# Do Cash Benefits Attract Asylum Seekers? Evidence from a German Constitutional Court Ruling

Joop Adema and Panu Poutvaara*

November 11, 2026

**Preliminary and incomplete — please do not cite or circulate.**

**Abstract**

In July 2012, Germany's Constitutional Court mandated a 54-percent increase in cash benefits for asylum seekers. We exploit this unexpected policy change to estimate the causal effect of cash benefits on the size, composition, and geographic allocation of asylum flows using a gravity framework and monthly data on asylum applications across Europe. The ruling led to a sharp increase in asylum applications to Germany relative to other European destinations. After accounting for the diversion of applicants from other EU countries, we estimate an application-benefit elasticity of 1.0–1.1. The increase was driven entirely by applicants from origin countries with low asylum recognition rates. As a consequence, the number of recognized refugees in Germany rose only modestly, implying a substantially smaller refugee-benefit elasticity of approximately 0.24. Combining asylum application data with monthly Eurodac records on secondary asylum movements, we estimate that at least 56 percent of the additional applications were diverted from other EU destinations rather than newly attracted to the EU. Our findings suggest that more generous asylum benefits primarily reallocate asylum seekers across destination countries, with substantially smaller effects on refugee inflows to the European Union as a whole.

**JEL Codes:** F22, H53, I38, J15

**Keywords:** welfare magnet, asylum seekers, difference-in-differences, Germany, diversion

## 1 Introduction

The idea that generous social transfers attract immigrants—the *welfare magnet hypothesis* (Borjas, 1999)—has influenced the European political debate for decades and motivated restrictive benefit reforms in Denmark, Austria, the United Kingdom, and elsewhere (Giulietti, 2014; Allen, Fernández-Reino and Ruiz, 2025).[1] Yet credible causal evidence on whether, by how much, and through which margins benefits affect migration remains scarce. Governments that cut migrant benefits typically do

*Adema: University of Innsbruck (email: joop.adema@uibk.ac.at). Poutvaara: LMU Munich and ifo Institute (email: poutvaara@ifo.de). We thank Carmen Cano for excellent research assistance. We thank Sebastian Schirner and participants at Silvaplana conference and the internal seminar at the University of Innsbruck for helpful comments.

[1] Migrant and refugee benefits also have consequences for integration beyond dignitary considerations. For Denmark, Dustmann, Landersø and Andersen (2024*a*) find that a benefit cut raised short-run employment but also poverty and subsistence crime, and Dustmann, Landersø and Andersen (2024*b*) document harms to affected refugees' children; for the United States, LoPalo (2019) finds that more generous TANF benefits in the arrival year raise later wages without reducing employment.

so under political pressures that simultaneously change rhetoric, enforcement, and a host of related policies, so that the financial incentive itself is hard to disentangle from the policy bundle (Docquier and Vasilakis, 2024). Two additional dimensions are particularly under-studied: how the *composition* of asylum applicants changes with the level of benefits, and to what extent unilateral policy changes *attract* new applicants to Europe and to what extent they *divert* asylum seekers across alternative destinations. Whether more generous benefits draw genuinely new asylum seekers into Europe or merely redistribute them across member states—and whether they attract people who will ultimately be recognized as refugees—is central to the recurring policy debate over a "race to the bottom" in reception standards.

This paper exploits an unexpected judicial ruling that sidesteps the standard identification problem. In Germany, asylum seekers receive housing and financial assistance under the Asylum Seekers' Benefits Act (German: *Asylbewerberleistungsgesetz*) while their claims are processed, a period that can last from several months to several years. After the Act was enacted in 1993, the nominal level of financial assistance remained unchanged for nearly two decades, eroding its real value by about 19 percent. On July 18, 2012, the German Federal Constitutional Court (*Bundesverfassungsgericht*) ruled that the benefit levels were unconstitutionally low because they no longer guaranteed a subsistence minimum. The unexpected ruling was unrelated to asylum policy and required an immediate increase in benefits. For a single adult living outside a reception center, the monthly cash allowance increased from 225 to 346 euros (54 percent), while benefits for a family of two adults and two young children increased from 736 to 1,100 euros (49 percent). As a result, Germany's ranking in asylum seeker benefit generosity among European countries rose from 13th to 2nd. This abrupt change generated a large and plausibly exogenous increase in Germany's attractiveness relative to other European destinations. Applicants whose claims are rejected but who cannot be deported continue to receive these benefits, implying that the present value of a benefit increase is larger for individuals less likely to be granted asylum but who cannot be deported. This feature of the German asylum system implies that any response to higher benefits should be concentrated among applicants with relatively low expected recognition rates.

We combine this natural experiment with monthly Eurostat data on first-time asylum applications from 168 origin countries to 25 destinations (henceforth EU+).[2] We estimate a theory-driven difference-in-differences (DiD) gravity model on a symmetric window spanning 11 months before and after the ruling (August 2011–June 2012 vs. August 2012–June 2013, with July 2012 excluded). The specification includes origin-time and origin-destination pair fixed effects, so identification comes from changes in Germany's attractiveness relative to other EU+ destinations for applicants from the same origin country. Origin-time fixed effects absorb all time-varying origin-specific shocks common to EU+ destinations, including events such as the Arab Spring and the early stages of the Syrian civil war. The short estimation window further limits the scope for contemporaneous policy responses in other destination countries or broader structural changes in asylum flows. Under the identifying assumption that, absent the ruling, Germany's share of asylum applications from a given origin country would have remained unchanged, the estimated treatment effect captures the causal impact of higher benefit levels on destination choice. The richness of the data further allows us to examine heterogeneity across

[2]The EU+ set comprises the 25 destination countries with consistent monthly Eurostat coverage over the sample window: the 2012 EU member states excluding Austria, Cyprus, Hungary, and Malta, together with Norway and Switzerland, which participate in the EU asylum (Dublin) system.

origin countries and to distinguish between responses in asylum applications and recognized refugees, thereby revealing how benefit levels affect both the scale and composition of asylum inflows.

