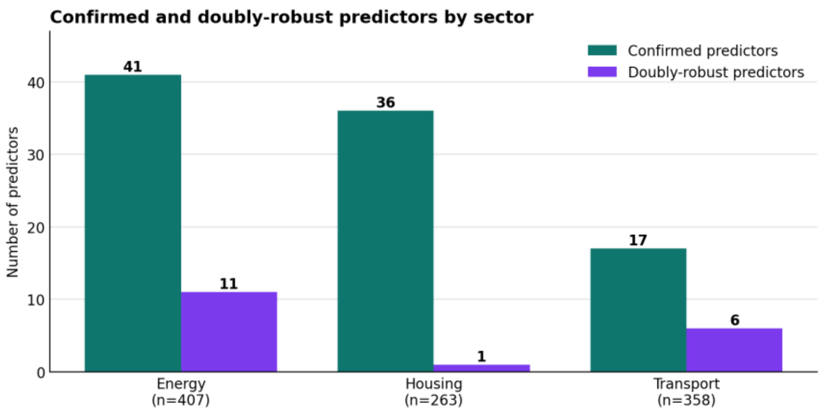

Fig. 2. Confirmed predictors and doubly robust predictors (reliably selected and reliably significant under resampling) by sector. The robustness gradient runs Transport > Energy > Housing, tracking sample size and country homogeneity.

One pattern runs through all three sectors, and it is initially counter-intuitive: the most universally shared concerns are the hardest to predict from individual characteristics. The hazards rated most concerning by the largest share of respondents—heating and cooling costs and electricity bills in Energy and Housing, fuel and repair costs in Transport—have the weakest predictive models, because when almost everyone reports the same concern level there is little variation for a model to explain. Hazards where concern is split more evenly produce the strongest models. An analogy helps: trying to predict who fears rising heating bills is like trying to predict who dislikes traffic jams—almost everyone does, so personal characteristics barely distinguish people. But trying to predict who worries about, say, missing out on solar savings does separate people meaningfully, because that concern depends on circumstances such as home ownership, dwelling age and prior intentions. The practical implication is that high-mean-concern hazards are universal concerns that policy must address broadly, whereas moderate-mean-concern hazards are where targeted policy can find traction. Figure 3 shows the highest-scoring hazards on the 1–25 scale.

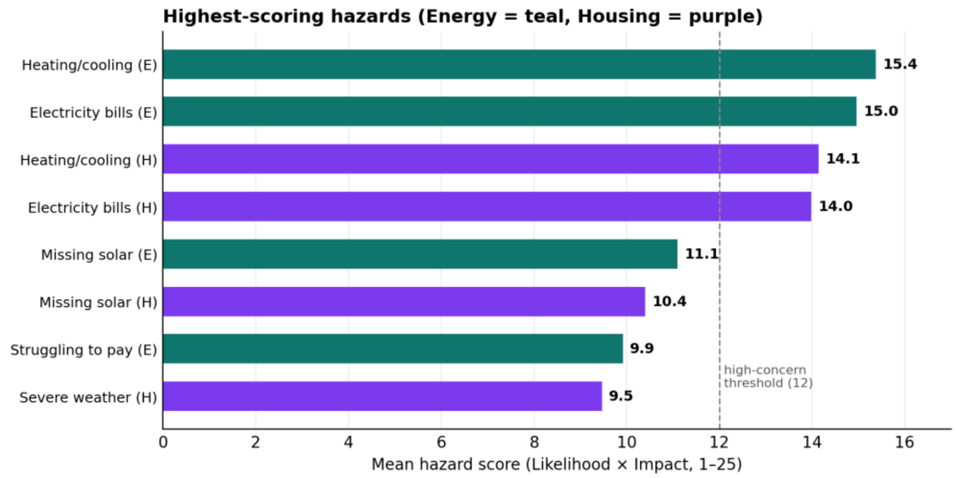

Fig. 3. Highest-scoring hazards by mean concern (Likelihood × Impact, 1–25). Cost-of-living hazards dominate in both Energy and Housing; the dashed line marks the high-concern threshold of 12.

3.2 Energy sector (n = 407)

The Energy sample spans six countries (Germany, Italy, Spain, Portugal, Ireland, Hungary) with α=0.82. The seven Energy hazards are dominated by cost-of-living concerns: heating and cooling costs (mean 15.4/25, 73% high concern) and electricity bills (15.0/25, 74%) top the ranking, while reliability and regulatory hazards (power cuts, new taxes) sit lower. Across the seven hazards, the financial-pressure hazard "struggling to pay bills each month" produces the strongest model (McFadden R^2=0.15, 21 confirmed predictors), while the near-universally shared "heating and cooling costs" produces the weakest (R^2=0.04, 73% already in high concern)—the first appearance of the concern–predictability inversion discussed in Section 3.5. The composite model reaches the "good" threshold (McFadden R^2=0.11; Nagelkerke 0.32), and all seven hazard models pass the global likelihood-ratio test. Of the 41 confirmed Energy predictors, 11 are doubly robust.

