Discussion Paper Series

IZA DP No. 18674

May 2026

Labor-Market Consequences of Cross-Border Employment: A Machine Learning Approach

Andrea Albanese
Luxembourg Institute of Socio-Economic Research (LISER), Ghent University, UCLouvain, IZA@LISER and GLO

David Marguerit
Luxembourg Institute of Socio-Economic Research (LISER) and University of Luxembourg

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18674 May 2026

Labor-Market Consequences of Cross-Border Employment: A Machine Learning Approach*

Abstract

Cross-border work is expanding in the EU, yet its labor-market effects on the cross-border workers themselves remain largely undocumented. Using linked Belgian administrative registers that identify cross-border spells in Luxembourg, we estimate the effects of cross-border employment on post-return labor-market outcomes through dynamic double machine learning. Returnees face a short-run employment penalty that fades with cross-border tenure and time since return. They are also more likely to receive Belgian unemployment benefits than comparable stayers, with higher daily benefit levels among recipients.

JEL classification

J61, J65, J64, I38, C21, C14

Keywords

cross-border commuting, return migration, unemployment insurance, EU labor mobility, administrative data

Corresponding author

David Marguerit

david.marguerit@liser.lu

* We thank Julio Garbers, Christina Gathmann, Martin Huber, Francis Kramarz, Lukáš Lafférs, Joël Machado, and Konstantinos Tatsiramos for constructive comments. We are also grateful to participants of the 2023 Workshop on Cross-border Labor Mobility & Data Linkage, the 17th IAAEU Workshop on Labor Economics, and the LISER internal seminar. Andrea Albanese acknowledges financial support from the CORE program of the Luxembourg National Research Fund (FNR) (CrossEUwork – 11700060). This paper uses confidential administrative data from the Belgian Crossroads Bank for Social Security (BCSS) under contract no. ART5/18/033, which precludes redistribution. Researchers may apply for access through the BCSS Information Security Committee, and the authors will provide all code and instructions to replicate the analysis from the raw extracts. Claude (Anthropic) was used for language editing; all content, analysis, and final wording are the authors' responsibility.

1 Introduction

Cross-border work (CBW) has become a structural feature of the European Union labor market. In 2024, 2.3 million EU residents were employed outside their country of residence, more than double the 2005 figure (Eurostat, 2024). Flows are concentrated in regional clusters where wage differentials, geographic proximity, and the social-security coordination framework make daily commuting across a national border economically viable. Despite the growing scale of intra-EU cross-border work, little is known about the labor-market consequences for cross-border workers themselves.

The existing literature examines the consequences of CBW for local labor markets in destination and origin countries. Prior studies analyze effects on native workers and firms in the destination country (Dustmann et al., 2017; Beerli et al., 2021), on regional labor markets, firms, and wages in origin countries (Dicarlo, 2022; Bütikofer et al., 2024; Hafner and Hedtrich, 2024; Illing, 2025), and on non-movers in sending regions (Dodini et al., 2024, 2025). None, however, examines the consequences for cross-border workers themselves upon return.

This study addresses three questions. First, we identify predictors associated with entry into and exit from CBW. Second, we estimate the employment impact upon return from CBW. Third, we examine the integration of returning cross-border workers into the welfare system of their country of residence. Addressing these questions requires longitudinal data capable of tracking individuals through entry into, duration of, and exit from cross-border employment, data that are rarely available.

Using linked Belgian administrative registers, we study the Belgium-Luxembourg corridor, where nearly 30 percent of the active population in the Belgian Province of Luxembourg commutes abroad (Eurostat, 2024) – the highest outgoing cross-border employment rate in the EU. We identify cross-border work spells in Luxembourg through compulsory sickness-insurance registries, a source rarely exploited for outbound cross-border employment because such spells are typically unrecorded in origin-country social-security systems. The sample covers approximately three-quarters of individuals born between 1973 and 1990 who resided near the Luxembourg border at some point between 2008 and 2013, with quarterly labor-market histories from 2003 to 2017.

Selection into CBW is non-linear and operates on several margins, as revealed by random-forest analysis. Commuting time to Luxembourg City accounts for the largest share of explained variation in entry, with a sharp non-linear gradient. Predicted entry is elevated at both the bottom and the top of the Belgian earnings distribution, consistent with the coexistence of a necessity-driven margin and a skill-driven margin. Prior CBW exposure of the worker, the partner, and other household members predicts entry with concave returns suggestive of learning and network effects. Selection at the exit margin shows analogous structure: commuting time predicts return with the same non-linearity but reversed sign, the last Belgian salary shows a similar U-shape, and prior labor-market histories distinguish instability-driven exits from constrained persistence in CBW. These patterns make standard low-dimensional conditioning inadequate for causal analysis and suggest the need for more advanced techniques.

We rely on causal machine learning to provide the first evidence on the labor-market consequences of CBW upon return. We estimate the average treatment effect of CBW spells of one to three years followed by return, on outcomes measured up to three years after return. Identification rests on sequential conditional independence given a covariate vector that includes ten years of quarterly labor-market history, household and household-member CBW exposure, and detailed socio-demographic characteristics, totaling more than 200 predictors. We use dynamic double machine learning (DML) to flexibly control for this high-dimensional history without imposing a parametric function (Bodory et al., 2022; Bia et al., 2024). The credibility of the design hinges on the assumption that, conditional on this history, treatment assignment in each period is independent of potential outcomes. We discuss this assumption explicitly, report covariate-balance and common-support diagnostics, and quantify robustness to unmeasured confounding through E-values (VanderWeele and Ding, 2017).

Three main findings emerge. First, the raw post-return employment gap between returnees and stayers exceeds 30 pp, but the conditional dynamic DML estimate is around 10 pp after one year of CBW. About three quarters of the raw difference between returnees and stayers reflects observable selection on labor-market histories, household structure, and geographic location, rather than a causal consequence of CBW. Studies that do not condition on detailed pre-period history are likely to misattribute selection to the CBW experience itself. The remaining causal employment penalty is short-lived: it shrinks with CBW duration and fades to close to zero after three years of CBW or three years after return. Second, returnees are substantially more likely to access the residence-country unemployment insurance (UI) system, and this probability rises with time spent in CBW; among recipients, returnees receive significantly higher daily benefits than stayers. Longer spells progressively shift returnees from inactivity into registered unemployment — consistent with stronger labor-force attachment — so that by a three-year spell, reduced inactivity accounts for more than half of the total unemployment effect. Third, we find no sizeable wage premia upon return.

The paper contributes to three literatures. First, we provide individual-level evidence on post-return labor-market outcomes for cross-border workers. The return-mobility literature has studied post-return employment, wages, and occupational mobility for residential migrants (e.g., Co et al., 2000; Reinhold and Thom, 2013; Abramitzky et al., 2019). However, mechanisms emphasized in these studies – relocation costs, savings under host-country uncertainty, and re-entry frictions from prolonged absence – do not apply to daily commuters who retain home networks, consumption, and local labor-market exposure throughout the spell (Dustmann et al., 2017; Beerli et al., 2021). Commuters and migrants also differ systematically in age, education, and household structure (Huber, 2014), so existing estimates do not translate.

Second, we provide, to our knowledge, the first evidence in any return-mobility literature on returnees' uptake of residence-country welfare benefits. The welfare-magnet and public-transfer-assimilation literatures study migrants' benefit receipt in destination countries (e.g., De Giorgi and Pellizzari, 2009; Agersnap et al., 2020; Suari-Andreu and van Vliet, 2023). Neither this literature nor the return-migration literature speaks to returnees' uptake of residence-country benefits funded by contributions accrued abroad. Our estimates quantify this channel.

Third, we analyze the predictors of CBW entry and exit from individual-level administrative histories, using random forests on a high-dimensional covariate set. Existing work on CBW determinants relies on cross-sectional surveys, aggregate flows, or stated intentions with parametric specifications (e.g., Ahrens et al., 2020), as does the return-mobility literature on residential migrants (Dustmann, 1997, 2003; Dustmann et al., 2011; Bijwaard et al., 2014). We document non-linear selection on commuting time and earnings on both margins, which standard parametric specifications would miss.

The paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 presents the empirical strategy. Section 4 reports results. Section 5 concludes.

2 Setting and Data

2.1 Institutional setting

Cross-border commuting is concentrated in several regional clusters across Europe, most prominently the Greater Region of Luxembourg, the Swiss border regions, and parts of the Central and Eastern European frontier.[1] We conduct our analysis in the Greater Region of Luxembourg, where cross-border workers account for roughly half of Luxembourg employment, and for 15 to 30 percent of the active population in the neighboring NUTS-2 regions of France, Belgium, and Germany (Eurostat, 2024; STATEC, 2025). Within this cluster, the Belgian Province of Luxembourg displays the highest outgoing CBW rate in the EU, with almost 30 percent of its active population commuting abroad.

The corridor we study is characterized by low cross-border frictions, including a shared official language, and large wage differentials. In 2016, the median net monthly salary for a Belgian cross-border worker aged 25 to 43 was around €3,000, against €1,800 for a resident employed in Belgium's Province of Luxembourg (IGSS, BCSS data). Commuting times to Luxembourg are substantial during peak hours (47.5 minutes on average against 24.4 for local workers), generating a trade-off between higher wages abroad and time costs at home (Godefroy et al., 2021; Albanese et al., 2022).

Cross-border employment in the EU is governed by Regulation (EC) 883/2004, which assigns insurance and contributions to the country of work. A Belgian resident employed

[1] Substantial outgoing CBW rates are observed in Lorraine in France and Trier in Germany (both around 14 percent of the active population), Vorarlberg in Austria (10 percent), Western Transdanubia in Hungary (10 percent), Franche-Comté, Grand Est, and Alsace in France (around 10 percent), Lubuskie in Poland (8 percent), and Freiburg in Germany (6 percent) (Eurostat, 2024).

in Luxembourg therefore pays Luxembourg contributions and accrues Luxembourg pension and healthcare rights, while receiving healthcare in Belgium through a sickness fund financed by the Luxembourg system. Unemployment insurance is the exception. Under Article 65, unemployed cross-border workers must claim benefits from the country of residence, not the country of work, and insurance periods accrued abroad are aggregated to satisfy residence-country eligibility.[2] A Belgian-resident worker who loses a Luxembourg job therefore registers as unemployed in Belgium and claims Belgian benefits, with Luxembourg work history counted toward Belgian eligibility thresholds. Article 62(3) further requires that the benefit amount be computed on the gross salary received in Luxembourg. The Belgian benefit schedule during our sample window is degressive, applied to a reference wage capped at approximately €2,598 per month in 2017. Because Luxembourg gross wages systematically exceed Belgian gross wages in comparable occupations, applying Belgian replacement rates to a Luxembourg reference wage pushes returnees to the Belgian ceiling more often than stayers.

2.2 Data and sample

We use Belgian administrative data from the Crossroads Bank for Social Security (BCSS), which links Belgian National Registries and Social Security institutions. The data span 2003 to 2017 and include detailed quarterly labor-market histories and demographic information. We supplement them with average commuting times by car and public transport from each individual's neighborhood to Luxembourg City during peak hours, computed using Google API Directions and OpenTripPlanner. Cross-border employment is identified through Belgian health-insurance registries (mutualities), with which CBW workers are required to register; this information is rarely available in origin-country social-security data, where outbound CBW spells typically leave no contribution trace. The data do not include the occupation held in Luxembourg, which limits our ability to measure within-spell job characteristics. We observe job characteristics and earnings only upon

[2] Article 65 requires the country of last employment to reimburse the residence country for the first three months of payments (five if the worker completed at least twelve months of employment). Because Belgian unemployment-benefit duration during our sample period was effectively open-ended, Belgium absorbs the bulk of the fiscal cost of returnee unemployment benefits. Under the EU coordination revision agreed in April 2026 (Council of the European Union, 2026), cross-border workers with at least 22 uninterrupted weeks of activity in the country of last employment would instead claim unemployment benefits from that country.

return.

The population of interest comprises individuals born between 1973 and 1990 who resided in the Belgian provinces of Luxembourg or Liège at some point between 2008 and 2013. The sampling frame is stratified geographically, oversampling residents close to the Luxembourg border, with design weights equal to the inverse of the inclusion probability (Albanese et al., 2024). All estimates use these weights. The initial sample contains 86,500 individuals, about 74 percent of the targeted population. We restrict attention to individuals living in the study area at the end of any year between 2013 and 2015 and not yet employed in Luxembourg, and we observe their CBW status in subsequent years and their post-return outcomes for up to three years. Pooling the 2013, 2014, and 2015 cohorts yields a final sample of 147,102 observations on 48,975 individuals.

About 2.4 percent of the sample transitions into CBW within a year. Among CBW workers, 19.4 percent return to Belgium after one year. On average, cross-border workers earn substantially less in their last observed Belgian job than stayers (daily wage of €68 against €91), are more likely to be in the residual inactivity category in the year before entry (51 percent against 15 percent), and have lower prior employment stability. They are also younger and more likely to be male. Full descriptive statistics are available in Appendix Table A1.

3 Empirical Strategy

3.1 Treatment sequences and estimand

Each individual's CBW status is observed at the end of each year and takes three values: 0 if the individual lives in Belgium and does not work in Luxembourg, 1 if the individual lives in Belgium and works in Luxembourg, and 2 if the individual no longer appears in the Belgian registers due to out-migration or death.[3] The treatment of interest is a sequence $\underline{d}_T = [d_1, \ldots, d_T]$ that combines a CBW episode and a post-return horizon, for example $[0,1,0]$ (one year of CBW, outcome measured one year after return), $[0,1,1,0]$

[3] The end-of-year frequency limits the explosion of possible treatment paths. CBW spells that are not held at year-end account for about 15 percent of all CBW transitions, generating a modest undercount. Contamination of the control group is small (0.4 percent).

(two years of CBW, one year after return), or $[0,1,1,1,0]$ (three years of CBW, one year after return). The control sequence $\underline{d}_T^* = [0,0,\ldots,0]$ corresponds to individuals who never leave Belgium.