An important empirical concern is that the increased attractiveness of Germany diverted applications from destinations that constitute the control group: we show that the resulting bias is a function of the share of marginal applicants induced by the ruling that would otherwise have applied in another EU+ country – the diversion share – but that the bias is limited. The diversion share is also of primary interest, as it informs whether the policy induced additional applications to the EU+ or only reshuffled asylum applications across the EU+. To bound the diversion share, we proceed in two steps. First, we use novel monthly data from Eurodac, the European fingerprinting system for asylum applicants. This data describes the number of individuals who applied for asylum in a receiving country (e.g. Germany) who applied previously in a sending country (e.g. Greece) in a given month. The increase in onward asylum migration to Germany allows us to put a lower bound on the diversion share. In the second step, assuming the absence of any other shocks to the attractiveness of the EU+, we corroborate that bound for a subset of origin countries with relatively stable conditions by comparing the change in total applications in the EU+ (corrected for double counting using the Eurodac data) to the change in Germany using our DiD gravity estimate.

The paper has four core findings.

First, applications to Germany rose sharply after the ruling. Relative to a no-ruling counterfactual, applications increased by 61% if all marginal applicants would otherwise not have applied in the EU+ and by 45% if all marginal applicants would otherwise have applied in another EU+ country. Event study estimates show no pre-trend and a large, persistent post-period jump starting two months after the ruling. An extensive set of robustness tests shows that this effect is present in many countries of origin, is not driven by a few destination countries in the control group, does not strongly depend on the exact time window of analysis, does not appear in other years, and that Germany did not become more attractive for non-asylum migrants. Population-weighted linear regression in rates, synthetic control, and synthetic difference-in-differences (synthetic DiD) specifications yield comparable estimates.

Second, the response is driven by origin countries with low pre-ruling recognition rates: the marginal effect declines steeply in the recognition rate and is small for origins whose recognition rate in Germany exceeds about one in three (0.33). This pattern is not surprising given the larger present discounted value of benefits for applicants whose claims are ex ante less likely to succeed. In line with this, we also find that applicants from origins with longer procedures respond more strongly. This is consistent with the welfare magnet hypothesis of Borjas (1999): asylum applicants in Germany from countries with low recognition rates tend to be negatively selected (Guichard, 2020). Further heterogeneity analyses show that the strongest effects are found in origins that are safer, face few visa restrictions to travel to Germany, and are closer to Germany—consistent with low migration costs amplifying the benefit pull for applicants with weak protection claims. These are mostly Western Balkan countries, Russia, and parts of sub-Saharan Africa.

Third, as already suggested by the strong heterogeneity by recognition rate, the effect on the number of *recognized* refugees is much smaller than the effect on applications. Direct evidence from Eurostat's quarterly first-instance decision data shows a similar change in the number of total first-instance decisions ($\hat{\beta} = 0.494$), but only a small post-ruling change in *positive* decisions ($\hat{\beta} = 0.119$).

Fourth, onward asylum migration to Germany increased sharply after the ruling. We estimate that at least 56% of marginal applicants attracted by the ruling had already applied in another

EU+ country, which gives a lower bound on the diversion share of $\hat{\rho}_{D1} = 0.56$. A complementary stable-origin comparison indicates that the true share lies above that bound, placing the additional unobserved diversion at no more than 0.44, so that the total diversion share lies between 0.56 and roughly one. Two case studies (the Western Balkans and Chechens from Russia) make a considerable attraction effect plausible, but these are special cases given the low risk and migration costs of entering the EU. Knowing this range lets us pin down the magnitudes: over it, the corrected application-benefit elasticity is 1.0–1.1, while the recognized-refugee response is far weaker, with an implied refugee-benefit elasticity of about 0.24. In terms of the destination-choice model we use, this range implies that EU+ destinations are close substitutes for the marginal applicant (a dissimilarity parameter between 0.00 and 0.39, where one would mean independent tastes across destinations).

**Literature and contribution.** This paper contributes to three strands of the literature on the determinants of migrants' and asylum seekers' migration and location decisions.

*Push and pull factors of asylum migration.* On the push side, political instability, conflict, and violence in countries of origin drive the bulk of cross-border asylum flows (Hatton, 2009*a*; Neumayer, 2004; Hatton, 2004). On the pull side, destination-country restrictiveness shapes both the level and the allocation of applications (Neumayer, 2004; Hatton, 2004, 2009*a*, 2016), and a growing empirical literature documents the footprint of specific policy levers: recognition rates, processing times, and the broader policy mix (Bertoli, Brücker and Fernández-Huertas Moraga, 2022), designation of origin countries as "safe" (Guichard and Machado, 2025), family-reunification rules (Brekke, Røed and Schøne, 2023), visa policies (Czaika and de Haas, 2017), and employment bans (Di Iasio and Wahba, 2024). Di Iasio and Wahba (2024) caution, however, that economic conditions and pre-existing co-national networks quantitatively dominate destination choice among refugees within the EU. A few studies exploit single-country reforms for cleaner identification: Andersson and Jutvik (2023) show that Sweden's 2013 shift from temporary to permanent residence permits for Syrians sharply raised Syrian applications without affecting flows from other nationalities, while Abdelaaty (2023) and Avdan, Rosenberg and Gelpi (2025) find essentially no effect of rights violations or border-wall construction on refugee flows.

We add to this literature a clean causal estimate of an underexplored policy lever—the cash benefit paid *during* the asylum procedure—using a court-mandated ruling that is unusually large, sharply dated, and orthogonal to the political economy of asylum policy. Methodologically, we decompose the response by origin-country recognition rate and show that benefits operate almost entirely through applicants with weak protection claims; recognition rates are therefore not a clean indicator of policy restrictiveness, since the composition of applicants itself responds to other pull factors (Bertoli, Brücker and Fernández-Huertas Moraga, 2022).