The dominant signal is financial vulnerability. Utility-bill arrears is the single most consistent predictor across the seven hazards, and in the overall (composite) Energy model it carries one of the largest effect sizes in the entire study: respondents who have fallen behind on bills twice or more have nearly six times the odds of high concern (OR 5.76, 95% CI 2.59–12.78; +35.4 percentage points on the probability of being in the highest concern band). At the level of individual hazards, the same predictor sits at the top of the odds-ratio ranking for four of the seven hazards, with odds ratios between 3.0 and 4.1, and low confidence in understanding the energy bill carries the largest single per-hazard odds ratio (OR 6.05 for "more frequent power cuts"). Women and non-binary respondents, ethnic-minority respondents and multi-generational households also carry significantly higher concern, while membership of a political group is the one consistent protective predictor (OR 0.28–0.44 across five hazards), interpretable as institutional trust mediating concern rather than as politics reducing real exposure. Figure 4 and Table 2 present the confirmed odds ratios.

Read concretely, the picture is coherent and human. A respondent who has fallen behind on utility bills twice or more is nearly six times as likely to be highly concerned about energy hazards as someone with no arrears, and a respondent who struggles to understand their energy bill is six times as likely to be highly concerned about power cuts—these are the people for whom any additional cost or disruption lands hardest. The protective association with political-group membership should not be misread as politics reducing real exposure; rather, it most plausibly reflects greater institutional trust and a sense of having a voice, which dampens perceived risk. Taken together, the Energy findings point unambiguously at financial precarity as the axis along which concern is organised.

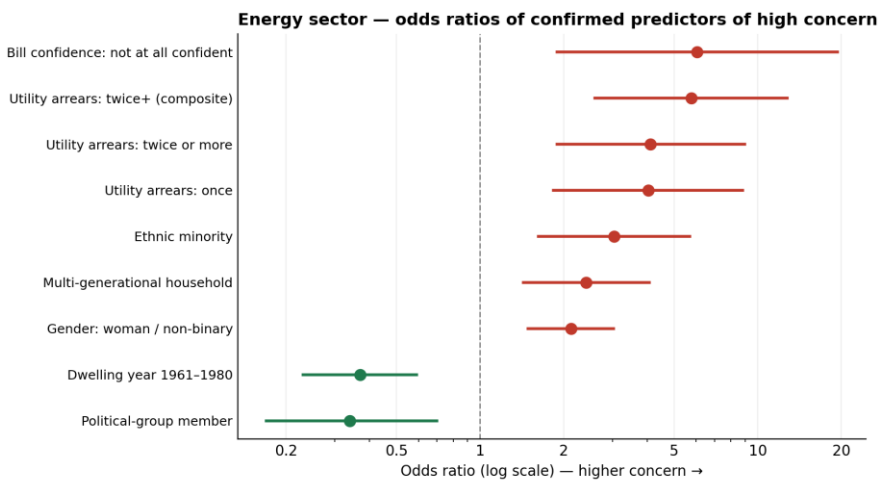

Fig. 4. Energy sector — odds ratios (95% CI, log scale) of confirmed predictors of high concern. Red = higher concern (OR>1); green = lower concern (OR<1). The dashed line marks OR=1.

Table 2. Energy sector — selected confirmed predictors with odds ratio, 95% CI and marginal effect on the probability of highest concern (MEM,

percentage points). The first row is from the overall composite model; the remainder are drawn from the per-hazard models.

Predictor	OR (95% CI)	p-value	MEM
Utility arrears: twice or more (composite)	5.76 (2.59–12.78)	<0.0001	+35.4
Bill confidence: not at all confident	6.05 (1.89–19.39)	0.0025	+39.0
Utility arrears: twice or more (per-hazard)	4.11 (1.89–8.96)	0.0004	+28.9
Utility arrears: once	4.03 (1.84–8.83)	0.0005	+25.5
Ethnic minority	3.04 (1.62–5.70)	0.0005	+21.3
Multi-generational household	2.41 (1.43–4.07)	0.0009	+13.6
Gender: woman or non-binary	2.13 (1.49–3.03)	<0.0001	+11.4
Political-group member	0.34 (0.17–0.70)	0.0035	−12.5

Headline for Energy: financial-pressure indicators—utility-bill arrears and low bill confidence—are the strongest predictors of Energy concern; respondents with two or more arrears episodes have nearly six times the odds of high concern in the overall model, one of the largest effect sizes in the whole study. Country differences in overall Energy concern are not statistically significant.