The target parameter is the average treatment effect on the population of first-period CBW entrants, $\Delta(\underline{d}_T, \underline{d}_T^*) = E[Y_T(\underline{d}_T) - Y_T(\underline{d}_T^*) \mid D_1 = 1]$. The estimated effects describe what happens to workers who go to Luxembourg and come back, and should be distinguished from the effect of permanent CBW.

3.2 Identification and dynamic DML

Identification rests on sequential conditional independence: in each period, treatment assignment is independent of potential outcomes conditional on the covariate path realized up to that point. The credibility of this assumption depends on the richness of the conditioning set. We condition on 213 predictors, derived from roughly 80 explanatory variables, including ten years of quarterly labor-market history (employment, self-employment, unemployment with and without job-search exemption, inactivity with and without transfers, occupation, sector, firm size, individual and firm fixed effects from an AKM model estimated on pre-2013 data, household income, partner and household-member CBW exposure), demographics (age, gender, nationality, education, migrant origin, Luxembourgish origin, household composition and number of children by age band), and geographic indicators (commuting time to Luxembourg City by car and public transport, local unemployment rate at the NUTS-3 level). Table 1 lists the full set of control variables.

Two arguments support the plausibility of sequential CIA in this setting. First, the persistence of unobservables is identified through their reflection in observable trajectories. Ability, motivation, attachment to the labor market, risk preferences, and informal job-search effort are typically stable over a decade, and over a ten-year quarterly window they translate into observable patterns of employment continuity, earnings levels and growth, occupational mobility, and benefit receipt (Lechner, 2010; Chabé-Ferret, 2015; Daw and Hatfield, 2018). We condition on all of these. Second, geographic sorting is observed: commuting time by car and public transport, and local unemployment rates, are explicit covariates. The remaining threats are post-baseline shocks specific to CBW workers in the last Belgian period before entry, in particular the arrival of a Luxembourg job offer. Such offers are unobservable by construction, but the dimensions along which Luxembourgish employers select workers are themselves captured in the ten-year history we condition on, limiting the scope for residual selection. Section 4.4 quantifies how strong such residual selection would have to be to overturn the main estimates.

We estimate the ATE with the dynamic DML procedure of Bodory et al. (2022), applied to sequences of up to four periods. At each transition node along the sequence, a period-specific propensity score is estimated conditional on the covariate vector available at that point, and the mean potential outcome along each path is recovered by reweighting the observations that actually followed the path by the inverse product of the propensity scores along it, augmented by nested outcome regressions delivering double robustness. To target the ATE on first-period CBW entrants, we follow the weighted-score representation of Bodory et al. (2022, Section 4) and multiply the efficient score by $g(X_0)/\Pr(D_1 = 1)$, where $g(X_0)$ coincides with the first-period propensity score. Both the outcome and the propensity-score components are estimated with random forests, with class weights inversely proportional to outcome frequencies at each node to address the rare-event nature of CBW transitions (Chen et al., 2004; Akbani et al., 2004); we report sensitivity to extreme gradient boosting in the appendix. K-fold cross-fitting ensures out-of-sample nuisance predictions, with the conditional and nested outcome means fitted in non-overlapping subsamples to avoid correlation between successive plug-in steps (Chernozhukov et al., 2018; Bodory et al., 2022). We cluster at the individual level since the same individual can appear in up to three cohorts.

When the outcome is observed only for a subset of individuals, as for wages (conditional on employment) or benefit amounts (conditional on UI receipt), conditioning on a post-treatment selection variable would bias the estimate. We address this within the dynamic framework using the equivalence in Bia et al. (2024), expanding the state space at the outcome period to include the relevant selection node. The credibility of the sample-selection correction rests on the same sequential CIA invoked for treatment assignment, since the main drivers of employment selection are themselves observed in the conditioning history.

The full formal exposition of the score, the weighted ATE construction, the dynamic selection diagrams, and the implementation details are in Appendix C.

4 Results

4.1 Selection into and out of cross-border employment

We first describe the selection into and out of cross-border employment, separately for entry (next-year transition from non-CBW to CBW, N=144,600 with a 2.4 percent base rate) and for exit (next-year transition from CBW to Belgian residence without CBW, N=3,708 with a 19.4 percent base rate). To interpret the random-forest predictions, we rely on variable-importance scores and partial-dependence plots (PDPs) estimated on the same covariate set as our DML propensity-score models. Full results are in Appendix B.

Figure 1 displays PDPs for the main predictors of entry and exit. PDPs trace how the predicted probability of the outcome varies with one predictor while averaging over the empirical distribution of all others, and should be read as descriptive net associations rather than causal effects (Hastie et al., 2009). Predictive performance is strong for entry (out-of-bag AUC = 0.867, with PR-AUC of 0.221 against a base rate of 0.024, an eight-fold improvement over chance) and modest for exit (AUC = 0.665), which is unsurprising because we observe Belgian-side covariates only and lack information on the Luxembourg job that would predict its termination.

Three patterns emerge from the PDPs, each highly non-linear. First, commuting time to Luxembourg City is the dominant predictor on both margins, with opposite signs: predicted entry declines steeply at short commutes and flattens beyond a threshold, while predicted exit is flat at short commutes and rises rapidly beyond it (Figure 1, panels (a) and (g)), consistent with a tipping point at which commuting burdens become difficult to sustain. Second, earnings-related variables predict both margins with U-shaped profiles. Predicted entry is elevated at both ends of the current Belgian labor-income distribution (panel b), consistent with the coexistence of a necessity-driven margin (workers pushed by limited domestic options) and a skill-driven margin (workers attracted by the Luxembourg wage premium). An inverted U-shape appears for ten-year average earnings (panel c). At exit, last salary in Belgium and annual household revenue show analogous U-shapes (panels h and i), with high return propensities at both ends of the distribution. Third, prior CBW exposure predicts entry with concave returns suggestive of learning effects: short initial spells substantially raise the predicted probability of subsequent entry, but additional exposure adds little (panel d).

Two further patterns reinforce these readings. Education shows a positive gradient at entry (panel e) and a negative gradient at exit (panel j), consistent with skill differentiation on both margins. Labor-market history differentiates the exit margin: prior residual inactivity predicts longer spells (panel k, declining gradient), while prior unemployment predicts shorter spells (panel l, rising gradient), distinguishing constrained persistence with limited outside options from instability-driven exit. Current-status indicators reinforce these readings: workers currently inactive at baseline are substantially more likely to enter CBW (panel f), consistent with the residual-category margin already absorbed by the conditioning set.

The implications for identification are direct. Selection into CBW is jointly governed by accessibility, dual-margin earnings dynamics, and current labor-market status. With 213 predictors and unknown functional forms, parametric specification would require ad hoc choices across nonlinearities and interactions. The DML procedure with random forests learns the functional form from the data without parametric restrictions on functional form.

4.2 Effect on labor-market states

Figure 2a-2c reports the effect of CBW experience on three mutually exclusive labor-market states one year after return, by CBW spell length. The three states are employment, registered unemployment (with or without UI receipt), and inactivity. Without controls, returnees show a raw employment penalty exceeding 30 pp (Panel a, blue bars). Once observable selection is accounted for through dynamic DML, the employment penalty shrinks substantially. After one year of CBW, returnees are about 10 pp less likely to be employed than comparable stayers. The penalty falls to about 7 pp after two years of CBW and is statistically indistinguishable from zero after three years.

The corresponding registered-unemployment effects (Panel b) are positive and significant across spell lengths: about +12 pp after one or two years of CBW and about +17 pp after three years. Inactivity effects (Panel c) are not statistically distinguishable from zero, although point estimates trend negative with spell length, from about -3 pp after one year to -12 pp after three.

Because employment, unemployment, and inactivity are exhaustive and mutually exclusive, the estimates approximately sum to zero in finite samples and yield a natural decomposition. For one-year CBW spells, about 80 percent of the increase in registered unemployment comes from reduced employment, with the remainder from reduced inactivity. For two- and three-year spells, this share falls to roughly 55 and 45 percent respectively. Under a pure job-loss story, the rise in unemployment would be matched one-for-one by a fall in employment, with no contribution from inactivity. The data depart from this benchmark, with a falling employment share and an increasingly negative point estimate on inactivity, consistent with individuals detached from the labor force becoming visible in the UI register after a longer CBW spell.

Figure 3 extends the analysis up to three years after return for one-year CBW spells, and up to two years for two-year spells. For one-year CBW, the employment penalty is roughly -10 pp at one and two years after return and attenuates to about -7 pp by the third year. For two-year CBW spells, the penalty falls from about -7 pp at one year to a value indistinguishable from zero at two years. Unemployment effects similarly attenuate. Inactivity effects remain insignificant throughout. Sample sizes shrink with the horizon, particularly for two-year spells, so medium-run estimates should be read with caution.

These patterns are similar to post-return employment dips for residential migrants (Saarela and Finnås, 2009; Piracha and Vadean, 2010; Xia et al., 2025), but the underlying mechanisms differ. Because commuters live in Belgium throughout the spell, they do not lose the broader home-country capital that residential migrants forfeit during prolonged absence, nor are their choices during the spell shaped by planned-return dynamics that drive residential migrants' savings and labor-supply decisions (Chabé-Ferret et al., 2018). What commuters do lose is the employment side of this capital: domestic employer relationships, sectoral exposure, and recruitment channels interrupted during the CBW spell. Two mechanisms specific to this setting are consistent with the attenuation we estimate. First, attenuation with time since return reflects job-search frictions in the home labor market, with the penalty decaying as the worker rebuilds these connections (Topa, 2011). Second, attenuation with CBW duration reflects the value of the Luxembourg experience itself: longer matches involve greater accumulation of general and employer-specific human capital, and a completed multi-year contract is a stronger signal to Belgian employers than a brief commuting episode that may read as a failed match. Both channels point to frictional reallocation rather than structural depreciation.

4.3 Effect on working conditions and unemployment benefits

Working conditions. Conditional on being employed after return, dynamic DML estimates of the effect of CBW on log daily wages are negative for one-year spells (-0.08 log points), close to zero for two-year spells (-0.03), and positive for three-year spells (+0.07, Figure 2d). None is statistically distinguishable from zero, and confidence intervals widen sharply with spell length, reflecting the small employed-returnee samples (203, 74, and 29 observations respectively).

Despite the imprecision, the data rule out large positive wage premia at short spell lengths: the upper bound of the 95% CI is +0.03 log points for one-year spells and +0.10 for two-year spells. The three-year estimate is too imprecise to discriminate, with an upper bound of +0.24 log points. Effects on part-time work are also indistinguishable from zero (Figure 2e). Tables and longer-horizon plots are in the Appendix.

The absence of a large wage premium at short spell lengths is consistent with the residential return-migration literature. Abramitzky et al. (2019), in linked Norwegian-US census data, find no premium for one-to-three-year return migrants once premigration occupation is controlled for, and El-Mallakh and Wahba (2021) document substantial heterogeneity in returnees' occupational trajectories, with both upward and downward shifts depending on origin context. Premia in the residential literature accumulate over longer horizons: Wahba (2015) estimates a peak at six years of 27 percent for Egyptian returnees, and Reinhold and Thom (2013) finds 2.2 percent per year of US experience for Mexican returnees. The wage-premium mechanisms in this literature, foreign skill acquisition that compounds with duration (Dustmann et al., 2011) and occupation-specific experience matched to home-country jobs (Reinhold and Thom, 2013), are not detectable at the durations we observe.

Unemployment benefits. Figures 2f to 2h report the effects on UI outcomes one year after return. Among non-employed individuals, returnees are 8 pp more likely than comparable stayers to receive unemployment benefits after one year of CBW, 12 pp after two years, and 21 pp after three years. The duration of benefit receipt among recipients (in log days) shows no significant effect of CBW experience. Daily benefit levels in logs are about 0.11 higher for one- and two-year CBW spells (significant at the 10 percent level) and imprecise at three-year CBW spells due to the smaller sample of returnees on benefits.

The pattern, higher receipt rates and higher daily benefit levels, is consistent with the Article 62(3) wage-portability rule, under which Belgian benefits are computed on the Luxembourg gross wage that exceeds comparable Belgian wages. The fiscal incidence of these claims falls on the residence country: Luxembourg reimburses the first three months of benefits, extendable to five months when the worker has been employed in Luxembourg for at least one year, and Belgium bears the cost thereafter. Our estimates therefore imply substantial fiscal exposure on the residence-country UI system under the current rule. The estimates are consistent with the labor-market-state decomposition above: longer Luxembourg spells likely satisfy the residence-country eligibility thresholds that shorter or fragmented prior histories may not have met, moving returnees who would otherwise have remained inactive into the formal UI register. The opposite direction is documented for residential migrants in destination countries, who tend to under-claim UI relative to natives (Strockmeijer et al., 2020).

4.4 Validation and sensitivity

We assess the results through a battery of diagnostics reported in Appendix D. Inverse-probability reweighting sharply improves covariate balance: the pooled mean absolute standardized mean difference (SMD) across 495 covariate-node pairs falls from 0.116 to 0.048, the median falls from 0.049 to 0.032, and the 95th percentile falls from 0.466 to 0.145, so 95 percent of all covariate-node pairs are below the standard 0.15 threshold after reweighting. Residual imbalances concentrate on spatial variables at the CBW-entry node, in particular peak commuting time to Luxembourg ($|\text{SMD}| = 0.41$), the local unemployment rate (0.24), and public-transport connectivity (0.18 to 0.20), reflecting strong geographic sorting near the border that is difficult to eliminate through reweighting alone. Common support is adequate along the treatment path. Along the control path, the joint generalized propensity score is more separated, which means the doubly robust estimator relies more on the outcome model in that part of the support; cross-fitting disciplines the outcome model's out-of-sample predictions and prevents overfitting to thin regions (Chernozhukov et al., 2018). Sensitivity to propensity-score trimming is small: more aggressive trimming delivers minor balance gains at substantial sample loss, including the loss of up to 48 treated units, and is therefore not adopted.