*Diversion across destinations and multilateral resistance.* Random-utility migration models imply that the attractiveness of one destination spills over onto others, so that estimates ignoring this joint structure conflate effects on the level of migration with effects on its allocation across countries (Bertoli and Fernández-Huertas Moraga, 2013; Beine, Bertoli and Fernández-Huertas Moraga, 2016). Bertoli, Brücker and Fernández-Huertas Moraga (2016) illustrate the empirical relevance of this point: using bilateral migration flows to Germany during 2006–2012 and 10-year sovereign bond yields as a proxy for forward-looking expectations about alternative destinations, they show that the euro-area crisis diverted EU-internal migration toward Germany, and that ignoring expectations about alternative destinations biases standard discrete-choice estimates of migration determinants. Within the asylum

context, where the joint EU framework makes within-EU reallocations a first-order concern, Brekke, Røed and Schøne (2017) propose comparing the focal country's applications to those in a set of other destinations. Their approach—like much of the subsequent literature—relies on aggregate restrictiveness indices and on aggregate Eurostat counts that suffer from double-counting when applicants apply in more than one EU country. More recent work tackles multilateral resistance directly: Guichard and Machado (2025) study the spillover effects of destination-specific "safe-country" designations, and Görlach and Motz (2021) model strategic interaction of EU asylum policies through recognition rates, allowing a more attractive policy in one destination to both divert applications from others and raise the option value of seeking asylum in Europe.[3] Bratu et al. (2020) document analogous spillovers across family-migration policies, and Hatton (2023) shows that EU directives have measurably harmonized recognition rates across member states.

We contribute to this literature by exploiting a sharp single-policy shock rather than an aggregate restrictiveness index, and use monthly Eurodac data on onward asylum migration to net out double-counting and point-identify the share of the German increase that comes from within-EU reallocation. Onward migration is quantitatively important: Artuc and Ozden (2018) document that 9 percent of recent immigrants to the United States arrived from a transit country rather than their country of birth, with the transit share exceeding 30 percent among arrivals from many OECD origins. The Eurodac data we use also have important advantages over aggregate Dublin-transfer-request statistics, which miss both unrecorded transits and cases where the receiving country chooses not to lodge a Dublin request. In the years we observe, total Dublin requests are considerably lower than the number of cross-border onward asylum hits in Eurodac: across the countries the Eurodac file covers, 51,952 Dublin requests were lodged in 2012 against 103,790 cross-border hits, and 78,612 against 151,088 in 2013.

*The welfare-magnet hypothesis.* The theoretical foundation is Borjas (1999), who shows that more generous welfare states can both attract migrants and select for lower-skilled inflows. Empirically, correlational studies on asylum (Neumayer, 2004; Hatton, 2009*b*) and on regular migration (De Giorgi and Pellizzari, 2009; Razin and Wahba, 2015; Kaushal, 2005) report positive associations between benefit levels and inflows, but cannot separate the financial pull from confounding shocks. Within-country location-choice studies—Mossaad et al. (2020) on US refugee resettlement, Ferwerda, Marbach and Hangartner (2024) on Switzerland, Wett, Gallegos Torres and Steinmayr (2024) on Austria—estimate small, sometimes zero, elasticities of subnational mobility with respect to local benefit variation. Stated-preference evidence with Ukrainian refugees in Adema et al. (2025) confirms that benefits matter for destination preferences but substantially less than wages and employment opportunities. At the international margin, Agersnap, Jensen and Kleven (2020) exploit Denmark's repeated benefit cuts and hikes for non-EU migrants (including recognized refugees) and report a large net-migration elasticity of approximately 1.3; Adema (2025) reanalyses the same setting and bounds the flow elasticity at around 0.3 once contemporaneous policy bundles are accounted for, while documenting a substantially larger response on the application margin.

We provide, to the best of our knowledge, the first quasi-experimental estimate of a welfare-magnet

[3] In this class of models, every country has a target number of recognized refugees. A more attractive policy in one destination can therefore induce both a strategic increase in recognition rates in alternative destinations (other North- and West-European countries) and a strategic decrease in typical first-arrival countries such as Greece. For nearby origin countries, however, the option-value channel is less plausible: Western Balkan nationals traveling visa-free typically cross all intermediate countries (Austria, Croatia, Slovenia) within a day and are unlikely to update their utility of those countries en route.

effect that is specific to asylum seeker benefits. The closest prior quasi-experimental evidence studies Denmark's repeated benefit reforms for non-EU migrants and recognized refugees (Agersnap, Jensen and Kleven, 2020; Adema, 2025); those reforms, however, bundled the benefit change with contemporaneous shifts in rhetoric and enforcement, whereas the court-mandated German ruling isolates the benefit level itself. Existing studies otherwise typically rely on aggregate welfare spending or generic social-protection indices, which are imperfect proxies for the entitlements asylum seekers actually face (Di Iasio and Wahba, 2024; Mossaad et al., 2020). We also document a key empirical regularity that reconciles the disparate elasticities in this literature: benefits move asylum *applications* strongly but the *stock* of recognized refugees only weakly—in line with the small refugee-flow and internal secondary-migration elasticities reported by Ferwerda, Marbach and Hangartner (2024), Wett, Gallegos Torres and Steinmayr (2024), Adema et al. (2025), and Adema (2025)—because filing an additional application in another EU country is cheap while relocating is costly and the marginal applicant is rarely granted protection. The large application elasticities that dominate the policy debate therefore overstate the effect of benefits on the refugee population, and—because the response is largely diversion across the EU+ rather than new entry—unilateral benefit changes mostly reshuffle applications between member states, pointing to a role for EU-level coordination of reception standards. Although our setting concerns benefits paid *during* the asylum procedure rather than after recognition, integration mechanisms similar to those documented in the recognized-refugee literature are plausible given Germany's long procedure length and the twelve-month employment ban during our period.

The remainder of the paper proceeds as follows. Section 2 describes the EU asylum framework, the German benefits system, and the 2012 ruling. Section 3 introduces the data. Section 4 motivates asylum applicants' location decisions and presents the empirical strategy. Section 5 reports the application-margin results and robustness. Section 6 translates the application response into refugee-grant terms and examines actual decision counts. Section 7 uses Eurodac to decompose the German increase into attraction and diversion. Section 8 discusses our findings in light of the broader literature, and Section 9 concludes.