The Step 5c robustness battery clarifies how far each Energy finding can be pushed. Eleven of the 41 confirmed predictors are doubly robust—reliably selected by LASSO and reliably significant under subsample refitting—and this anchor the sector's narrative: utility-bill arrears (confirmed in three hazards plus the composite), gender, bill confidence, income band, dwelling year and ethnic minority. A further 25 are stably selected but fragile in retention; these are best read at the construct level (financial pressure, household composition, country macro effects) rather than as individual point estimates. The macro variables in particular—national electricity-consumption and price bands—are consistently selected but consistently fragile, meaning the data prefers them as candidates but their precise significance is sensitive to which respondents are observed. Direction, however, is essentially invariant: sign consistency averages 0.99 across the confirmed set, so statements about who is more concerned are reliable even where magnitude is not.

3.3 Housing sector (n = 263)

The Housing sample is the most coherent of the three (α=0.84, meaning respondents have an unusually consistent sense of what "housing concern" means) but is heavily weighted toward Germany (75%), with Portugal (19%) and Ireland (7%); it is therefore best read as a predominantly German housing-perception space, a caveat that shapes how its country-level findings should be interpreted. Thirteen hazards were modelled, the widest catalogue of the three sectors, spanning cost-of-living worries, dwelling-quality risks (damp, mould, cold-related ill health) and regulatory risks (bans on selling non-renovated homes, insurer reclassification).

The strongest and most policy-relevant finding is clean and striking: respondents in the lowest housing-cost band carry two to three times the odds of high concern, confirmed across four separate hazards (cold/damp health effects, more expensive home insurance, more frequent power cuts, and the composite), always positive, always with $p<0.005$. This is among the most robust confirmed cross-hazard findings in the whole study by full-sample criteria. The interpretation is intuitive once unpacked: the lowest housing-cost band is not a marker of affluence but, in this sample, a proxy for occupying older, cheaper, often poorer-quality dwellings—exactly the homes most exposed to retrofit costs, damp and cold, and reclassification risk. Housing's small and nationally-concentrated sample means most of its 36 confirmed predictors fall into the "stable-selection, fragile-retention" quadrant, so the right level at which to read these findings is the construct—housing-cost burden, dampness, dwelling quality—rather than the precise magnitude of any single variable. Even so, the direction is rock-solid (sign consistency 1.00 on this finding across every resample).

Housing is also the one sector with a significant country difference: Portuguese respondents report higher mean concern than German respondents (Tukey HSD $p=0.005$). Because Portugal has the lowest heating-degree-days of the three Housing countries, this gap cannot be climate-driven—it points instead to differences in housing-stock quality, a reminder that national building stocks and housing systems still shape how transition risks are felt even where individual-level patterns are shared.

Headline for Housing: respondents in the lowest housing-cost band carry two to three times the odds of high concern across four hazards, pointing policy directly at the most vulnerable group in the Housing transition—likely occupants of older, less well-maintained dwellings.

3.4 Transport sector (n = 358)

The Transport sample is the most balanced and homogeneous (Italy 57%, Spain 43%; α=0.79) and, accordingly, has the strongest robustness profile of the three sectors: of its 17 confirmed predictors, six are doubly robust and none is fragile on both tests. The six hazards concern the move away from internal-combustion-engine (ICE) vehicles—rising fuel and repair costs, loss of resale value, new taxes, city-centre access restrictions, longer or more complex journeys, and pollution exposure. Concern is highest among the more exposed or more vulnerable: women, unemployed respondents, frequent travellers, and people who do not speak the national language.

Two stories stand out. The first is an equity signal that purely cost-focused analyses would miss: the language barrier. Non-native-language speakers have roughly 70% higher odds of high concern overall (OR 1.76), rising to more than double (OR 2.30) for "longer or more complex journeys". Transition measures such as new low-emission-zone signage, route diversions and EV-charging information impose a real navigational burden on those who do not read the national language fluently—and the data show they feel it. The second story is, at first glance, surprising: it is the people most enthusiastic about electric vehicles who worry most about ICE-related hazards. Respondents with the most positive EV perception carry concern about rising fuel and repair costs +8.55 points higher (on the 1–25 scale) than the most EV-sceptical, and the largest single effect in the entire study—pollution-exposure concern by EV-switch likelihood, η^2=0.119—follows the same logic. The explanation is straightforward once stated: people who have understood the case for electrification are precisely those most alert to the downsides of staying with ICE vehicles. Finally, Italy and Spain show mirror-image patterns—Italians worry more about pollution exposure, Spaniards more about being restricted from city centres—driven by the same underlying reliance on cars.