E-values (VanderWeele and Ding, 2017) report the minimum strength of association on the risk-ratio scale that an unmeasured confounder would need to have with both treatment and outcome, beyond the observed covariates, to fully explain away the estimated effects. For the one-year employment penalty, the CI-bound E-value is 1.75; for the unemployment effect after one year of CBW, 2.76; for the unemployment effect after two years of CBW, 2.61; for the UI receipt effect, 1.56 to 1.89 depending on spell length. These values can be benchmarked against the strongest observed predictor of CBW entry, commuting time to Luxembourg, which moves predicted entry probabilities by a factor of about 2.8 across its support. To overturn our headline conclusions, residual confounding would need to (i) correlate with treatment at least as strongly as the most informative observable in our data, (ii) correlate with the outcome by a comparable margin, and (iii) be orthogonal to the 213 conditioning variables, which span ten years of quarterly labor-market history, household structure, and geographic controls. A confounder satisfying all three conditions seems implausible in this setting. Wage and part-time effects are imprecisely estimated regardless of confounding, with CI-bound E-values equal to one. Replicating the analysis with extreme gradient boosting yields estimates aligned with the random-forest baseline, indicating the findings do not hinge on the choice of learner.

5 Conclusion

The labor-market consequences of cross-border work for cross-border workers themselves have been largely undocumented, in part because outbound CBW spells are not systematically recorded in origin-country registries. Using linked Belgian administrative data that identify cross-border spells through compulsory sickness-insurance registries, and a dynamic DML design conditioning on ten years of quarterly labor-market history, this paper provides the first causal estimates of post-return labor-market outcomes for cross-border workers. To our knowledge, this is also the first evidence in any return-mobility or welfare-receipt literature on returnees' uptake of origin-country welfare benefits.

Three findings stand out. First, returnees face a short-run employment penalty that attenuates with both CBW duration and time since return. Second, returnees are more likely to receive Belgian unemployment benefits than comparable stayers, with higher daily benefit levels among recipients. Third, we find no evidence of a sizeable post-return wage premium in the short run, consistent with the residential return-migration literature. Two further selection patterns emerge from the analysis. Commuting time is a first-order determinant of entry, persistence, and return, suggesting that transport investments and hybrid-work arrangements would shift both the incidence and the stability of cross-border employment. Cross-border employment also operates as a mobility channel at both ends of the skill distribution, providing an employment pathway for workers with weaker domestic prospects alongside the high-skill margin already documented in the migration literature.

Two caveats bound the study. Sequential CIA is untestable; the E-value bounds quantify the confounding strength required to overturn the main results, but cannot rule it out. Sample sizes shrink at longer spells and for wage outcomes, so those estimates should be read as exploratory.

Intra-EU mobility is increasingly central to EU labor-market integration: 2.3 million EU residents were employed outside their country of residence in 2024 (Eurostat, 2024). The post-return frictions we document are likely to extend across this population. Under the current EU social-security coordination framework, unemployed cross-border workers claim benefits from their country of residence, with only the first few months reimbursed by the country of work. The provisional EU agreement of April 2026 (Council of the European Union, 2026) shifts competence to the country of last employment, aligning the financing of benefits with the location in which contributions were paid, thereby addressing only the fiscal channel we document in this study. The worker-side cost, however, remains. The employment penalty we identify calls for a separate response, such as reintegration policies, so that the costs of free movement within the EU do not fall on the workers themselves.

References

Abramitzky, R., Boustan, L., and Eriksson, K. (2019). To the new world and back again: Return migrants in the age of mass migration. *ILR Review*, 72(2):300–322.

Agersnap, O., Jensen, A., and Kleven, H. (2020). The welfare magnet hypothesis: Evidence from an immigrant welfare scheme in Denmark. *American Economic Review: Insights*, 2(4):527–542.

Ahrens, A., FitzGerald, J., and Lyons, S. (2020). Commuting across the Irish border. *Economics Letters*, 190:109060.

Akbani, R., Kwek, S., and Japkowicz, N. (2004). Applying support vector machines to imbalanced datasets. In *European conference on machine learning*, pages 39–50. Springer.

Albanese, A., Cockx, B., and Dejemeppe, M. (2024). Long-term effects of hiring subsidies for low-educated unemployed youths. *Journal of Public Economics*, 235:105137.

Albanese, A., Nieto, A., and Tatsiramos, K. (2022). Job location decisions and the effect of children on the employment gender gap. IZA Discussion Paper 15353, Institute of Labor Economics (IZA).

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

Beerli, A., Ruffner, J., Siegenthaler, M., and Peri, G. (2021). The abolition of immigration restrictions and the performance of firms and workers: Evidence from Switzerland. *American Economic Review*, 111(3):976–1012.

Bia, M., Huber, M., and Lafférs, L. (2024). Double machine learning for sample selection models. *Journal of Business & Economic Statistics*, 42(3):958–969.

Bijwaard, G. E., Schluter, C., and Wahba, J. (2014). The impact of labor market dynamics on the return migration of immigrants. *Review of Economics and Statistics*, 96(3):483–494.

Bodory, H., Huber, M., and Lafférs, L. (2022). Evaluating (weighted) dynamic treatment effects by double machine learning. *The Econometrics Journal*, 25(3):628–648.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

Broersma, L., Edzes, A., and Van Dijk, J. (2022). Commuting between border regions in The Netherlands, Germany and Belgium: an explanatory model. *Journal of borderlands studies*, 37(3):551–573.

Bütikofer, A., Løken, K. V., and Willén, A. (2024). Building bridges and widening gaps. *Review of Economics and Statistics*, 106(3):681–697.

Chabé-Ferret, B., Machado, J., and Wahba, J. (2018). Remigration intentions and migrants' behavior. *Regional Science and Urban Economics*, 68:56–72.

Chabé-Ferret, S. (2015). Analysis of the bias of matching and difference-in-difference under alternative earnings and selection processes. *Journal of Econometrics*, 185(1):110–123.

Chen, C., Liaw, A., Breiman, L., et al. (2004). Using random forest to learn imbalanced data. *University of California, Berkeley*, 110(1-12):24.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.

Clark, W. A. V., Huang, Y., and Withers, S. (2003). Does commuting distance matter?: Commuting tolerance and residential change. *Regional Science and Urban Economics*, 33(2):199–221.

Co, C. Y., Gang, I. N., and Yun, M.-S. (2000). Returns to returning. *Journal of Population Economics*, 13(1):57–79.

Council of the European Union (2026). Council and parliament strike provisional deal on social security coordination. Press release, 22 April 2026.

Daw, J. R. and Hatfield, L. A. (2018). Matching and regression to the mean in difference-in-differences analysis. *Health Services Research*, 53(6):4138–4156.

De Giorgi, G. and Pellizzari, M. (2009). Welfare migration in Europe. *Labour Economics*, 16(4):353–363.

Dicarlo, E. (2022). How do firms adjust to negative labor supply shocks? Evidence from migration outflows. IZA Discussion Paper 14994, IZA Institute of Labor Economics.

Dodini, S., Løken, K. V., and Willén, A. (2024). Labor market competition and its effect on firms and local communities. Working paper, Norwegian School of Economics.

Dodini, S., Lundborg, P., Løken, K. V., and Willén, A. (2025). The fatal consequences of brain drain. Working paper.

Dustmann, C. (1997). Return migration, uncertainty and precautionary savings. *Journal of Development Economics*, 52(2):295–316.

Dustmann, C. (2003). Return migration, wage differentials, and the optimal migration duration. *European Economic Review*, 47(2):353–369.

Dustmann, C., Fadlon, I., and Weiss, Y. (2011). Return migration, human capital accumulation and the brain drain. *Journal of Development Economics*, 95(1):58–67.

Dustmann, C., Schönberg, U., and Stuhler, J. (2017). Labor supply shocks, native wages, and the adjustment of local employment. *Quarterly Journal of Economics*, 132(1):435–483.

El-Mallakh, N. and Wahba, J. (2021). Upward or downward: Occupational mobility and return migration. *World Development*, 137:105203.

Eurostat (2024). Employment and commuting by sex, age and NUTS 2 regions. Eurostat, Code: Lfst_r_lfe2ecomm. Retrieved on the 4th of February 2026.

Godefroy, S., Klein, S., Delloye, J., Bredel, C., Lang-Karevski, V., Jacquot, M., and Schiebel, J. (2021). *Exploitation harmonisée des enquêtes de déplacements sur le périmètre MMUST*. AGAPE-Agence d'urbanisme et de développement durable Lorraine Nord.

Gottholmseder, G. and Theurl, E. (2007). Determinants of cross-border commuting: Do cross-border commuters within the household matter? *Journal of Borderlands Studies*, 22(2):97–112.

Hafner, F. and Hedtrich, C. (2024). The equilibrium effects of workers' outside employment options: Evidence from a labour market integration. Working paper.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The elements of statistical learning*. Springer New York.

Huber, P. (2014). Are commuters in the EU better educated than non-commuters but worse than migrants? *Urban Studies*, 51(3):509–525.

Huber, P. and Nowotny, K. (2013). Moving across borders: Who is willing to migrate or commute? *Regional Studies*, 47(9):1462–1481.

Illing, H. (2025). Crossing borders: Labor market effects of European integration. *Journal of Human Resources*.

Lechner, M. (2010). The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics*, 4(3):165–224.

Mathä, T. Y. and Wintr, L. (2009). Commuting flows across bordering regions: A note. *Applied Economics Letters*, 16(7):735–738.

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

Nowotny, K. (2014). Cross-border commuting and migration intentions: The roles of risk aversion and time preference. *Contemporary Economics*, 8(2):1–123.

Piracha, M. and Vadean, F. (2010). Return migration and occupational choice: Evidence from Albania. *World Development*, 38(8):1141–1155.

Reinhold, S. and Thom, K. (2013). Migration experience and earnings in the Mexican labor market. *Journal of Human Resources*, 48(3):768–820.

Rouwendal, J. and Meijer, E. (2001). Preferences for housing, jobs, and commuting: A mixed logit analysis. *Journal of Regional Science*, 41(3):475–505.

Rouwendal, J. and Rietveld, P. (1994). Changes in commuting distances of Dutch households. *Urban Studies*, 31(9):1545–1557.

Saarela, J. and Finnås, F. (2009). Return migrant status and employment in Finland. *International Journal of Manpower*, 30(5):489–506.

STATEC (2025). Employment: percentage by economic activity and residence/nationality.

Strockmeijer, A., de Beer, P., and Dagevos, J. (2020). Explaining differences in unemployment benefit takeup between labour migrants and Dutch native workers. *International Social Security Review*, 73(2):75–95.

Suari-Andreu, E. and van Vliet, O. (2023). Intra-EU migration, public transfers and assimilation. *Economica*, 90(360):1229–1264.

Topa, G. (2011). Labor markets and referrals. In Benhabib, J., Bisin, A., and Jackson, M. O., editors, *Handbook of Social Economics*, volume 1B, pages 1193–1221. Elsevier.

Tsiopa, A., Gerber, P., and Caruso, G. (2024). Framing the cross-border commuting literature: A systematic review and bibliographic analysis. *Transport Reviews*, 44(4):889–911.

VanderWeele, T. J. and Ding, P. (2017). Sensitivity analysis in observational research: Introducing the e-value. *Annals of Internal Medicine*, 167(4):268–274.

Wahba, J. (2015). Selection, selection, selection: The impact of return migration. *Journal of Population Economics*, 28(2):535–563.

Xia, W., Turunen, J., Aradhya, S., and Saarela, J. (2025). Returns to returning with sibling comparisons: Register-based evidence from Finland. *Stockholm Research Reports in Demography*.

Table 1: Full list of explanatory variables

Demographics:	Age, Gender, Being dependent in household, Nationality, Migrant origin (parents/grandparents), Luxembourgish origin (parents/grandparents), Internal migrant, Years in Belgium, Education level
Household:	Household composition, Household size, Nr children by age group (0, 1–5, 6–17 years), Work intensity, Partner cross-border worker, Any household member cross-border worker, Other household earners
Labor-Market History:	*Work history (1–10 years):* Average quarterly remuneration conditional on working (log), Quarters as: employed, self-employed, unemployed (with/without job search exemption), inactive (with/without transfers, minimum integration income); Months as: cross-border worker (self, partner, household); share of quarters in: public sector, blue-collar, white-collar
	Previous job: Type (no job, employee, self-employed), Collar (blue, white, public servant), Part-time status, Firm size, Industry sector (NACE 2-digit), Last salary (log), No last salary indicator, Firm fixed effects, Individual fixed effects
	Current status: Private sector employee, Self-employed, Unemployed (with/without job search exemption), Inactive (with/without transfers, minimum integration income), Public Servant, Current quarterly remuneration (log), Current total revenue (log), Unemployment benefit days (log)
Geography:	Distance to Luxembourg City (car), Distance by public transport, Missing public transport indicator, Local unemployment rate

Figure 1: Selection into and out of cross-border employment: partial dependence plots