## 2 Institutional Background

### 2.1 Asylum Migration in the EU+

Under the 1951 Refugee Convention, signatory countries must examine asylum claims and cannot return applicants to places where they face persecution. The Common European Asylum System (CEAS) adds minimum reception, procedure, and qualification standards across the EU. Four non-member states—Norway, Iceland, Switzerland, and Liechtenstein—participate in the Dublin and Eurodac framework through association agreements; of these, Norway and Switzerland are part of our EU+ sample. The Dublin system (Dublin II throughout our sample window; its successor, Dublin III, entered into force in July 2013 and applies to applications lodged from January 2014) determines which member state is responsible for a given claim; in the default case, responsibility lies with the first country of irregular entry, and other member states may issue a "Dublin request" to transfer an applicant. Enforcement relies on the Eurodac fingerprint database, which we exploit in Section 7. In practice, Dublin transfers are leaky: humanitarian clauses, Dublin refusals, expiring deadlines, and low effective transfer rates mean that applicants retain meaningful destination choice within the EU+. At

the time of the ruling, effective Dublin transfer rates were on the order of 25 percent, with substantial cross-country heterogeneity (Appendix B.3).

### 2.2 The Asylum Seekers' Benefits Act

Germany's system of financial support for asylum seekers is governed by the *Asylbewerberleistungsgesetz* (AsylbLG), enacted in 1993. The Act was explicitly designed to provide a lower standard of support than regular social assistance. Benefits were set below the regular welfare level and, importantly, were not indexed to inflation.[4] Asylum seeker benefits consist of two components. The first, a personal allowance (*notwendiger persönlicher Bedarf*, colloquially *Taschengeld*), is paid in cash to every eligible adult and amounted to EUR 40.90 before the ruling. The second covers basic necessities, chiefly food (*notwendiger Bedarf*), and is provided either in cash or in kind depending on whether meals are supplied in the shelter; it made up the remainder of the EUR 225 monthly amount for a single adult.[5] The real value of both components eroded continuously from 1993 onwards (a real decline of roughly 19 percent over the period).

To understand how eligibility for asylum seeker benefits depends on the asylum trajectory, Figure 1 schematically shows three typical cases. Once registered as an asylum applicant, an individual receives asylum seeker benefits. If the claim is granted—after about seven months for initial decisions, or about fifteen months including appeals over the period we study (BAMF, 2017)—the now recognized refugee becomes eligible for regular social benefits (SGB II), which are set independently of the asylum seeker benefits. If the claim is rejected, the outcome depends on whether the individual is found to be deportable. If deported, benefit receipt ends; if the individual obstructs deportation, sanctions reducing the benefit level can be imposed. If the individual is not deportable (and may receive a tolerated status, *Duldung*), she continues to receive asylum seeker benefits. In 2012 the maximum term on the basic rate was 48 months: rejected but non-deportable applicants then move to so-called equivalent benefits (*Analogleistungen*), tied to the amount payable under SGB II.

Benefit receipt is conditional on registration: an arriving asylum seeker must first register with the authorities (recorded in the EASY system, Appendix G) before entering the asylum procedure, and only then becomes eligible for AsylbLG support. Without registration there are no benefits and no formal application on file, so the individual's case is not being processed at all. This gives arriving applicants a strong private incentive to register as quickly as possible rather than to delay or remain informal, which is also why EASY registrations—rather than the eventual formal application, which can lag behind at moments of high administrative load—are the more accurate real-time measure of when someone actually arrived (Appendix G).

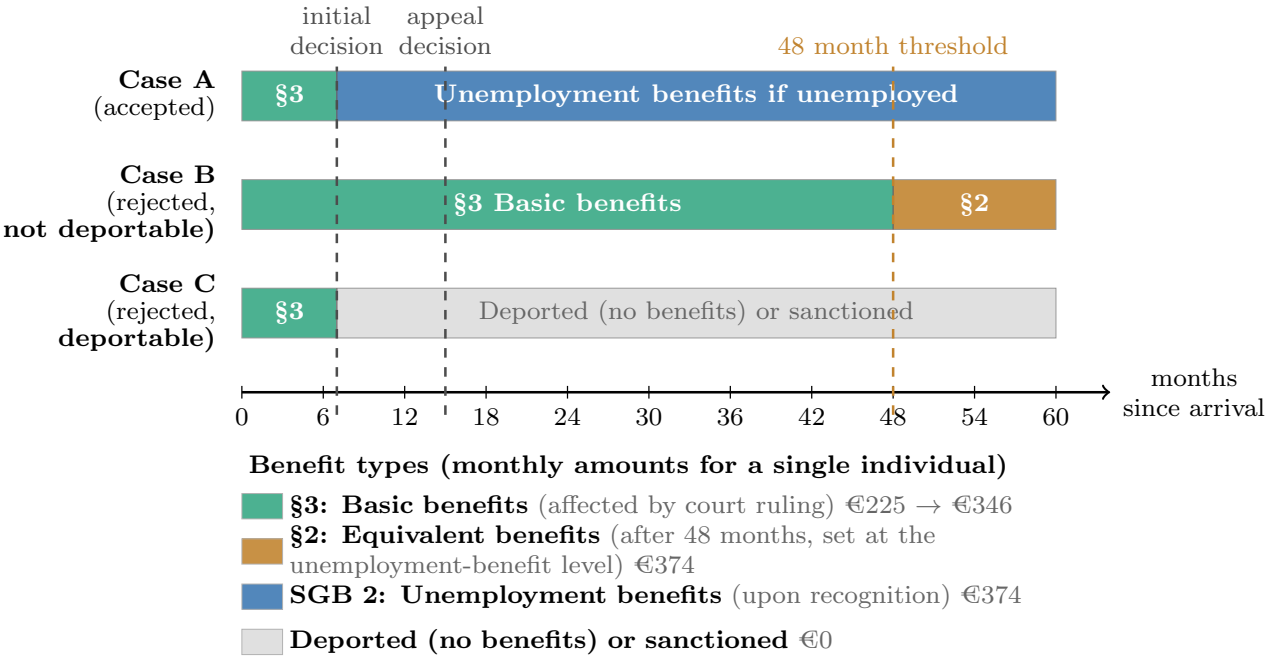
The discounted value of asylum seeker benefits is therefore larger for applicants with longer expected procedures and those more likely to be rejected without being deportable. Germany is comparatively generous to this latter group: rejected-but-tolerated applicants are moved onto benefits at the

[4] During the period studied, applicants were subject to an employment ban in their first twelve months and faced near-100 percent effective benefit-withdrawal rates against earnings after the employment ban.