Headline for Transport: women, unemployed respondents, frequent travellers and non-native-language speakers all carry significantly higher Transport concern, while remote workers carry less; inclusive, multilingual communication around EV and city-access measures is the clearest policy implication.

Table 3. Transport sector — selected confirmed effects from the per-hazard ANOVA (Step 7): the highest- and lowest-concern category of each predictor and the gap on the 1–25 concern scale.

Predictor (hazard)	Highest concern	Lowest concern	Gap
EV perception (fuel/repair costs)	Most positive (5)	Least positive (1)	+8.55
EV likelihood (pollution exposure)	Likely to switch	Unlikely to switch	+7.95
Travel frequency (longer journeys)	3+ days/week	Works from home	+3.48
Language (longer journeys)	Non-national speaker	National speaker	+2.75
Vehicle age (city restrictions)	Pre-2010 car	2010–2019 car	+4.14

Two readings follow. The EV-attitude gradients are intuitive and policy-meaningful: prospective EV adopters and EV-positive respondents are the most acutely aware of internal-combustion-engine downsides (rising fuel and repair costs, pollution exposure), so they report the highest concern—identifying the population segment most engaged with the clean-transport transition. The language-barrier and concrete-exposure signals (non-national speakers; owners of pre-2010 vehicles that are most likely to fail emission-zone criteria; daily commuters) are the most directly actionable, because the affected population can be identified from observable characteristics and reached with targeted, multilingual communication.

3.5 Cross-sector conclusions

Three findings hold across all three sectors. First, financial vulnerability is the strongest cross-sector pattern: although the specific variable differs by sector—utility arrears and low bill confidence in Energy, the lowest housing-cost band in Housing, unemployment in Transport—the substantive meaning is identical: respondents already under economic strain perceive transition hazards as more serious. This makes financial vulnerability the single most important factor for mitigation to address. Second, the concern–predictability inversion: the hazards with the highest mean concern have the weakest predictive models, because near-universal concern leaves little between-respondent variation to explain (Figure 5). High-mean-concern hazards therefore call for broad reassurance; moderate-mean-concern hazards are where targeted policy gains traction. Third, country differences are mostly absent: Energy and Transport show no significant cross-country differences in overall concern, and only Housing shows one (Portugal > Germany). The country a person lives in explains only about 1–4% of their concern—far less than their personal circumstances—so mitigation can be designed around socio-economic profiles and applied broadly across countries with similar EU context.

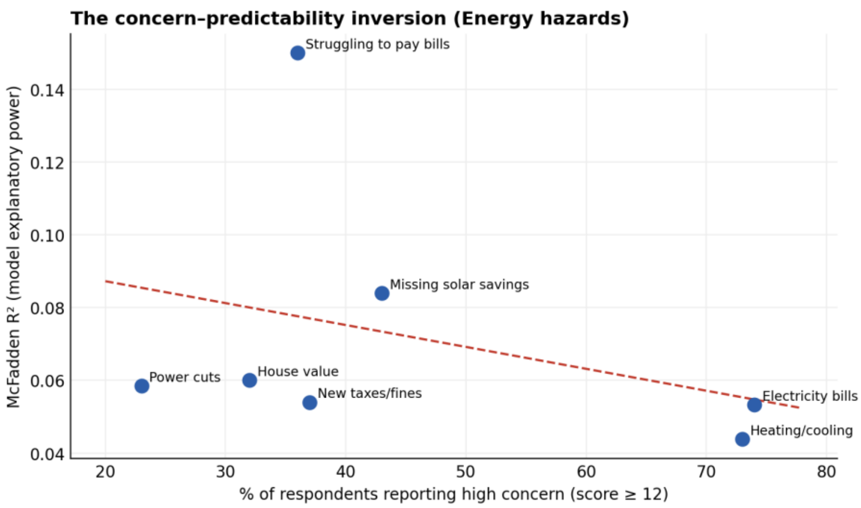

Fig. 5. The concern–predictability inversion (Energy hazards): the more widely a hazard is feared (x-axis), the less its concern can be explained by individual characteristics (McFadden R², y-axis).

A robustness gradient orders the sectors: Transport (6 of 17 confirmed predictors doubly robust, 0 fragile on both) is strongest, Energy (11 of 41 doubly robust) intermediate, and Housing (1 of 36 doubly robust) weakest—tracking sample size and country homogeneity (Table 4). Across all three sectors the direction of every confirmed effect is robust (sign consistency ≥ 0.94), so claims about which way a predictor pushes concern are defensible everywhere; claims about effect magnitude at the individual-predictor level are fully defensible only for the doubly robust predictors. Effects flagged with events-per-predictor warnings are reported for completeness but are not used as headline findings.