Panel A. Entry into CBW

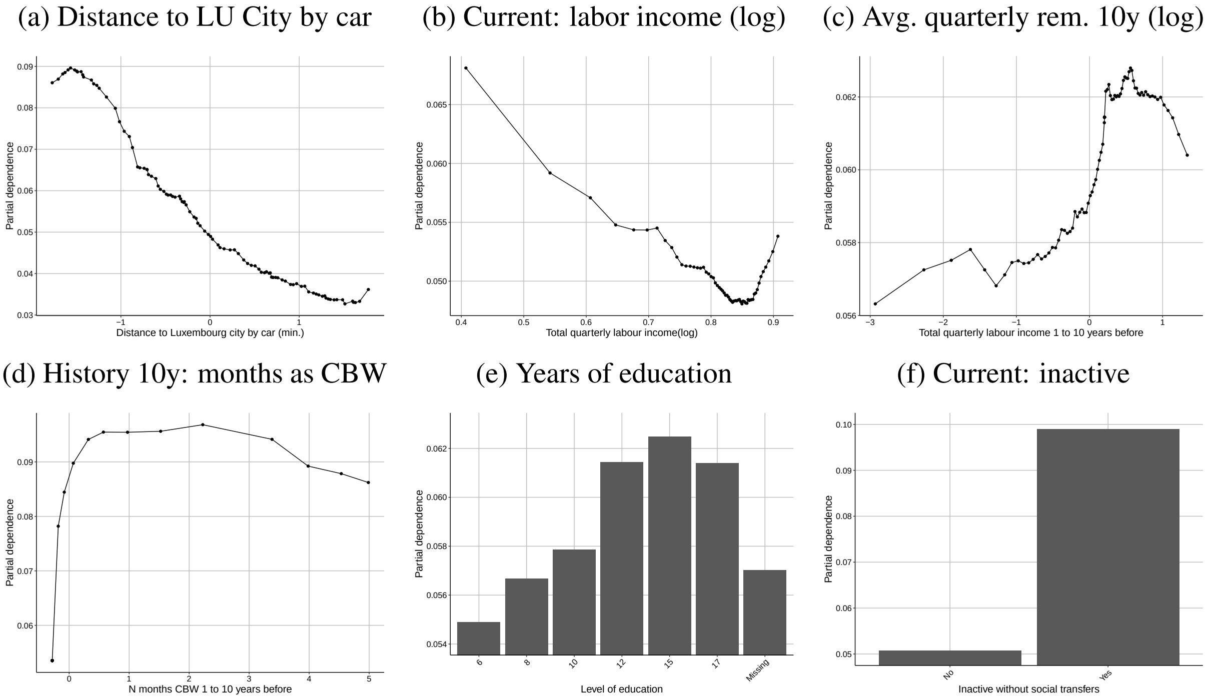

Panel B. Return from CBW

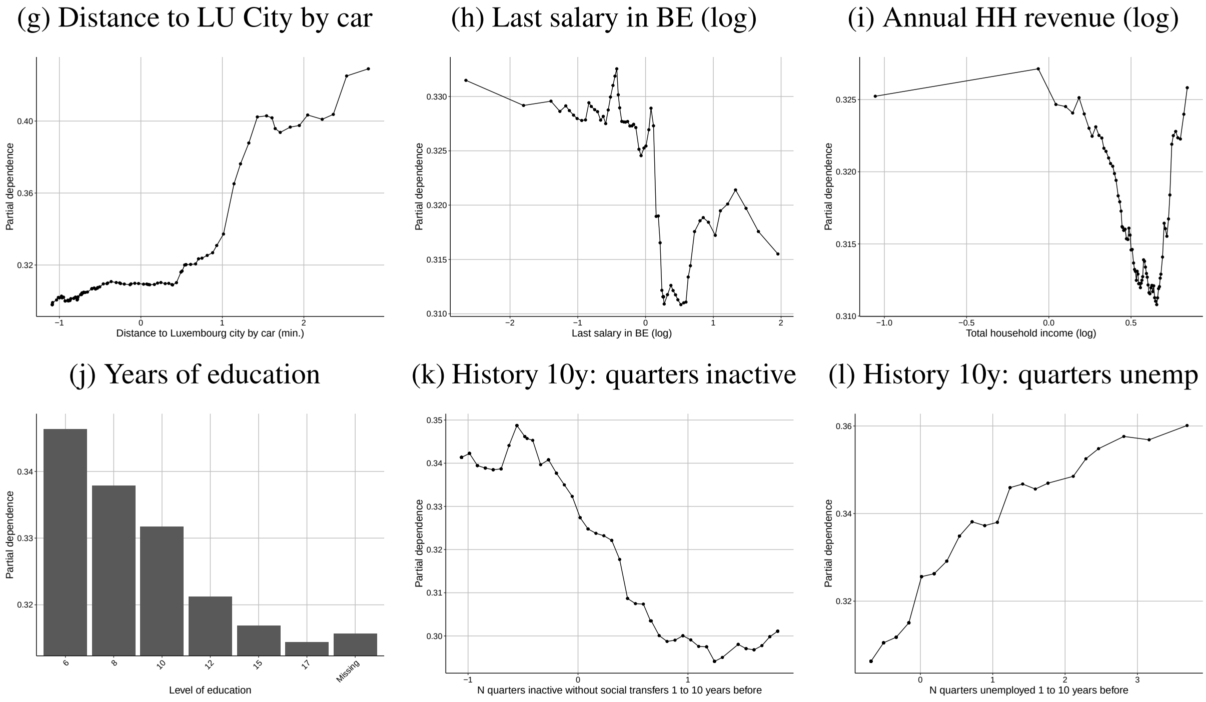

Note: Partial dependence plots for the strongest predictors of CBW entry (Panel A) and return within one year of entry (Panel B), estimated by random forest. Panel A: pooled 2013-2015 sample, $N = 144{,}600$. Panel B: pooled 2014-2016 sample of individuals in CBW one year earlier, $N = 3{,}708$. Each curve traces the model's predicted probability as the predictor on the horizontal axis varies, averaging over the joint empirical distribution of the remaining covariates. Continuous covariates standardized to mean zero and unit standard deviation. PDPs are descriptive net associations, not causal effects. Additional predictors and variable-importance rankings in Appendix B.

Figure 2: Effects One Year after Return, by CBW Spell Length

(a) Employment

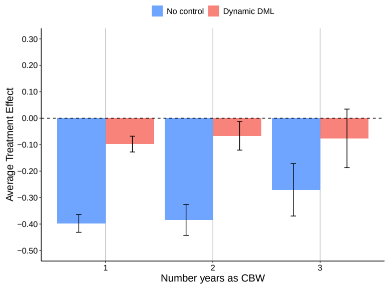

(b) Registered unemployment

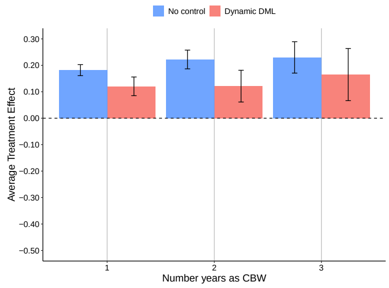

(c) Inactivity

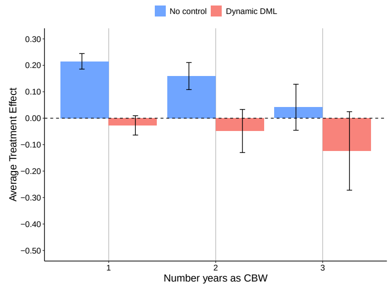

(d) Log Daily Wages (employed)

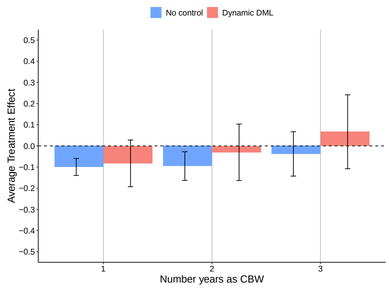

(e) Part-time Working Hours (employed)

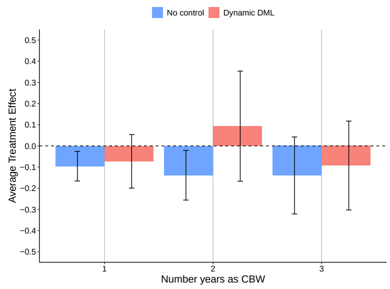

(f) Receiving UI (non-employed)

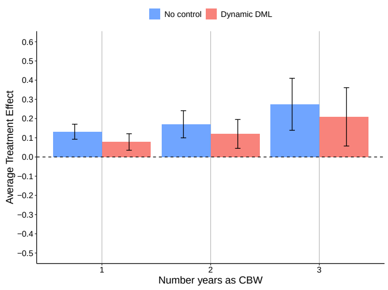

(g) Days Receiving UI (log - unemployed)

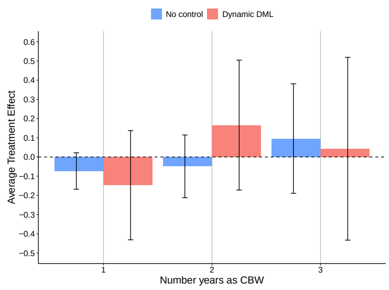

(h) Daily UI Amount in Euros (log - unemployed)

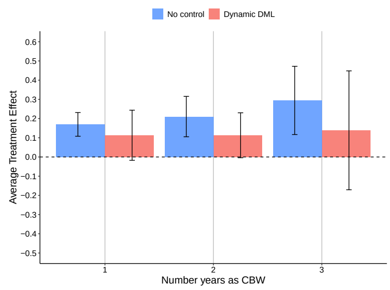

Note: Point estimates and 95% confidence intervals for the effect of CBW one year after return, by spell length. Top row: labor-market states (employment, unemployment, inactivity). Middle row: job quality conditional on employment (log daily wages, part-time status). Bottom row: UI outcomes (receipt conditional on non-employment, log days of receipt, log daily benefit level among recipients). Pooled 2013–2015 samples for one year of CBW; 2013–2014 for two years; 2013 for three years. Blue bars: raw differences. Red bars: dynamic DML with random forests. Standard errors clustered at the individual level. Stayer (returnee) sample sizes: top row: 136,646 (740), 88,734 (241), 43,428 (79); middle row: 81,178 (203), 53,826 (74), 26,731 (29); panel (f): 38,257 (498), 23,421 (155), 10,972 (42); panels (g) and (h): 9,273 (187), 5,601 (64), 2,598 (22).

Figure 3: Evolution of effects on labor-market states up to three years after return

(a) 1 year of CBW: Employment

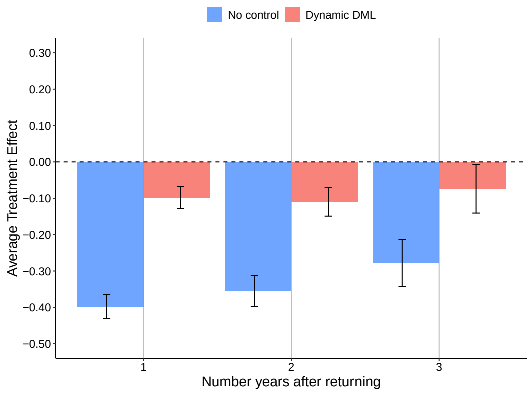

(b) 2 years of CBW: Employment

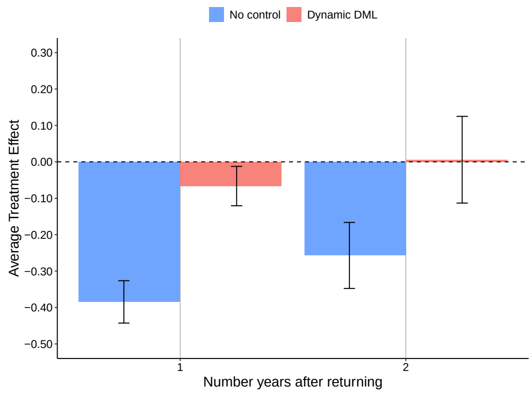

(c) 1 year of CBW: Registered Unemployment

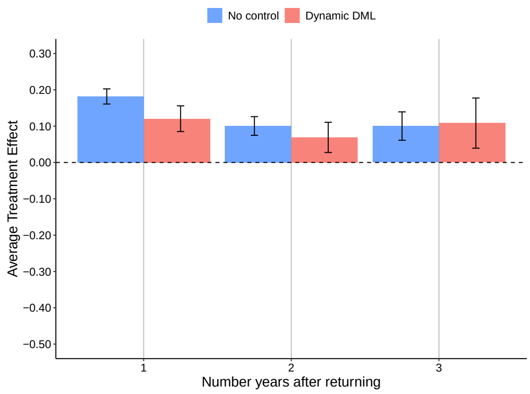

(d) 2 years of CBW: Registered Unemployment

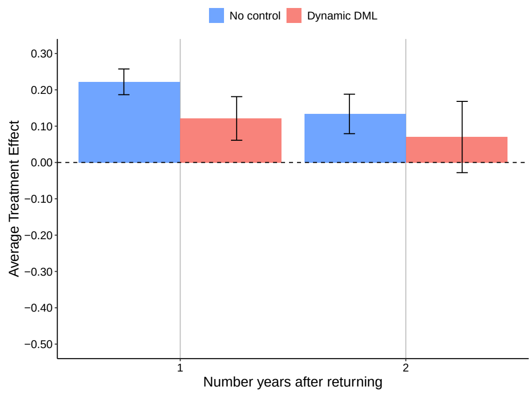

(e) 1 year of CBW: Inactivity

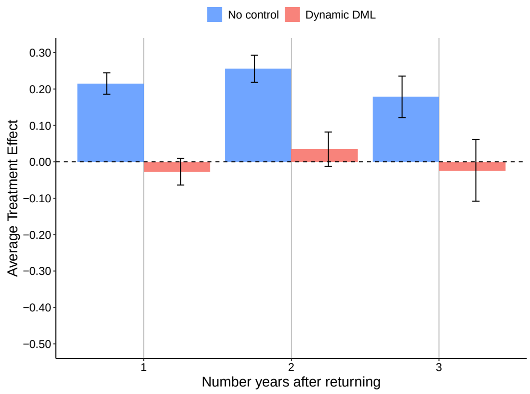

(f) 2 years of CBW: Inactivity

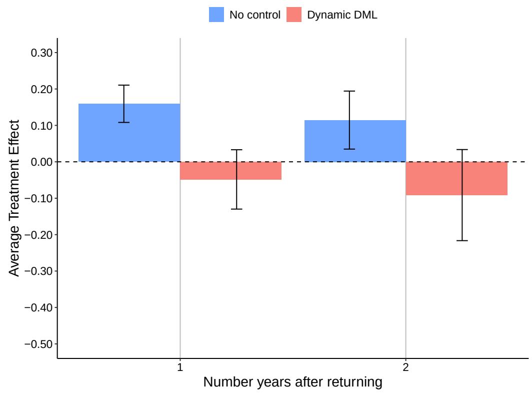

Note: Point estimates and 95% confidence intervals up to three years after return for one-year spells (left column), and up to two years for two-year spells (right column). Blue bars: raw differences. Red bars: dynamic DML with random forests. Standard errors clustered at the individual level. For one year of CBW: stayer (returnee) sizes are 136,646 (740) at one year after return, 88,734 (446) at two years, and 43,428 (184) at three years. For two years of CBW: 88,734 (241) and 43,428 (98).