[5] In December 2011 about 128,000 people received AsylbLG benefits, of whom roughly 95,000 received them in cash. The cash/in-kind split is closely tied to accommodation type: newly arrived applicants first stay in initial reception centers (*Erstaufnahmeeinrichtungen*) for up to three months, where meals are typically provided in kind, after which they are allocated across the federal states (Müller, 2013). A substantial share are then housed in collective accommodation centers (*Gemeinschaftsunterkünfte*), where food and basic necessities are frequently provided in kind and the cash allowance is correspondingly reduced, whereas those placed in decentralized housing (*Dezentrale Unterkünfte*) typically receive the full cash benefit. The prevalence of collective versus decentralized accommodation—and hence of in-kind versus cash provision—varied widely across the Länder during this period.

regular social-assistance (SGB II) level after 48 months, whereas neighboring destinations such as the Netherlands and Austria push rejected applicants toward departure-oriented, often in-kind, support rather than continued cash assistance. In addition, Kosyakova and Brücker (2020) show that asylum procedures are considerably longer for applicants eventually rejected than for those who receive a status. This suggests that asylum applicants who believe they are likely to be rejected—but who also do not expect to be deported, and so would continue drawing benefits rather than lose them—may be more responsive to changes in benefit levels.

Figure 1: AsylbLG regimes by trajectory: accepted, rejected-not-deportable, rejected-deportable




*Notes:* Schematic benefit trajectory by months since arrival for the three cases; amounts are the monthly cash entitlement for a single adult applicant living outside a collective reception center (Appendix A.1). § 3 basic benefits (*Grundleistungen*) comprise accommodation, food, heating, clothing, health-care products, household goods, and a monthly cash allowance for personal needs; they are paid during the procedure and are the benefits affected by the July 2012 ruling. On recognition (Case A) the applicant moves to SGB 2 unemployment benefits (unaffected by the ruling), while a rejected-but-not-deportable applicant (Case B) stays on § 3 until the 48-month threshold and then moves to § 2 equivalent benefits (*Analogleistungen*), set at the SGB 2 unemployment-benefit level. Neither basic nor equivalent benefits are conditional on job search, unlike unemployment benefits.

## 2.3 The Constitutional Court Ruling

On July 18, 2012, the Federal Constitutional Court delivered a landmark ruling that the level of AsylbLG benefits was incompatible with the constitutional guarantee of a dignified existence and the fundamental right to a minimum subsistence level. The court held that parliament's failure to index benefits to the general welfare standard over nearly two decades rendered the amounts constitutionally inadequate (Bundesverfassungsgericht, 2012*b*). The case (Az. 1 BvL 10/10 and 1 BvL 2/11) was brought by an Iraqi asylum seeker and a Liberian asylum seeker; the proceedings were focused on human-rights standards and bore no connection to migration policy debates.

The case's timing followed the Court's own procedural calendar rather than contemporaneous developments in asylum policy. Both underlying disputes originated as ordinary first-instance benefit claims: the case later docketed as 1 BvL 2/11 concerned benefits paid to the Liberian-born claimant for January–November 2007, and the case later docketed as 1 BvL 10/10 concerned benefits paid to

the Iraqi claimant for January 2009; the Sozialgericht dismissed both claims before they were appealed (Bundesverfassungsgericht, 2012*b*). On appeal, the Landessozialgericht Nordrhein-Westfalen held the basic AsylbLG rates unconstitutional and referred the question to the Federal Constitutional Court for binding review (*konkrete Normenkontrolle*), on July 26, 2010 and November 22, 2010 respectively (Landessozialgericht Nordrhein-Westfalen, 2010*b*,*a*). The referrals were then pending before the Court for close to two years: the Court publicly announced an oral hearing on May 30, 2012, held it on June 20, 2012, and delivered judgment four weeks later, on July 18, 2012 (Bundesverfassungsgericht, 2012*a*,*b*). Because the underlying claims, referral, hearing, and judgment dates were all fixed by ordinary litigation and constitutional-review procedure well before the 2011–2013 study window, the ruling's exact timing is plausibly unrelated to contemporaneous asylum-migration developments. Consistent with this, public attention to the AsylbLG itself only arose with the ruling rather than building up beforehand: weekly Google Trends search interest in the term *Asylbewerberleistungsgesetz* in Germany is low and sporadic throughout 2010–2012 and spikes sharply only at the ruling, settling at a persistently higher baseline afterward (Appendix A.3).