Table 4. Three-sector robustness comparison across confirmed predictors (Step 5c). MB = Meinshausen–Bühlmann bootstrap; SS = Shah–Samworth complementary-pairs subsampling.

Metric	Energy	Housing	Transport
Confirmed predictors	41	36	17
Sample size	407	263	358
Countries	5	3	2
SS mean selection stability	0.90	0.65	0.95
Mean retention	0.38	0.15	0.42
Mean sign consistency	0.99	0.94	0.99
Doubly robust	11	1	6
Fragile on both	4	23	0

In short, the three sectors tell one coherent story with sector-specific detail. Energy is driven by direct financial pressure (arrears, bill confidence) with a clear gender and ethnic-minority dimension; Housing concentrates concern in the lowest housing-cost band and is the one sector where national context matters; Transport is structured by EV attitudes and by a distinctive language-barrier equity signal. The common thread—financial vulnerability—and the common caveat—universal concerns resist targeting—hold in every sector, which is exactly what a transferable, sector-agnostic methodology should reveal.

The near absence of country effects deserves emphasis because it is good news for policy design. If concern were strongly shaped by nationality, every finding would need re-checking country by country before it could be acted on. Instead, the patterns of who is concerned and why—financial pressure, housing quality, language barriers—operate consistently across the EU contexts sampled, which means a measure that addresses energy-bill arrears or multilingual transport communication is likely to be relevant wherever those conditions exist, not just where it was first identified. This is the practical pay-off of representing countries through transferable macro indicators rather than through nationality: the analysis describes the conditions that produce concern, and conditions travel across borders in a way that passports do not. The one documented exception—higher Housing concern in Portugal than Germany—marks where national tailoring genuinely is warranted, and even there the analysis points to a transferable cause (housing-stock quality) rather than to nationality as such.

4. From prediction to policy: co-creation and the FITTER platform

The quantitative pipeline identifies which hazards are most feared and which groups are most concerned, but it does not, by itself, produce policy. The framework closes this gap through a participatory phase in which the risk evidence is handed to stakeholders to co-create mitigation measures, and an LLM-based agent then supports their qualitative assessment and tailoring.

Co-creation. A co-creation laboratory was run in each of the six FITTER use-case countries (Germany, Hungary, Ireland, Italy, Portugal and Spain), each following a five-meeting design—three in-person working sessions and two online validation sessions—held between September and December 2025. Labs engaged intermediary organisations representing disadvantaged groups, together with policymakers and civil-society actors, using challenge-led system mapping and backcasting from a 2050 just-transition vision. Across the six countries, 26 lab meetings engaged 237 external participants and produced a portfolio of 82 co-created mitigation measures (adjustments to existing policies and new proposals), classified by impact and feasibility (Figure 6). The measures converge on the same priorities the quantitative analysis identifies cost redistribution, progressive and automatic social tariffs, targeting of financially stretched and low-housing-cost households, administrative simplification, and inclusive, multilingual transition communication.

The design of the labs deserves a word, because it is what turns survey findings into policy. The first session was diagnostic: using realistic persona profiles and a “pentagonal problem” frame (environment; society and culture; economy and finance; policy and governance; science, technology and infrastructure), participants mapped the causes, actors and barriers behind sectoral inequalities. The second moved from diagnosis to design, working backwards from a 2050 vision of a fair transition to the measures needed to reach it. The third finalised and scored each measure on impact and feasibility. Engaging intermediary organisations—bodies that represent or work closely with disadvantaged groups—rather than vulnerable individuals directly was a deliberate choice: it secured the necessary policy expertise while avoiding the power imbalances and burden that direct participation can impose on already-stretched people. The striking result is how closely the stakeholders’ priorities track the statistics: a process that began from lived experience and a process that began from ordinal regressions arrived at the same conclusion—that affordability and inclusion are where mitigation must concentrate.