Appendix

A Variable construction

Explanatory variables are grouped into four blocks: demographics, household characteristics, labor-market history and current status, and geography (for a literature review, see Tsiopa et al., 2024). Table A1 reports descriptive statistics on the analysis sample.

Demographics. Age, gender, educational attainment, nationality, migrant and Luxembourgish origin, internal migration within Belgium, and years of residence in Belgium. These capture flexibility, expected lifetime earnings, household responsibilities, and transaction costs through language skills, networks, and institutional familiarity (Huber and Nowotny, 2013; Gottholmseder and Theurl, 2007; Mathä and Wintr, 2009).

Household characteristics. Household composition (type, size, and children by age band), household work intensity, the presence of another earner, and indicators for CBW by the partner and any other household member. Household structure and childcare responsibilities are expected to influence commuting choices by raising time constraints, while CBW among household members may facilitate transitions through information sharing and network effects (Gottholmseder and Theurl, 2007; Albanese et al., 2022).

Labor-market history and current status. Using the BCSS nomenclature, we assign each individual a quarterly socio-economic position and aggregate over ten years to compute quarters spent in employment (private and public), self-employment, unemployment, and inactivity. We also measure the share of quarters by occupational category (blue-collar, white-collar, public servant), months in CBW (individual, partner, household), and average remuneration (log) conditional on employment. We include characteristics of the most recent Belgian job: employment type, occupational category, part-time status, firm size, industry (NACE 2-digit), last wage (log), a missing-wage indicator, and worker and firm fixed effects estimated on pre-2013 data (Bonhomme et al., 2019). Last-quarter status includes labor-market position, minimum-income receipt, unemployment-benefit days (log), current remuneration (log), and household income (log) (Nowotny, 2014).

Geographical factors. Commuting time to Luxembourg City by car and by public transport, and the local unemployment rate in the municipality of residence (Rouwendal and Meijer, 2001; Clark et al., 2003; Mathä and Wintr, 2009; Broersma et al., 2022).

In the causal analysis, we use the full specification with 213 parameters as potential confounders. For the predictive analysis, we adopt a more parsimonious specification with 159 parameters, removing or aggregating highly correlated variables.[1] This enhances the interpretability of partial dependence plots in the presence of multicollinearity, since they depict how the predicted probability of CBW varies with a given explanatory variable while keeping the remaining covariates fixed, which becomes difficult to interpret with highly correlated explanatory variables.

Table A1: Descriptive Statistics

	All	Time 1		For CBW in time 1, time 2	
		Stayers	CBW	Remain CBW	Returnees
Women	0.516	0.521	0.392	0.377	0.451
No previous CBW exp.	0.886	0.891	0.649	0.660	0.600
Employee	0.580	0.598	0.211	0.213	0.212
Self-employed	0.092	0.094	0.042	0.041	0.047
Unemployed	0.097	0.094	0.210	0.195	0.265
Inability	0.035	0.036	0.007	0.005	0.015
Minimum income	0.014	0.014	0.003	0.002	0.006
Other unempl.	0.016	0.016	0.014	0.014	0.015
Other inactivity	0.167	0.148	0.512	0.529	0.439
Daily salary (€)	89.7	91.2	68.4	67.1	76.7
Age	32.6	32.6	31.4	31.4	31.3
Years of education	12.0	12.0	12.3	12.5	11.9
CBW in next year	0.024	0.000	1.000	1.000	1.000
Returning from CBW in 2			0.194	0.000	1.000
Remain CBW in 2			0.782	1.000	0.000
N	147,102	140,805	3,795	2,968	740

Pooled 2013-2015 cohorts, design weights applied. "Time 1" refers to CBW status at the end of the first observation year. "For CBW in time 1, time 2" restricts to those in CBW at end of year 1, then splits by year-2 status. "Inability" denotes disability, "Other unempl." denotes unemployment without active job search.

[1] In an OLS regression on the transition to CBW with the explanatory variables, the mean VIF is reduced by a factor of 2.7 in the parsimonious specification, from 9.21 to 3.44.

B Predictive Analysis

We implement Random Forest classification (Breiman, 2001) to predict transitions into and out of CBW within a one-year horizon. We cluster individual identifiers to ensure that the same individual is not included in both the training and test samples, and reweight observations to balance the outcome distribution given the rare-event nature of CBW transitions (Chen et al., 2004; Akbani et al., 2004). The model achieves an out-of-bag area under the ROC curve (AUC) of 0.867 for entry (PR-AUC of 0.221 against a base-rate prevalence of 0.024, an eightfold improvement over chance) and 0.665 for exit. The lower performance for exit is unsurprising: our covariates are measured on the Belgian side, and we lack information on the Luxembourg job (contract type, occupation, sector, wages) that is likely relevant for return decisions.

B.1 Variable importance

Figure A1 reports permutation-based variable-importance scores for all predictors of entry (Panel a) and exit (Panel b). Importances reflect predictive contribution rather than causal relevance.

B.2 Partial dependence plots: entry

Figure A2 reports PDPs for predictors of entry that complement those displayed in the main text. The patterns are consistent with the discussion in Section 4.1 of the main paper.

Four additional patterns emerge. Public transport time to Luxembourg shows a similar negative gradient to car commuting time, confirming that mobility frictions operate through both modes. Local unemployment rates at the NUTS-3 level are positively associated with entry, consistent with a push effect of weaker domestic opportunities (Gottholmseder and Theurl, 2007). Entry propensities are higher at younger ages and among men, consistent with longer horizons for amortizing fixed mobility costs and gender differences in commuting constraints (Rouwendal and Rietveld, 1994; Clark et al., 2003; Albanese et al., 2022). The presence of children, particularly infants, is negatively associated with entry, consistent with higher commuting time costs (Rouwendal and Rietveld, 1994; Clark et al., 2003; Albanese et al., 2022). Individuals of Luxembourgish origin are

Figure A1: Variable importance scores

(a) Entry into CBW (b) Return from CBW

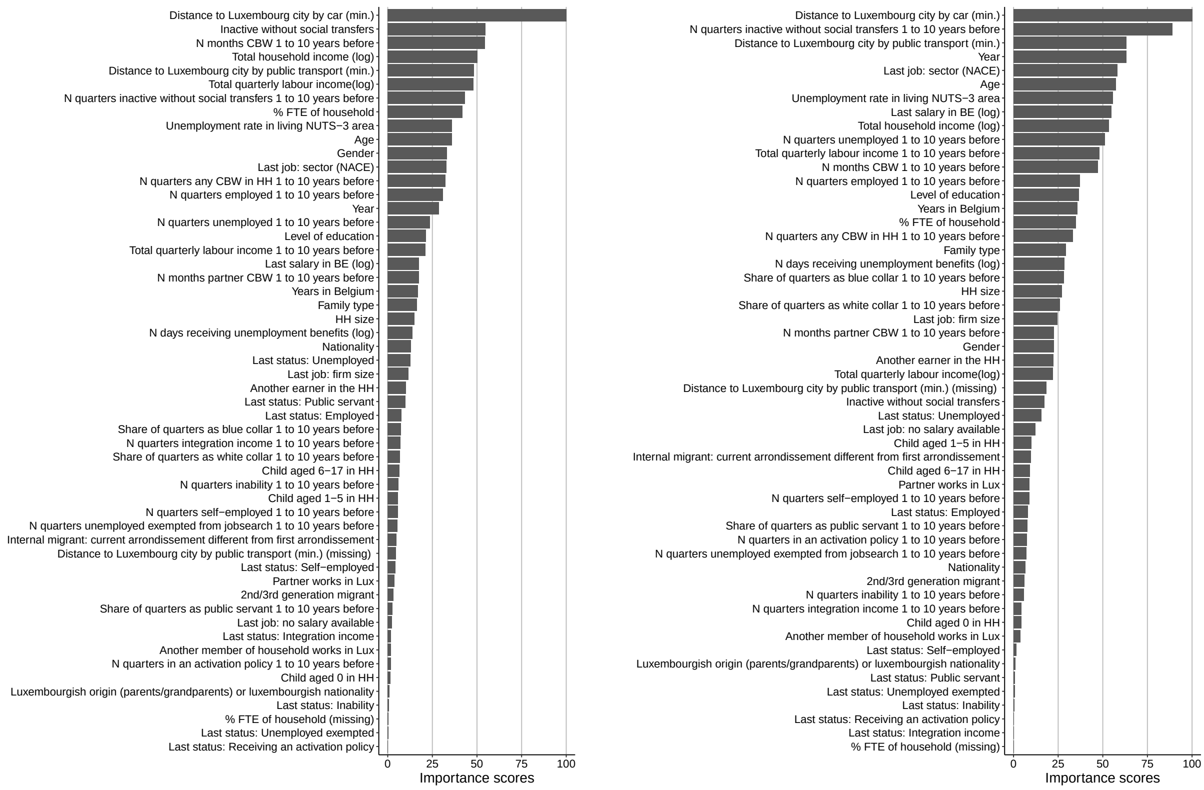

Higher values indicate a stronger contribution to predicting the outcome at a one-year horizon. Panel (a): pooled sample 2013-2015, $N = 144{,}600$. Panel (b): pooled sample 2014-2016 of individuals in CBW one year earlier, $N = 3{,}708$. Clustering by individual identifiers.

more likely to enter CBW, suggesting that language and institutional familiarity reduce transaction costs (Huber, 2014). Finally, household income shows a negative gradient consistent with a financial-pressure margin, reinforcing the necessity-driven channel discussed in the main text. The sectoral profile (last NACE division of the Belgian job) reinforces the dual-margin reading: predicted entry is concentrated in construction and market-oriented activities, remains elevated through transport and market services, and is lowest in public administration, education, and health and social work, consistent with the negative gradient on public-sector employment histories. Predicted entry is also elevated among workers from extraterritorial organisations, reflecting Luxembourg's concentration of European and other international institutions.

B.3 Partial dependence plots: exit

Figure A3 reports PDPs for additional predictors of return.

Figure A2: Partial dependence plots: additional predictors of entry into CBW

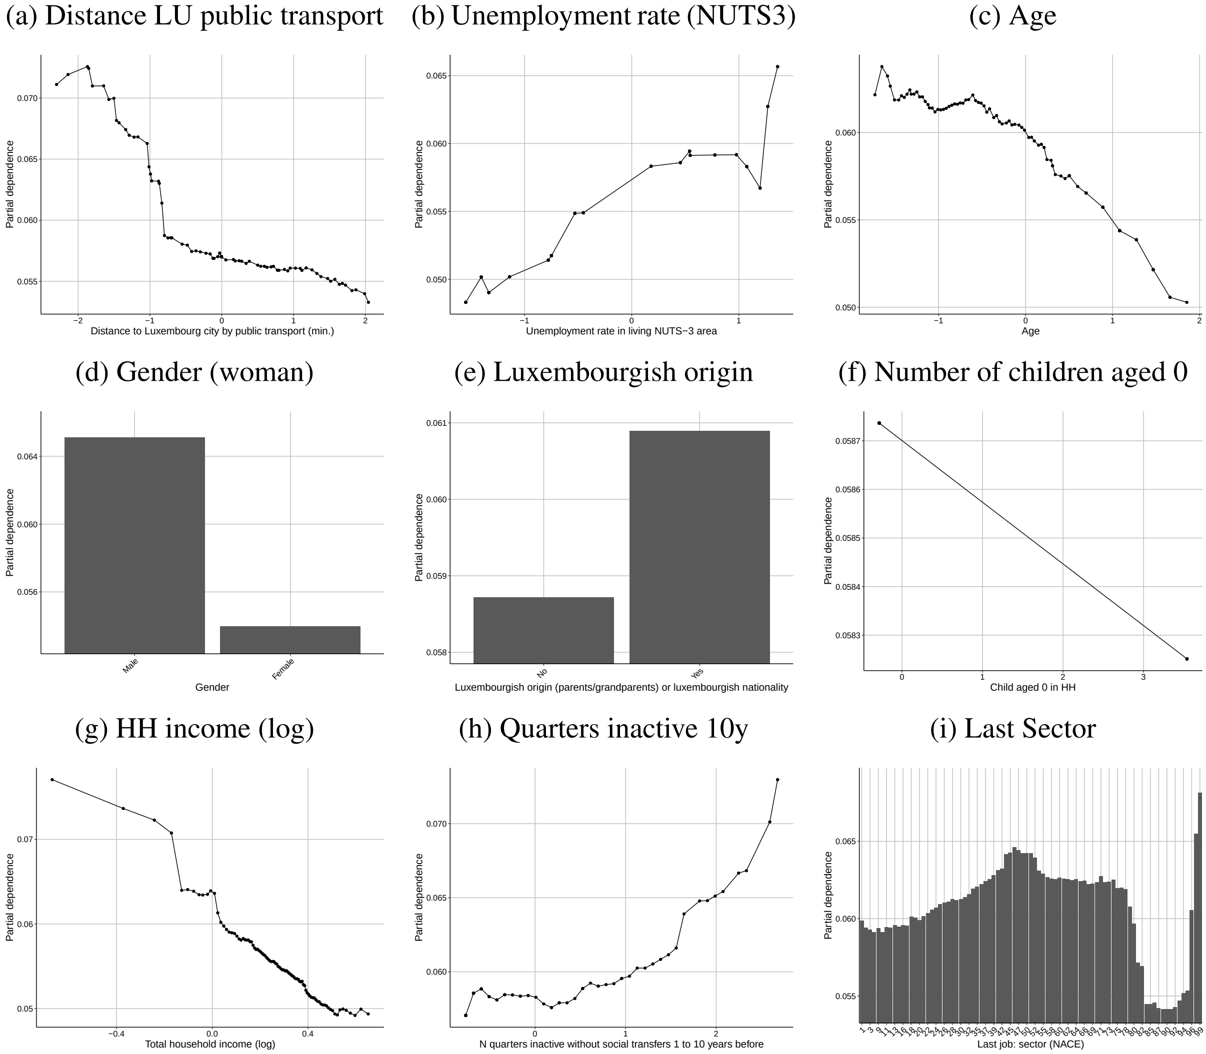

Note: Partial dependence plots for additional predictors of entry into CBW. Continuous covariates standardized to mean zero and unit standard deviation. PDPs are descriptive net associations, not causal effects.