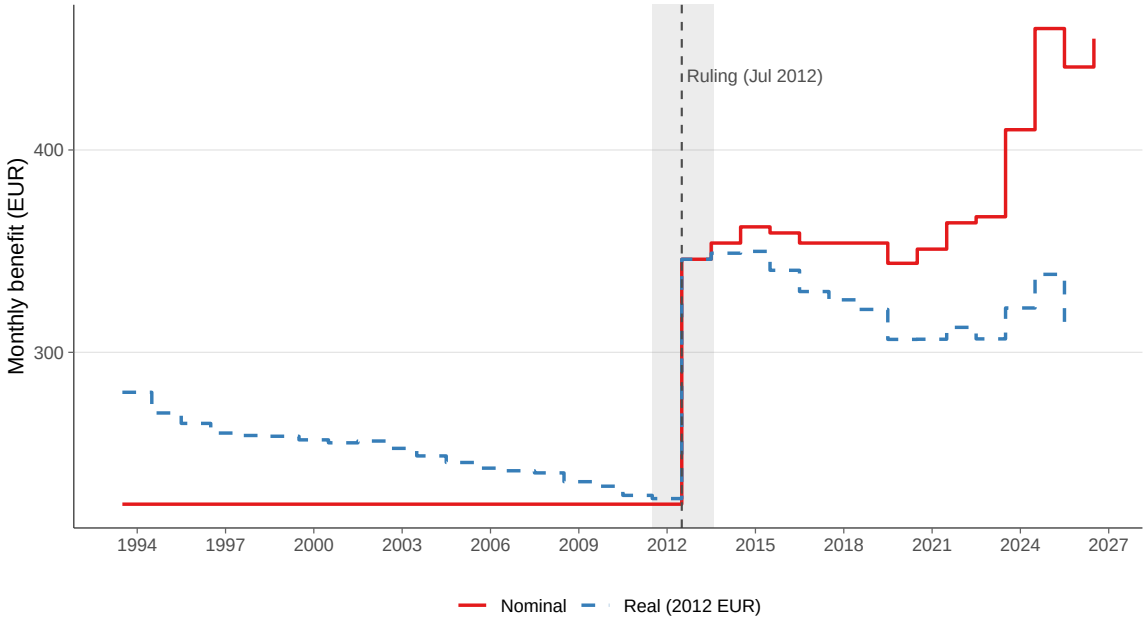
The ruling required benefits to be raised with immediate effect. State authorities applied the constitutionally mandated higher amounts from the date of the ruling, and parliament put them on a statutory footing later in 2012 (Flüchtlingsrat Berlin, 2012). The increase was permanent rather than a one-off adjustment: the single-adult rate stayed near or above its 2012 post-ruling level in every subsequent year, well above the frozen pre-ruling amount (Figure 2). For a single adult outside a collective center, the monthly cash allowance rose from 225 to 346 euros, an increase of 54 percent. For a family of two adults and two children (aged 5 and 10), total monthly benefits rose from 736 to 1,100 euros (49 percent).[6] Figure 2 traces benefit levels over time for representative household types. Appendix Table A1 documents the full schedule by need level from 1993 to 2026.

To compare Germany's pre- and post-ruling benefit level to those of other countries in the EU+, we collected information on personal and food allowances. To make the numbers comparable across countries, we calculate an equivalent amount for a single person without children (in the case of Germany, relative changes in benefit levels are very comparable across different family types), imputing a value for the food allowance for countries that provide food strictly in kind. For countries without an observed cash food allowance we add the cross-country mean gap between the cash-for-food and the full-in-kind regime in 2012, EUR 170 per month; Germany's own allowance is observed directly and not imputed. Figure 3 shows that Germany's pre-ruling benefit level was near the bottom of the cash-benefit distribution and that the ruling moved it close to the top. All monetary amounts in the paper are expressed in 2012 euros.[7]

[6] A consequential choice to make is to determine the base to calculate the percentage change in benefits, as it will be used to calculate the asylum application-benefit elasticity. We choose to use total potential monetary benefits, consisting of both the food and personal allowance. We choose to include the food allowance but exclude the value of housing. We do so because the amount of the first is within the locus of control and the latter because countries in the CEAS need to provide housing, most of which do so in-kind.

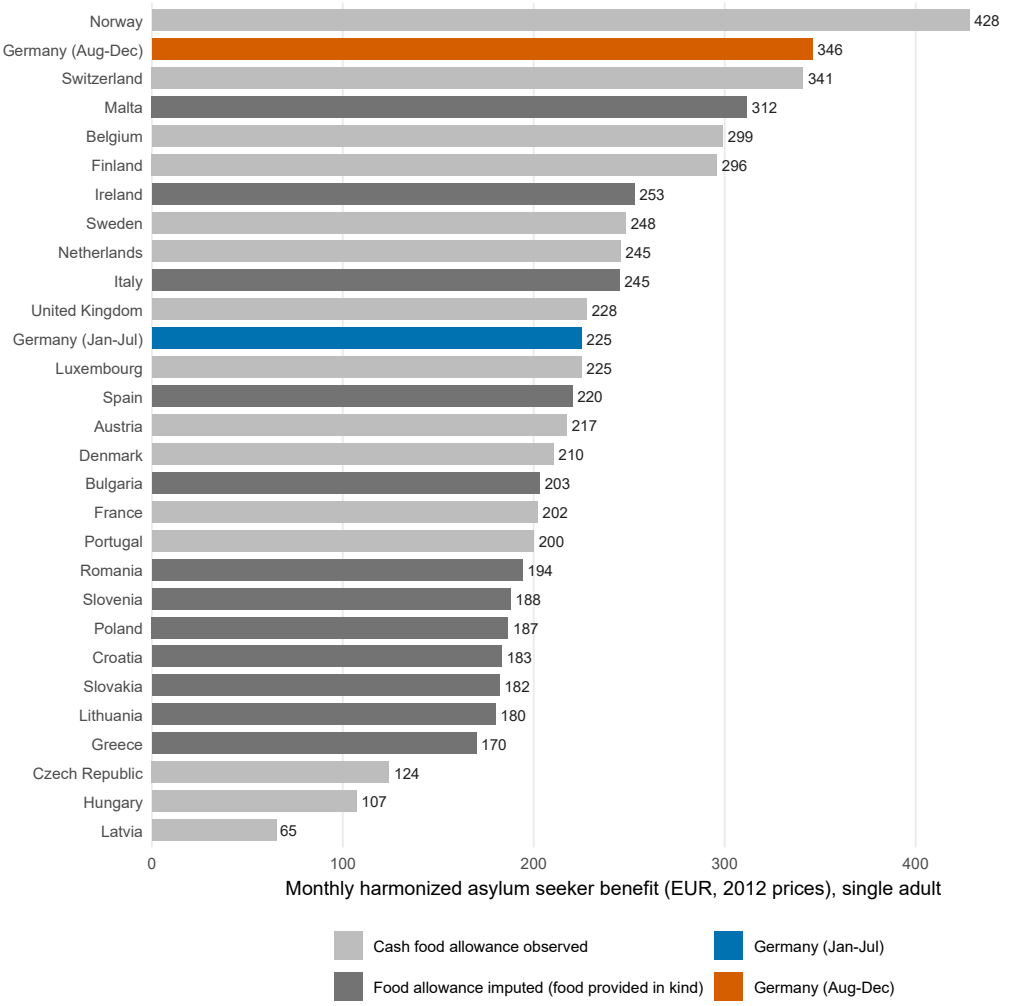
[7] Appendix A.2 reports the cross-country benefit levels we constructed for all EU+ over 2000–2025, and shows the same levels re-expressed in German 2012 prices (purchasing-power adjusted).