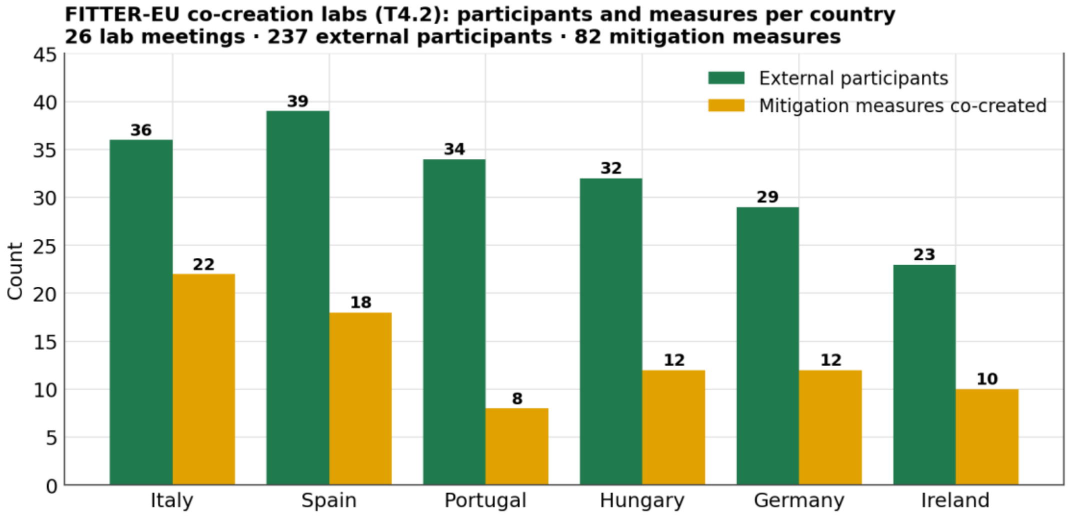

Fig. 6. FITTER-EU co-creation labs: external participants and co-created mitigation measures per country (September–December 2025). The six national labs engaged 237 participants across 26 meetings and produced 82 measures.

Qualitative assessment via an LLM agent. The co-created measures and the quantitative evidence are made navigable through “Dr. Transition”, a conversational agent within the FITTER Digital Platform. The agent draws on a web-based knowledge base and is fine-tuned on FITTER-EU outputs so that the sources it surfaces align with the project’s eco-social-justice goals. For a given sector and territory, it surfaces the leading hazards, identifies the groups facing significantly higher perceived exposure, and helps a policymaker reach the root causes and tailor mitigation measures (Figure 7). The detailed design of this LLM-guided process is reported in a companion paper in this Special Session [3].

The reason for adding a conversational layer is practical: the statistical findings of Section 3, however rigorous, are not in a form a busy policymaker can readily act on. A forest plot of odds ratios and a quadrant of robustness classifications carry exactly the nuance the analysis demands—but they are not how policy questions are usually asked. The agent translates between the two. A user can ask, in plain language, which groups are most concerned about a given hazard in their region; the agent answers from the confirmed evidence, flags where the evidence is firm and where it is only suggestive and helps think through which co-created measures fit the situation. In doing so it preserves the discipline of the underlying analysis—never presenting a fragile finding as a firm one—while removing the barrier that technical output would otherwise place between the evidence and its use. The agent assists reasoning; it does not replace it, and it issues no prescriptive verdicts.

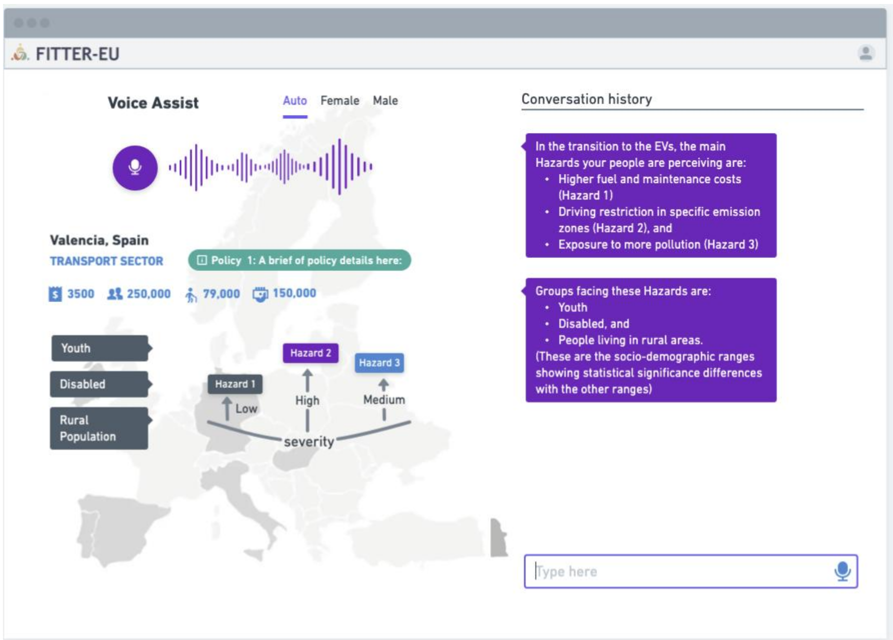

Fig. 7. Illustrative dialogue with the “Dr. Transition” LLM agent, linking the quantitative hazard and disadvantaged-group evidence to a sector- and territory-specific mitigation discussion.