The age profile declines into mid-career and then flattens, suggesting that younger individuals hold less stable Luxembourg positions. Women are more likely to return, consistent with household constraints that limit the ability to remain in CBW. Family constraints are most visible for infants: the presence of children aged zero predicts higher return, consistent with tighter time constraints and higher value of proximity to home. Household structure and cross-border networks contribute to persistence, consistent with information advantages that reduce search and adaptation costs. Prior CBW exposure of the worker, partner, or other household members is negatively associated with return, mirroring the positive association with entry: the same network and learning channels that pull workers into CBW also stabilize them within it. Finally, the sectoral profile is asymmetric to

Figure A3: Partial dependence plots: additional predictors of return from CBW

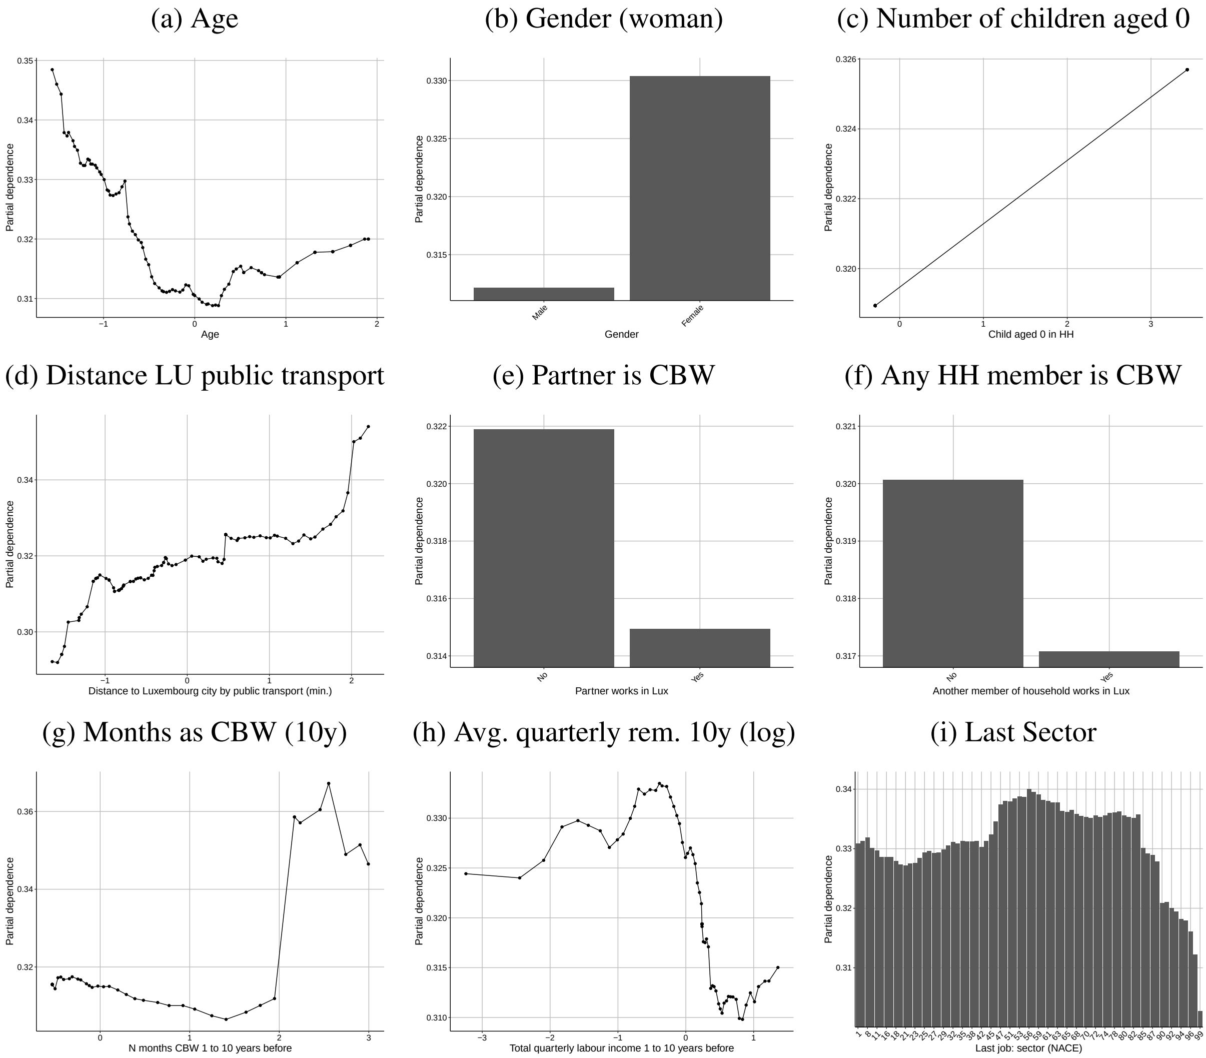

Note: Partial dependence plots for additional predictors of return from CBW. Continuous covariates standardized to mean zero and unit standard deviation. PDPs are descriptive net associations, not causal effects.

the entry margin: predicted return is highest among workers whose last Belgian job was in market-oriented activities and public administration, and lowest in cultural and recreational services and extraterritorial organisations, consistent with greater persistence among entrants from sectors with weaker domestic outside options and with high match stability in international-institution employment.

C Dynamic Double Machine Learning Framework

This appendix collects the formal definitions and implementation details underlying the dynamic DML estimator described in Section 3 of the main paper.

C.1 Treatments, sequences, and the ATE

We consider the time-varying treatment D_t, which can take three values d_t at time t, defined as the status at the end of the year: 0 if the individual lives in Belgium while not working in Luxembourg, 1 if the individual lives in Belgium while working in Luxembourg, and 2 if the individual no longer lives in Belgium or is deceased.[2]

The key object of interest is the average treatment effect (ATE) of a sequence of treatments $\underline{D}_T = [D_1, \ldots, D_T]$. This is retrieved by comparing the potential outcome $Y_T(\underline{d}_T)$ under a specific treatment sequence $\underline{d}_T = [d_1, \ldots, d_T]$ to the potential outcome $Y_T(\underline{d}_T^*)$ under a control sequence $\underline{d}_T^*$. In the potential outcome framework, the ATE of two distinct treatment sequences is defined as

$$ATE_T = \Delta(\underline{d}_T, \underline{d}_T^*) = E[Y_T(\underline{d}_T) - Y_T(\underline{d}_T^*)]. \tag{A1}$$

In our setting, the control sequence corresponds to individuals who never leave Belgium, so $\underline{d}_T^* = [0, 0, \ldots, 0]$. To assess how the effects vary with spell length and time since return, we evaluate sequences $[0,1,0]$, $[0,1,0,0]$, $[0,1,0,0,0]$ (one year of CBW; outcome measured one, two, three years after return); $[0,1,1,0]$, $[0,1,1,0,0]$ (two years of CBW; outcome one and two years after); and $[0,1,1,1,0]$ (three years of CBW; one year after return).[3]

C.2 Sequential propensity scores and dynamic selection

The dynamic DML estimator models selection at each transition node along the sequence, with the number of nodes equal to the length of the treatment sequence T. At the first node, we model selection into cross-border employment ($D_1 = 1$ vs. $D_1 \in \{0, 2\}$),

[2] The analysis is conducted at the end-of-year frequency to avoid an explosion in the number of possible treatment paths. Individuals who change CBW status within a calendar year, whether entering, exiting, or briefly interrupting a spell, but do not hold that status at year-end are classified according to their December status, which leads to a modest undercount of short or interrupted spells. Short spells transitioning to CBW within the year but not maintained at year-end account for around 15 percent of all transitions to CBW. Contamination of the control group is minimal at 0.4 percent.

[3] We do not report treatment effects for sequences involving exit from the Belgian administrative registers (out-migration or death), as outcomes are not observed. We model exit as a separate treatment value at each period when estimating propensity scores, so selection into exit is accounted for in the reweighting across in-sample sequences. Cross-border workers are somewhat more likely to exit subsequently (87/3,795 = 2.3 percent against 1,713/140,805 = 1.2 percent per period), but these shares are small.

estimating a propensity score conditional on the baseline covariate vector X_0. At each subsequent node $t = 2, \ldots, T$, a further propensity score is estimated conditional on the updated covariate vector X_{t-1}, which is progressively enriched with labor-market outcomes observed during the preceding period. For the treatment sequences involving CBW, these subsequent nodes model continuation in CBW versus exit ($D_t = 1$ vs. $D_t \in \{0,2\}$) during the spell, and the return decision ($D_t = 0$ vs. $D_t \in \{1,2\}$) at the end of the spell, with additional nodes appended for each post-return year over which the outcome is measured. The same sequential logic applies symmetrically to the control sequence: at each node $t = 1, \ldots, T$, conditional on having remained in Belgium, a propensity score is estimated to account for selection into continued domestic employment versus exit ($D_t = 0$ vs. $D_t \in \{1,2\}$). For each sequence of interest, including the control sequence, the mean potential outcome is estimated from the observations that actually followed that path, reweighted by the inverse of the product of period-specific propensity scores along the path and augmented by nested outcome regressions. Figure A4 graphically represents the dynamic causal paths for the baseline two-period case.

C.3 Weighted ATE on first-period entrants

The estimates we identify are average treatment effects on the subpopulation of first-period CBW entrants, $S = \mathbb{1}\{D_1 = 1\}$, rather than on the full population of potential movers. This is the policy-relevant estimand, as it describes the effect of cross-border experience among workers who actually take up CBW, and it is also the population on which common support is strongest. Estimation follows the weighted-score representation of Bodory et al. (2022, Section 4): the efficient score $\psi^{\underline{d}_T}$ is multiplied by $g(X_0)/\Pr(S = 1)$, where $g(X_0) = \Pr(S = 1 \mid X_0)$ is the conditional probability of belonging to the subgroup. In our setting, S is a deterministic function of D_1, so $g(X_0) = p^{d_1=1}(X_0)$ coincides with the first-period propensity score we already estimate at Node 1, and no additional nuisance parameter is introduced. Mechanically, the weighted score leaves treated-path observations reweighted by inverse propensity-score products as usual, while control-path observations are reweighted to match the covariate distribution of first-period entrants. Under Assumption 4.1 of Bodory et al. (2022), which reduces in our case to the rate conditions imposed

Figure A4: Dynamic selection for labor-market states with 2013-2015 samples

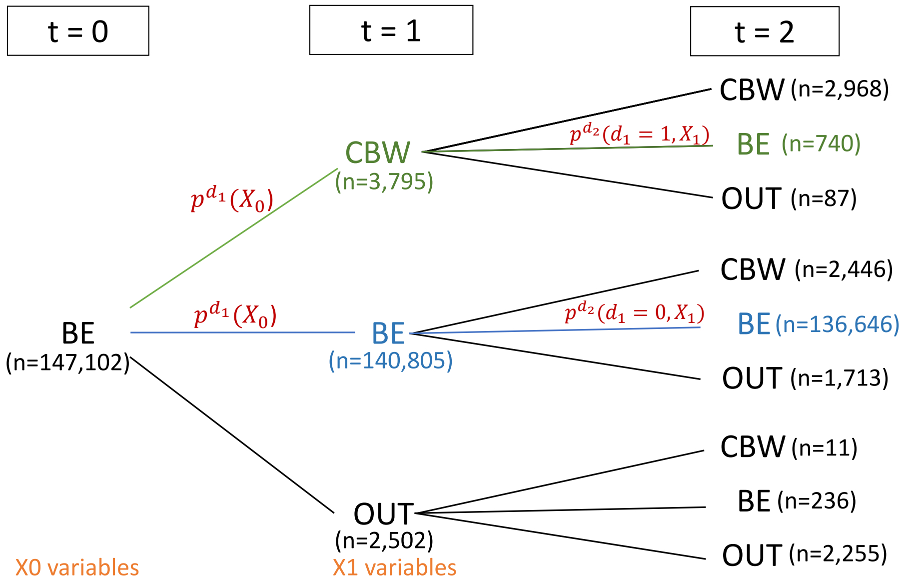

The figure illustrates the empirical setup for the baseline case of one year of CBW experience ($\underline{d}_T = [0,1,0]$), used to analyze the effect on labor-market states (employment, unemployment, or inactivity). At the end of each year t, individuals are observed in Belgium (BE, d_t=0), employed in Luxembourg (CBW, d_t=1), or no longer in the Belgian registers due to migration or death (OUT, d_t=2). Propensity scores $p^{d_1}(X_0)$, $p^{d_2}(d_1 = 1, X_1)$, and $p^{d_2}(d_1 = 0, X_1)$ are estimated at the indicated nodes. The variables X_0 and X_1 denote the observed covariates at the corresponding time points. For longer treatment sequences, the same structure is extended by adding further nodes, each associated with an additional propensity score estimated on the correspondingly updated covariate vector. Sample sizes for each group are indicated by n.

on the period-specific propensity-score and outcome estimators, Theorem 4.1 of that paper delivers $\sqrt{n}$-consistency and asymptotic normality for the resulting weighted ATE.

C.4 Sample selection correction

When the outcome is observed only for a subset of individuals, such as wages or unemployment-benefit amounts, an additional selection challenge arises. Since CBW itself affects subsequent employment, conditioning on post-return employment is conditioning on a post-treatment variable, which can introduce sample selection bias. We address this within the existing dynamic DML framework by exploiting the equivalence, shown in Bia et al. (2024), between modeling sample selection as a separate propensity score and incorporating it directly into the dynamic treatment sequence. As Proposition 4 of that paper demonstrates, the efficient score function under sequential conditional independence for sample selection has the same structure as the score for dynamic treatment effects. We expand the state space at the outcome period to include employment (or unemployment) status, so that the sequential propensity scores estimated at each node jointly correct for selection into the treatment path and selection into employment or unemployment. The credibility of this approach rests on the same sequential CIA invoked for treatment assignment: conditional on the full covariate history $\underline{X}_{t-1}$ and the treatment sequence $\underline{d}_T$, potential outcomes are independent of the selection indicator. Given the richness of the conditioning set, including ten years of quarterly employment, earnings, and benefit histories, this is plausible because the main drivers of employment selection are themselves observable in the trajectories we control for.