Figure 2: Asylum seeker benefits over time




*Notes:* Monthly cash benefit entitlements under the AsylbLG for asylum seekers living outside collective reception centers, in nominal and real terms. Single-adult benefits rose from 225 to 346 euros per month following the July 2012 constitutional court ruling.

Figure 3: Harmonized asylum seeker benefits in 2012 (nominal, 2012 EUR)




*Notes:* Monthly harmonized asylum seeker benefit for a single adult in 2012 euros (at market exchange rates), equal to the personal allowance plus the food allowance, sorted from highest to lowest. Bar shading distinguishes countries for which the food allowance is *observed* as cash (lighter gray) from those that provide food strictly in kind, for which we *impute* the food allowance (darker gray): for the latter we add to the observed pocket-money amount the cross-country mean gap between the cash-for-food and the full-in-kind regime among countries reporting both in 2012 (EUR 170). Germany's pre-ruling (Jan–Jul, blue) and post-ruling (Aug–Dec, black) levels are observed, not imputed, and are highlighted. Because this is a 2012 cross-section, the amounts are already in 2012 prices. Benefit levels were hand-collected from national legislation, AIDA country reports and EMN/OECD sources; Appendix A.2 lists the per-country sources and shows the same levels purchasing-power adjusted to German 2012 prices.

# 3 Data

## 3.1 Asylum Applications

The main outcome variable is the monthly number of first-time asylum applications filed in each EU destination country by applicants from each origin country, from Eurostat's asylum and managed migration statistics (Table `migr_asyappctzm`). We focus on first-time asylum applications (instead of total applications, as those include subsequent applications submitted after rejection as well) as it provides the best measure of asylum applicants who have recently come to the country of asylum and filed for an application. The data provide consistent coverage from January 2008 onwards. Eurostat additionally reports first-time applications broken down by applicant age and sex, but only at the destination × month level with origins aggregated; we use this breakdown for the demographic heterogeneity in Section 5.1.

Figure 4 plots normalized monthly asylum applications in Germany and in the rest of the EU+ combined.[8] Applications in Germany began to diverge around the ruling and rose sharply through 2012, dropped thereafter, but picked up again in 2013. By contrast, asylum applications in other EU countries rose only slightly. This suggests that the ruling increased the relative attractiveness of Germany, which we examine in detail in subsequent sections.

Over the course of 2014, asylum applications in Germany and the rest of the EU increased strongly—largely driven by Syrians fleeing the civil war, but also by other nationalities. Because the structure of asylum flows to Europe may have changed during the *European refugee crisis*, we restrict the sample period to a narrow window: the gray shaded area. This is a symmetric 11-month window: the pre-treatment period covers August 2011–June 2012 and the post-treatment period August 2012–June 2013. We exclude July 2012 as the ruling fell in the middle of the month. This window balances pre- and post-ruling months and keeps the estimation tight around the ruling.

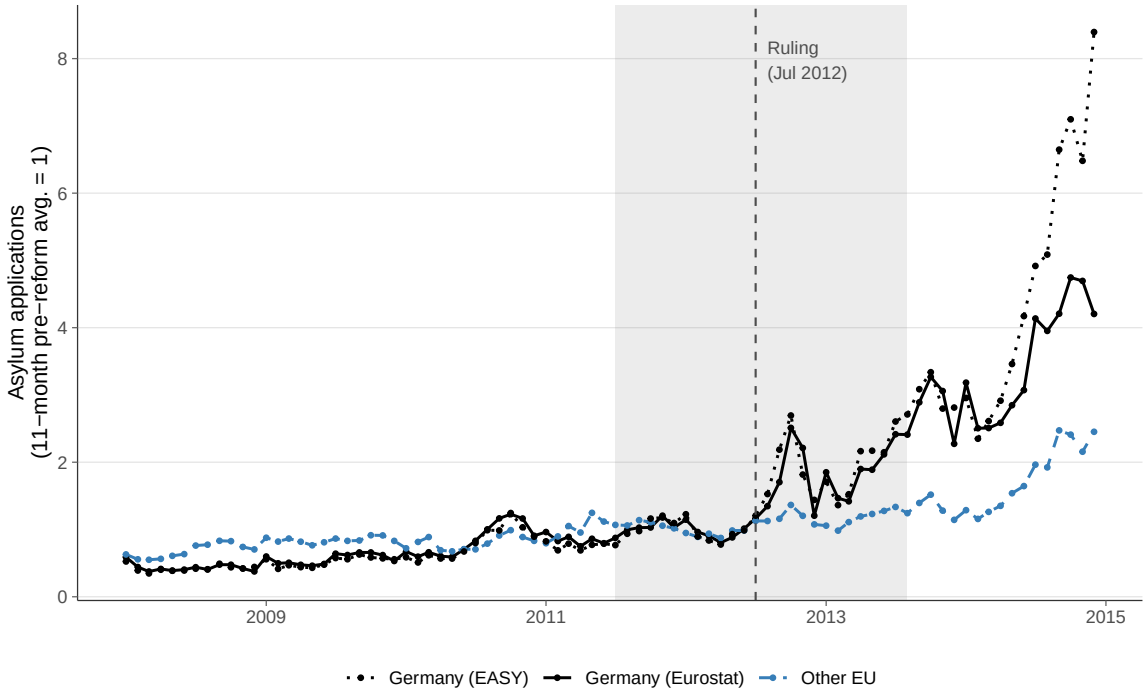
We exclude five destination countries throughout. Croatia is dropped because it acceded to the EU on 1 July 2013. Malta and Cyprus are excluded as small island destinations with highly variable flows that are unrepresentative of the broader EU+ pool. Austria is excluded because Eurostat does not disaggregate first-time from repeat applications for Austria over the sample window (only the total of first-time and repeat applications is reported). Hungary is excluded for the same reason: Eurostat does not report the first-time/repeat breakdown for Hungary until January 2013, leaving usable data for only the final six months of our estimation window—too short and too lopsided (entirely post-ruling) to serve as a control destination. On the origin side, we keep individual countries of origin and drop the Eurostat aggregate and residual citizenship categories—*Total*, *Unknown*, *Stateless*, and *Recognised non-citizens*—none of which corresponds to a well-defined country of origin; the same exclusions are applied to the first-instance decisions data. The resulting estimation sample covers 168 origin countries and 25 EU+ destination countries.