5. Discussion

Financial vulnerability is the structuring axis of perceived risk. Across Energy, Housing and Transport, indicators of financial strain are the most consistent predictors of elevated concern. This aligns the findings with the distributive dimension of energy justice [13] and with the vulnerability pathways identified in the FITTER framework (D2.1, D3.1): respondents experiencing economic strain have fewer resources to absorb transition-related costs and therefore perceive transition hazards as more serious. The consistency of this pattern across three independent analyses—different sectors, different questionnaires, different confirmed variables, but the same underlying meaning—gives it real evidential weight and suggests that perceived hazard of the twin transition is concentrated among precisely those least able to absorb additional costs. Affordability and accessibility thus emerge as central dimensions of a just transition, closely connected to wider questions of inclusion and fair access to opportunity.

Housing concern and environmental justice. The Housing sector provides the single most actionable signal in the study: respondents in the lowest housing-cost band carry two to three times the odds of high concern across four hazards, consistently and in the same direction. Although the Housing sample's size and German concentration mean this is best read at the construct level, it points directly at the most structurally vulnerable group—likely occupants of older, less well-maintained dwellings—and at a clear policy target: protecting this group from transition-related cost shocks (retrofits, insurance reclassification, low-emission heating requirements). The finding that Portuguese respondents are more concerned than German respondents despite Portugal's milder climate reinforces that housing-stock quality, not weather, drives much of this concern.

Transport concern and spatial/communicative justice. The Transport findings extend the justice lens beyond affordability to communication and access. The language-barrier effect—non-native-language speakers reporting markedly higher concern about longer journeys and city-centre restrictions—identifies an equity dimension that cost-focused analyses miss transition measures such as new low-emission-zone signage, route diversions and EV-charging information impose a disproportionate navigational burden on linguistic minorities. Inclusive, multilingual transition communication is therefore not a peripheral courtesy but a substantive equity measure.

Universal concerns and targeted concerns require different policy logics. The concern–predictability inversion has a direct policy reading: the cost-of-living hazards that worry almost everyone (heating, electricity, fuel) need broad, population-wide reassurance and cost protection, whereas the moderate-concern hazards—where individual characteristics distinguish the concerned from the unconcerned—are where targeted, group-specific measures can be efficiently aimed. Treating these two classes of hazard with the same instrument would either spread targeted resources too thin or leave universal concerns under-addressed.

Context matters less than circumstance. The general absence of country effects indicates that the patterns of who is concerned and why operate consistently across EU national contexts, so policy informed by these findings is likely to transfer broadly rather than needing country-specific tailoring. The one exception—higher Housing concern in Portugal than Germany—likely reflects real differences in housing-stock quality and merits attention in Housing-specific national policy.

What the analysis can and cannot support. The pipeline identifies associations, not causes, and its confirmation regime is deliberately conservative: only predictors surviving selection, estimation, multiple-testing correction and resampling are reported. The most defensible findings are those confirmed across multiple hazards with consistent sign and $p<0.005$ and no instability warning—the lowest-housing-cost-band finding being the exemplar. Effects based on very small subgroups are flagged and not used for policy claims. The robustness gradient (Transport > Energy > Housing) means Transport and Energy support predictor-level claims for their doubly robust predictors, while Housing's findings are best read at the construct level.

Limitations. The achieved samples fall below the design target and are unevenly distributed across countries (Housing is 75% German), so country-level and small-subgroup inferences are less precise; one use-case country (Hungary) is absent from the quantitative dataset owing to panel composition. Quota sampling via an online panel introduces self-selection bias mitigated, but not eliminated, by EU-SILC quotas, and the analysis measures perceived rather than realised hazard. On the participatory side, engagement was uneven across countries and intermediary organisations represented rather than embodied disadvantaged groups. These constraints motivate targeted oversampling, complementary national-panel recruitment, and integration of the validated indices into the FITTER Digital Platform for automated, country-aware policy briefs.

6. Conclusions and further developments

This paper has presented and empirically applied a rigorous, reproducible nine-step methodology that predicts and addresses the hazards European citizens perceive in the twin transition, with explicit attention to disadvantaged groups. By combining LASSO shortlisting, ordinal proportional-odds modelling, likelihood-ratio confirmation with false-discovery-rate control and an events-per-predictor cap, four resampling robustness tests, and odds-ratio and ANOVA interpretation, the pipeline converts harmonised, EU-SILC-aligned perception surveys into interpretable, profile-specific and honestly bounded risk intelligence. Applied identically across Energy, Housing and Transport with 1,028 quality-controlled responses, it shows that financial vulnerability is the strongest cross-sector predictor of concern, that the most widely shared hazards are the least group-specific, and that national context explains only a small share of concern. The participatory phase—26 national-lab meetings, 237 participants, 82 co-created measures—and the "Dr. Transition" agent demonstrate how the evidence can be operationalised participatively and made navigable for inclusive policy design.