Figure A5 illustrates the dynamic paths for wage outcomes.

Figure A5: Dynamic selection for working conditions, 2013-2015 samples

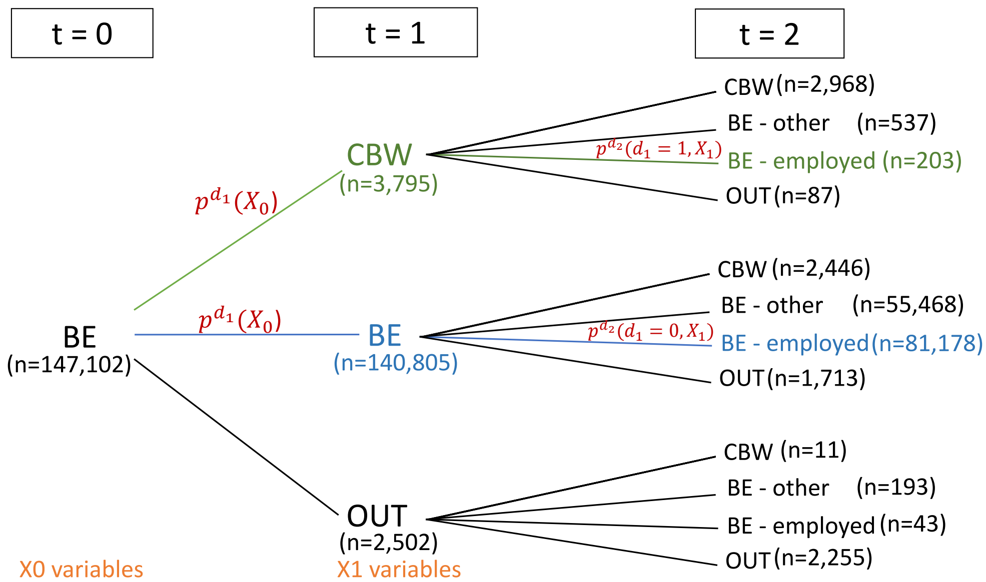

Empirical setup for the effect of one year of CBW on log daily wages and part-time employment. At t=2, employment status enters as an additional node: BE individuals split into employed (BE, employed) and other (BE, other) to correct for sample selection into the wage outcome.

C.5 Neyman-orthogonal score

The dynamic DML estimator is doubly robust through the Neyman-orthogonality condition, achieved by combining the outcome and treatment models. The potential outcome $E[Y_T(\underline{d}_T)]$ is estimated by the score function $\psi^{\underline{d}_T}$, which we extend to sequences of up to four periods, building on the general multi-period formula of Bodory et al. (2022, footnote

5).[4]

$$\psi^{\underline{d}_T} = \frac{\prod_{t=1}^{T} I\{D_t = d_t\}\left[Y_T - \mu^{Y_T}(\underline{d}_T, \underline{X}_{T-1})\right]}{\prod_{t=1}^{T} p^{d_t}(\underline{d}_{t-1}, \underline{X}_{t-1})} + \frac{\prod_{t=1}^{T-1} I\{D_t = d_t\}\left[\mu^{Y_T}(\underline{d}_T, \underline{X}_{T-1}) - \nu^{Y_T}(\underline{d}_T, \underline{X}_{T-2})\right]}{\prod_{t=1}^{T-1} p^{d_t}(\underline{d}_{t-1}, \underline{X}_{t-1})} + \cdots + \frac{I\{D_1 = d_1\}\left[\xi^{Y_T}(\underline{d}_T, \underline{X}_1) - \gamma^{Y_T}(\underline{d}_T, X_0)\right]}{p^{d_1}(X_0)} + \gamma^{Y_T}(\underline{d}_T, X_0) \tag{A2}$$

where $I\{D_t = d_t\}$ is an indicator for units in treatment status d_t at time t. The propensity score of treatment d_t, $p^{d_t}(\underline{d}_{t-1}, \underline{X}_{t-1})$, is conditional on the sequence of covariates until the previous period $\underline{X}_{t-1}$ and the realized treatment sequence $\underline{d}_{t-1}$. The expected outcome at T, conditional on the treatment sequence $\underline{d}_T$ and the covariate sequence $\underline{X}_{T-1}$, is $\mu^{Y_T}(\underline{d}_T, \underline{X}_{T-1})$, with nested conditional means represented by $\nu^{Y_T}(\underline{d}_T, \underline{X}_{T-2})$ down to $\xi^{Y_T}(\underline{d}_T, \underline{X}_1)$ and $\gamma^{Y_T}(\underline{d}_T, X_0)$. For the weighted ATE on first-period entrants, the score is multiplied by $g(X_0)/\Pr(S = 1)$.

D Validation and Sensitivity

This appendix reports validation diagnostics summarized in Section 4.4 of the main paper.

D.1 Covariate balance

Under sequential CIA, the relevant balance diagnostic is node-specific: at each decision node, the period-specific propensity score should equalize the covariate distributions of the groups that diverge at that node.[5] Table A2 reports the mean, median, and 95th percentile of the absolute standardized mean differences (SMD) at each node and pooled across all 495 covariate-node pairs. After inverse-probability reweighting, the pooled mean |SMD| falls from 0.116 to 0.048, the median from 0.049 to 0.032, and the 95th percentile from

[4] We update the R package *causalweight* from the two-period framework accordingly.

[5] At Node 1, the comparison is $D_1 = 1$ vs. $D_1 = 0$, reweighted by the first-period propensity score, with balance assessed on X_0. At Node 2, the comparison is $D_2 = 0$ vs. $D_2 \neq 0$ conditional on D_1, reweighted by the second-period propensity score, with balance assessed on X_1.

0.466 to 0.145, indicating that 95 percent of all covariate-node pairs exhibit imbalance below the standard 0.15 threshold. Balance is strongest at Node 2 conditional on $D_1 = 0$, where the 95th percentile drops from 0.492 to 0.094. Figure A6 confirms these patterns: the distribution of |SMD| shifts sharply toward zero after reweighting, with the right tail largely eliminated.

Table A2: Covariate balance at each decision node

		Raw			IPW-weighted		
Node	Covariate set (K)	Mean	Med.	P95	Mean	Med.	P95
Node 1: =1*vs.* =0	(170)	0.135	0.057	0.512	0.057	0.032	0.178
Node 2 | D_1=1: =0*vs.* ≠0	(155)	0.075	0.048	0.254	0.049	0.035	0.130
Node 2 | D_1=0: =0*vs.* ≠0	(170)	0.134	0.049	0.492	0.037	0.028	0.094
Pooled	All (495)	0.116	0.049	0.466	0.048	0.032	0.145

Mean, Med., P95: mean, median and 95th percentile of |SMD| across covariates. Raw: design weights W_i^s only. IPW: $W_i^s \times$ inverse node-specific propensity score. Categorical covariates are expanded into indicator dummies (one base omitted).

The residual imbalances at Node 1 are concentrated among spatial variables, peak commuting time to Luxembourg ($|\text{SMD}| = 0.41$), local unemployment rate (0.24), and public-transport connectivity (0.18 to 0.20), reflecting strong geographic sorting of cross-border workers into Belgian municipalities near the Luxembourg border, which is difficult to eliminate through reweighting alone. At Node 2, residual imbalances are smaller and dispersed across demographic and sectoral indicators.

D.2 Common support

Figure A7 shows kernel density estimates of the joint generalized propensity score (GPS), computed as the product of period-specific propensity scores, for the treatment sequence $(1,0)$ and the control sequence $(0,0)$. Panel (a) shows substantial overlap for the treatment-sequence GPS: the two realized-path groups have similar distributional shapes over the region where the GPS places most mass. Panel (b) reveals that for the control-sequence GPS, the distributions are shifted: path-$(0,0)$ observations concentrate at high values (median around 0.70), while path-$(1,0)$ observations concentrate near 0.10, reflecting the fundamentally different covariate profiles of cross-border and purely domestic workers.

Figure A6: Distribution of absolute standardized mean differences before and after reweighting

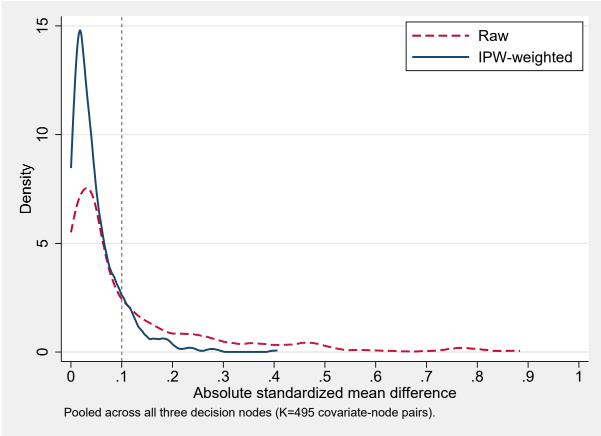

Distribution of absolute standardized mean differences across all covariate-node pairs, before (dashed) and after (solid) inverse-probability weighting. The vertical line marks the 0.10 threshold. The distribution pools 495 covariate-node pairs across the three decision nodes reported in Table A2.

Two remarks are in order. First, the separation in Panel (b) does not imply that inverse-probability weights are generally extreme. Each observation is weighted by the inverse of its own-path GPS: path-$(0,0)$ observations typically have joint GPS values near 0.70 and thus receive moderate weights. A left tail of control-path observations, predominantly individuals in border municipalities with low estimated probability of remaining purely domestic, has joint GPS below 0.10, generating larger weights. The effective sample size (ESS) for path 000 is approximately 44,000 (32 percent of 136,574 realized controls), against 463 (63 percent of 740) for path 010.[6] Second, the separation indicates that the doubly robust estimator relies partly on the outcome model to bridge the covariate gap between the two populations for the control counterfactual. Node-specific balance, reported above, provides reassurance: the propensity score equalizes covariate distributions within each period's selection margin, even though the joint GPS distributions appear separated. Cross-fitting further disciplines the outcome model's out-of-sample predictions and prevents overfitting to thin regions (Chernozhukov et al., 2018).

[6] ESS $= (\sum_i w_i)^2 / \sum_i w_i^2$, where $w_i = W_i^s / \widehat{\text{GPS}}_i$ combines design weights with inverse-probability weights. During estimation, 75 control observations with first-period predicted probability of exactly zero were excluded to avoid infinite weights.

Figure A7: Common support diagnostics

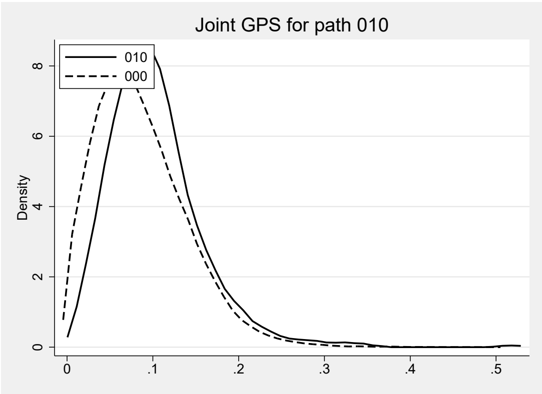

(a) Joint GPS for $\underline{d}_2 = (1,0)$

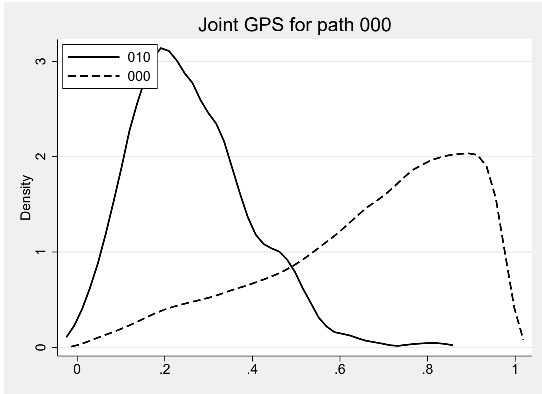

(b) Joint GPS for $\underline{d}_2^* = (0,0)$

Note: Each panel displays kernel density estimates of the joint generalized propensity score, $\hat{p}^{d_1}(X_0) \times \hat{p}^{d_2}(\underline{d}_1, \underline{X}_1)$, for individuals who followed the treatment sequence $(1,0)$ and the control sequence $(0,0)$.

D.3 Sensitivity to trimming

Following Bodory et al. (2022), we apply symmetric trimming on the normalized first-period propensity score at thresholds $\hat{p}_1 \in \{0.01, 0.02, 0.05, 0.10\}$. Table A3 reports the results. Mild trimming ($\hat{p}_1 \in [0.01, 0.99]$ and $[0.02, 0.98]$) removes 3,600 to 8,200 observations but slightly worsens balance (pooled P95 of 0.155 and 0.151 vs. 0.145 without trimming). More aggressive trimming ($\hat{p}_1 \in [0.05, 0.95]$ and $[0.10, 0.90]$) modestly improves balance (pooled P95 to 0.139 and 0.125), at the cost of excluding 27,000 to 63,000 observations, including up to 48 treated units, and discarding up to 44 percent of the sample. Given the modest balance gains and the sample loss, we retain the full untrimmed sample as the baseline. The doubly robust structure of the score provides an additional safeguard, as consistency requires only that either the propensity score or the outcome model be correctly specified (Chernozhukov et al., 2018).

D.4 Sensitivity to unobserved confounding (E-values)

We compute E-values (VanderWeele and Ding, 2017) as a partial sensitivity check on sequential CIA. The E-value answers the question: what minimum strength of association would an unmeasured confounder need to have simultaneously with treatment assignment and outcome, on the risk-ratio scale, to fully explain away a result? We convert additive ATEs to approximate risk ratios using the DML-estimated control potential outcome as

Table A3: Sensitivity of covariate balance to propensity-score trimming

Threshold	N	N_{treated}	N_{trimmed}	Pooled IPW Mean	Pooled IPW Med.	Pooled IPW P95	Node 1 IPW Mean	Node 1 IPW Med.	Node 1 IPW P95
None	147,027	740	0	0.048	0.032	0.145	0.057	0.032	0.178
$\widehat{p}_1^{d_1}(\underline{d}_0, \underline{X}_0) \in [0.01, 0.99]$	140,907	740	3,619	0.050	0.034	0.155	0.065	0.042	0.215
$\widehat{p}_1^{d_1}(\underline{d}_0, \underline{X}_0) \in [0.02, 0.98]$	136,349	740	8,177	0.050	0.034	0.151	0.063	0.040	0.195
$\widehat{p}_1^{d_1}(\underline{d}_0, \underline{X}_0) \in [0.05, 0.95]$	117,432	731	27,094	0.047	0.032	0.139	0.059	0.036	0.180
$\widehat{p}_1^{d_1}(\underline{d}_0, \underline{X}_0) \in [0.10, 0.90]$	81,529	692	62,997	0.044	0.032	0.125	0.049	0.031	0.164

Notes: Symmetric trimming is applied to the estimated first-period propensity score $\widehat{p}_1^{d_1}(\underline{d}_0, \underline{X}_0)$. N is the number of observations retained; N_{treated} counts treated-path observations (path 010) retained; N_{trimmed} counts excluded observations. Pooled IPW reports the mean, median and 95th percentile of |SMD| pooled across the three node-specific balance files after IPW reweighting. Node 1 IPW reports the same statistics for node 1 only.

the baseline risk. The headline magnitudes (CI-bound E-values of 1.75 for the one-year employment penalty, 2.76 for the one-year unemployment effect, 2.61 for the two-year unemployment effect, with the strongest observed predictor moving entry by a factor of 2.8) are reported in Section 4.4 of the main paper.

The breakdown by spell length and outcome reveals where confounding bites and where it does not. For UI receipt (Table A9), CI-bound E-values rise with spell length: 1.56 for one year of CBW, 1.70 for two years, 1.89 for three years. The longer-spell estimates are therefore the most resilient to unobserved confounding among the benefit outcomes. For inactivity (Table A6), wages and part-time work (Tables A7 and A8), and UI duration and amount (Tables A10 and A11), all CI-bound E-values equal 1, indicating these results are statistically fragile irrespective of confounding.

Two caveats. E-values are derived for a single unmeasured confounder in a static setting; in our sequential framework, multiple period-specific confounders with individually weaker associations could jointly produce equivalent bias, so the reported values are best read as a lower bound. For outcomes measured on a log scale, the risk-ratio conversion is approximate.

D.5 Alternative estimator: Extreme Gradient Boosting

We replicate the analysis using Extreme Gradient Boosting (XGBoost) instead of Random Forests. XGBoost is a boosting algorithm that constructs predictive models by sequentially combining a large number of weak learners, typically shallow decision trees, such that each new tree corrects the errors of the previous ensemble (Athey and Imbens, 2019; Mullainathan and Spiess, 2017). Whereas random forests primarily mitigate overfitting through tree decorrelation, XGBoost addresses both bias and variance by explicitly optimizing the loss function and incorporating regularization on tree depth and leaf weights, making it a well-suited alternative learner within dynamic DML (Chernozhukov et al., 2018; Bodory et al., 2022). Figure A8 reports the corresponding estimates. The findings mirror those from Random Forests, indicating that the results are not sensitive to the choice of machine learning algorithm.

Figure A8: One year after return: XGBoost replication

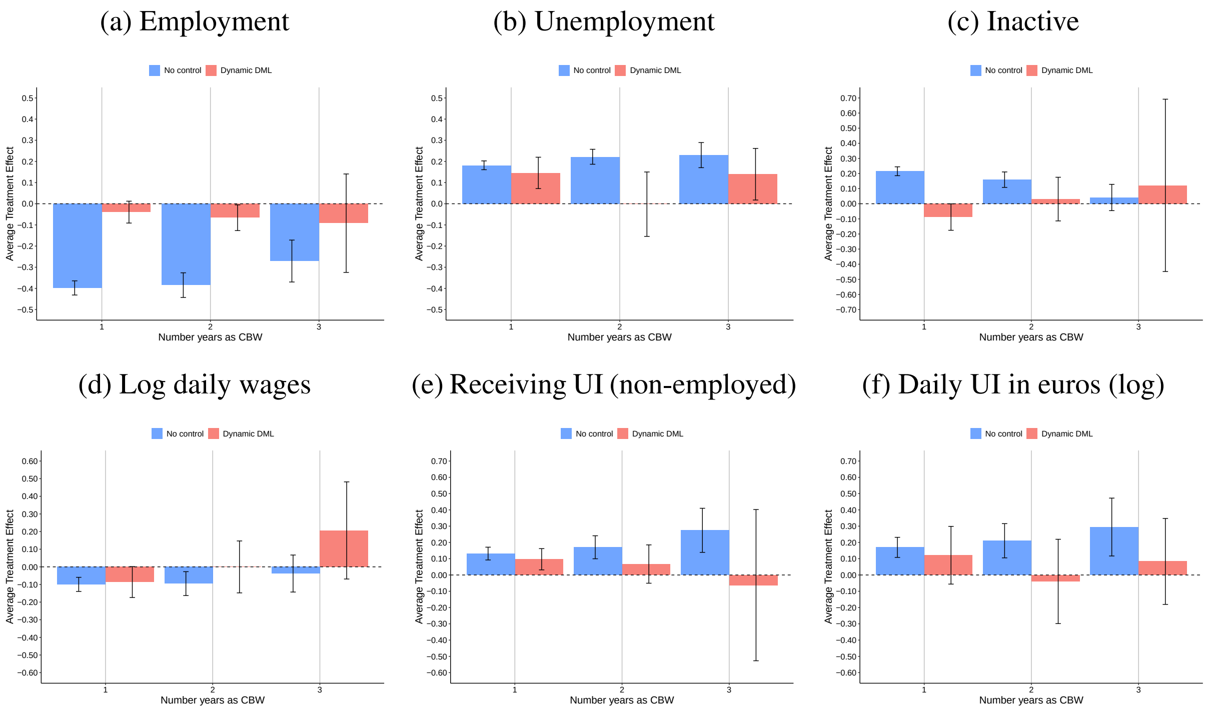

Effects of CBW one year after return, by spell length, estimated by dynamic DML with XGBoost. Top row: labor-market states (compare to top row Figure 2 of the main paper). Bottom row: log daily wages and unemployment benefits (compare to middle and bottom Figure 2 of the main paper).

E Full estimate tables

Tables A4 to A11 report point estimates, standard errors, p-values, and E-values for the main outcomes and treatment sequences considered in the paper. Long-run effects on log daily wages and part-time work follow the same patterns as the main-paper estimates and are not separately graphed: all such estimates are statistically indistinguishable from zero across the relevant horizons, and the corresponding numerical values are contained in the tables below.

Table A4: Causal DML: effect on Employment

	010	0100	0110	01000	01100	01110
Dynamic ATE	-0.098***	-0.109***	-0.067**	-0.074**	0.006	-0.076
Std. error	0.015	0.020	0.028	0.034	0.061	0.056
p-value	<0.001	<0.001	0.016	0.030	0.923	0.177
Mean outcome (treated)	0.274	0.279	0.322	0.312	0.391	0.310
Mean outcome (control)	0.371	0.388	0.388	0.386	0.386	0.386
N treated	740	446	241	184	98	79
N control	136,646	88,734	88,734	43,428	43,428	43,428
E-value (point estimate)	2.05	2.13	1.71	1.78	1.14	1.80
E-value (CI bound)	1.75	1.74	1.22	1.16	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Employment along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $** < 0.05$, * <0.10.

Table A5: Causal DML: effect on Unemployment

	010	0100	0110	01000	01100	01110
Dynamic ATE	0.121***	0.069***	0.121***	0.108***	0.070	0.165***
Std. error	0.018	0.021	0.031	0.035	0.050	0.050
p-value	<0.001	0.001	<0.001	0.002	0.161	0.001
Mean outcome (treated)	0.244	0.168	0.221	0.191	0.153	0.248
Mean outcome (control)	0.124	0.099	0.099	0.083	0.083	0.083
N treated	740	446	241	184	98	79
N control	136,646	88,734	88,734	43,428	43,428	43,428
E-value (point estimate)	3.36	2.78	3.87	4.06	3.10	5.44
E-value (CI bound)	2.76	1.87	2.61	2.32	1.00	3.01

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Unemployment along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $** < 0.05$, * <0.10.

Table A6: Causal DML: effect on Inactivity

	010	0100	0110	01000	01100	01110
Dynamic ATE	-0.027	0.035	-0.048	-0.024	-0.091	-0.124
Std. error	0.019	0.024	0.042	0.043	0.064	0.076
p-value	0.147	0.146	0.245	0.586	0.152	0.102
Mean outcome (treated)	0.399	0.469	0.386	0.433	0.365	0.333
Mean outcome (control)	0.426	0.434	0.434	0.457	0.457	0.457
N treated	740	446	241	184	98	79
N control	136,646	88,734	88,734	43,428	43,428	43,428
E-value (point estimate)	1.34	1.37	1.50	1.29	1.81	2.09
E-value (CI bound)	1.00	1.00	1.00	1.00	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Inactivity along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**$ $<$ 0.05, * <0.10.

Table A7: Causal DML: effect on Log of Daily Wages

	010	0100	0110	01000	01100	01110
Dynamic ATE	-0.083	-0.082*	-0.030	-0.172	0.091	0.067
Std. error	0.056	0.045	0.068	0.144	0.104	0.089
p-value	0.142	0.070	0.657	0.231	0.381	0.453
Mean outcome (treated)	4.132	4.161	4.214	4.098	4.362	4.337
Mean outcome (control)	4.214	4.243	4.244	4.270	4.271	4.270
N treated	203	150	74	71	34	29
N control	81,178	53,826	53,826	26,731	26,731	26,731
E-value (point estimate)	1.16	1.16	1.09	1.25	1.17	1.14
E-value (CI bound)	1.00	1.00	1.00	1.00	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Log of Daily Wages along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**$ $<$ 0.05, * <0.10.

Table A8: Causal DML: effect on Part-time Working Hours

	010	0100	0110	01000	01100	01110
Dynamic ATE	-0.073	-0.017	0.093	0.420	0.200	-0.093
Std. error	0.065	0.118	0.133	0.334	0.161	0.107
p-value	0.259	0.884	0.483	0.208	0.216	0.387
Mean outcome (treated)	0.245	0.264	0.374	0.717	0.498	0.206
Mean outcome (control)	0.318	0.281	0.281	0.297	0.298	0.298
N treated	203	150	74	71	34	29
N control	81,178	53,826	53,826	26,731	26,731	26,731
E-value (point estimate)	1.92	1.33	2.00	4.26	2.73	2.26
E-value (CI bound)	1.00	1.00	1.00	1.00	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Part-time Working Hours along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**<0.05$, * <0.10.

Table A9: Causal DML: effect on Receiving Unemployment Benefits

	010	0100	0110	01000	01100	01110
Dynamic ATE	0.078***	0.053**	0.120***	0.112**	0.118	0.209***
Std. error	0.022	0.027	0.038	0.048	0.073	0.078
p-value	<0.001	0.050	0.002	0.021	0.108	0.007
Mean outcome (treated)	0.313	0.273	0.340	0.311	0.317	0.408
Mean outcome (control)	0.235	0.219	0.220	0.199	0.199	0.199
N treated	498	272	155	96	50	42
N control	38,259	23,432	23,432	10,979	10,979	10,979
E-value (point estimate)	2.00	1.79	2.47	2.50	2.57	3.52
E-value (CI bound)	1.56	1.01	1.70	1.39	1.00	1.89

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Receiving Unemployment Benefits along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**<0.05$, * <0.10.

Table A10: Causal DML: effect on Number of Days Receiving Unemployment Benefits (log)

	010	0100	0110	01000	01100	01110
Dynamic ATE	-0.147	-0.024	0.166	-0.010	0.250	0.043
Std. error	0.145	0.125	0.173	0.197	0.206	0.243
p-value	0.311	0.846	0.336	0.958	0.227	0.860
Mean outcome (treated)	3.521	3.538	3.729	3.453	3.712	3.506
Mean outcome (control)	3.668	3.562	3.563	3.464	3.463	3.463
N treated	187	77	64	30	20	22
N control	9,274	5,603	5,603	2,598	2,598	2,598
E-value (point estimate)	1.25	1.09	1.27	1.06	1.35	1.12
E-value (CI bound)	1.00	1.00	1.00	1.00	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Number of Days Receiving Unemployment Benefits (log) along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**<0.05$, * <0.10.

Table A11: Causal DML: effect on Daily Unemployment Benefits in Euros (log)

	010	0100	0110	01000	01100	01110
Dynamic ATE	0.113*	-0.121	0.113*	-0.331**	-0.092	0.139
Std. error	0.067	0.082	0.060	0.165	0.174	0.158
p-value	0.089	0.143	0.058	0.045	0.597	0.380
Mean outcome (treated)	3.486	3.241	3.475	3.134	3.373	3.604
Mean outcome (control)	3.372	3.362	3.362	3.465	3.465	3.465
N treated	187	77	64	30	20	22
N control	9,274	5,603	5,603	2,598	2,598	2,598
E-value (point estimate)	1.22	1.23	1.22	1.45	1.19	1.24
E-value (CI bound)	1.00	1.00	1.00	1.05	1.00	1.00

Notes: This table reports Dynamic Double Machine Learning (DML) estimates using Random Forests for the effect of cross-border employment (CBW) on Daily Unemployment Benefits in Euros (log) along the treatment sequences (paths) shown in the column headings. Standard errors are clustered at the individual level. E-values report sensitivity to unobserved confounding and are computed from the ATE, its standard error, and the estimated control mean outcome. Significance: *** <0.01, $**<0.05$, * <0.10.