Future work will integrate the full validated dataset into the FITTER Digital Platform, extend the macro-variable layer for richer cross-country profiling, and broaden the framework to additional sectors and countries, strengthening FITTER-EU's contribution to scalable, anticipatory and participatory governance for a fair and inclusive twin transition across Europe.

Beyond the specific findings, the contribution we most wish to underline is methodological and practical. The twin transition will generate a steady stream of new policies, new hazards and new affected groups; a one-off study cannot keep pace with it. What endures is a transparent, reproducible procedure that can be re-run as conditions change—one that is honest about uncertainty, separates firm conclusions from suggestive ones, attributes concern to transferable circumstances rather than to nationality, and feeds directly into a participatory process and a decision-support tool. Used this way, the methodology is less a verdict than an instrument: a way for policymakers to keep asking, sector by sector and year by year, who feels at risk from the transition and how that risk can be fairly addressed. Returning to the household with which this paper began, the value of such an instrument is that the cascade of obligations it faces need no longer be designed in ignorance of how it lands—its concerns become visible, quantifiable and, ultimately, addressable.

7. Acknowledgements

This work is funded by the European Union under the project FITTER-EU (Grant Agreement No. 101132546). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them. The authors thank the FITTER-EU consortium partners and the stakeholders who participated in the national co-creation labs.

Funded by the European Union

8. References

1. Molero, G. D., Santarremigia, F. E., Kapoor, A., Zarehparast Malekzadeh, M., Leva, M. C., Clavero, S.: A methodological framework for a fair and inclusive twin transition in the transport sector. Transportation Research Arena (TRA) (2026).
2. Kapoor, A., Molero, G. D., Santarremigia, F. E., Zarehparast Malekzadeh, M., Malviya, A. K., Flynn, T.: An evolutionary tree framework for a fair and inclusive twin transition in Europe. ERSA Special Session (2025).
3. Kapoor, A., Zarehparast Malekzadeh, M., Molero, G. D., Santarremigia, F. E., Bjegojevic, B., MacNamara, N., Salomón, O.: Anticipating the future: a pedagogical policy-making process guided by AI. 65th ERSA Congress, Sofia (2026).
4. Peer, E., Brandimarte, L., Samat, S., Acquisti, A.: Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. Journal of Experimental Social Psychology 70, 153–163 (2017).
5. Memon, M. A., Ting, H., Cheah, J.-H., Ramayah, T., Chuah, F., Cham, T.-H.: Sample size for survey research: review and recommendations. Journal of Applied Structural Equation Modeling 4(2), i–xx (2020).
6. Tibshirani, R.: Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society, Series B 58(1), 267–288 (1996). [see also Hastie, Tibshirani & Friedman, The Elements of Statistical Learning, Springer (2009)]
7. McCullagh, P.: Regression models for ordinal data (with discussion). Journal of the Royal Statistical Society, Series B 42(2), 109–142 (1980).
8. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society, Series B 57(1), 289–300 (1995).
9. Brant, R.: Assessing proportionality in the proportional-odds model for ordinal logistic regression. Biometrics 46(4), 1171–1178 (1990).
10. McFadden, D.: Conditional logit analysis of qualitative choice behaviour. In: Zarembka, P. (ed.) Frontiers in Econometrics, pp. 105–142. Academic Press (1974).
11. Meinshausen, N., Bühlmann, P.: Stability selection. Journal of the Royal Statistical Society, Series B 72(4), 417–473 (2010).
12. Cronbach, L. J.: Coefficient alpha and the internal structure of tests. Psychometrika 16(3), 297–334 (1951).
13. Heffron, R. J., McCauley, D.: The concept of energy justice. Energy Policy 105, 658–667 (2017).
14. Sangiuliano, M., Salomón Balsamo, O.: Deliverable 4.2 — Methodology of the co-creation process and report on mitigation measures. FITTER-EU Project (2026).
15. Santarremigia, F. E., Molero, G. D., Zarehparast Malekzadeh, M., Kapoor, A.: Deliverable 4.4 — Computational analysis report. FITTER-EU Project (2026).
16. Muench, S., Stoermer, E., Jensen, K., Asikainen, T., Salvi, M., Scapolo, F.: Towards a green and digital future. Publications Office of the European Union (2022).