## 3.2 Asylum Decisions and Recognition Rates

To study the effect of the ruling on the number of refugee statuses granted and to examine heterogeneity of the main results for refugees across origin country recognition rates, we draw on quarterly data on first-instance asylum decisions from Eurostat (Table `migr_asydcfstq`). This data contains the number

[8] Appendix B plots the same series over the full 2008–2014 window (Figure A4) and lists the largest origin countries by application volume (Table A4).

Figure 4: Monthly asylum applications: Germany vs. rest of EU+




*Notes:* Monthly first-time asylum applications to Germany (black, solid) and to all other EU member states combined (blue, dashed) from Eurostat, and new asylum registration from Germany's asylum seeker registration system EASY, which is more accurate than the number of asylum applications when capacity to register asylum applications is exhausted (black, dotted). All three series are normalized to their respective 11-month pre-reform mean (August 2011–June 2012). The shaded area marks August 2011–June 2013 (the symmetric DiD estimation window); the vertical line marks July 2012.

of decisions by destination country, origin country, decision outcome, and quarter. We aggregate decision categories into *positive* (Geneva-Convention refugee, subsidiary protection, or humanitarian status) and *negative* (rejection). The Germany-specific pre-ruling recognition rate for origin $o$ is

$$R_o = \frac{\text{positive first-instance decisions in DE for origin } o,\ 2010\text{–H1 } 2012}{\text{total first-instance decisions in DE for origin } o,\ 2010\text{–H1 } 2012}.$$

$R_o$ varies enormously: Syrian, Iraqi, Somali, and Afghan applicants are recognized in the majority of cases, while Western Balkan recognition rates are below 5 percent throughout the period.

For the decisions analysis we use the quarters 2011 Q1–2013 Q4 (excluding the ruling quarter 2012 Q3), which brackets the monthly application window (August 2011–June 2012 versus August 2012–June 2013) used for applications. As destinations we keep Germany (the treated unit) and the other EU member states with consistent quarterly decision coverage, with two restrictions made for data reasons rather than design: we drop Italy, whose first-instance decisions surged in this window as it processed the 2011 Arab Spring cohort for reasons unrelated to the German ruling, and we exclude Syrian, stateless, and unknown origins from the aggregate, because German positive decisions for these categories spiked just before the ruling owing to a Syrian-conflict batch-processing decision rather than the ruling itself. These are the same exclusions applied when we estimate the decision-margin results in Section 6.

### 3.3 Onward Asylum Migration

Although the Dublin agreements (Section 2.1) should determine which EU country is responsible for an asylum applicant, it is common that asylum applicants apply in multiple countries. To study *onward* asylum migration (also referred to as secondary asylum migration, or asylum shopping) across the EU+, we use monthly aggregate Eurodac statistics obtained on request from EU-LISA.[9] The Eurodac system contains fingerprints of all asylum applicants, irregular border crossers, and apprehended irregular stayers in the EU.[10] When a receiving country queries a new applicant, the system reports any matches against records previously stored elsewhere. We use *Category 1–Category 1* hits as our measure of onward asylum migration—which indicate a new asylum applicant whose fingerprint matches an earlier asylum record elsewhere in the EU+.[11] This includes both recent asylum applicants in the sending country whose application is still pending and asylum applicants who were rejected in another country. As fingerprints for asylum applicants were stored for 10 years, this is plausibly a near-complete record of previous asylum applications.

The number of hits is reported on a monthly basis at the bilateral sender × receiving country level for 2008–2015. In this time period the Eurodac register unfortunately does not include information on the country of origin of the applicant. Appendix B.5 reports descriptive statistics on the cross-border hits—inward and outward hits per applicant by country and the quarterly hit share (Figures A7–A9). In 2012, cross-border Eurodac hits amounted to 34% of all asylum applications in the EU, the great majority of which (83%) are Category 1–Category 1 hits. In Germany, inward Category 1–Category 1 hits made up 12–25 percent of applications before the ruling, rising to 39–46 percent in the quarters afterwards, while outward Category 1–Category 1 hits were considerably smaller. This shows both the prevalence of onward migration and that Germany is a net receiver of onward asylum applications, indicating that it is a particularly favorable destination.

When destination countries receive an asylum application for which they claim another country is responsible, they can file a *Dublin request* at the sending country to take the applicant back. Cross-border onward applications are far more numerous than the Dublin requests that nominally govern them. A Dublin request is filed only when the receiving state actively seeks to transfer an applicant, and in practice the large majority of detected onward movements never result in a request—let alone an effective transfer (Fratzke, 2015). Over the years we observe, the yearly number of cross-border Category 1–Category 1 Eurodac hits substantially exceeds the yearly number of Dublin requests recorded in Eurostat's Dublin statistics (Tables `migr_dubri`/`migr_dubro`). Hence, the Eurodac data capture actual onward movement better and are available at higher frequency than the Dublin request data, which is crucial for our design.

[9] For institutional details of the Eurodac system, see eu-LISA (2015).

[10] In 2012 Eurodac comprised all 27 EU member states then in membership, together with the four associated states (Norway, Iceland, Switzerland, and Liechtenstein); Croatia joined only upon its EU accession in July 2013 and is therefore absent from the data we use.

[11] The data records the number of hits, not the number of individuals. However, as Category 1–Category 1 hits happen upon asylum application, it seems unlikely that this contains many duplicate individuals. Individuals have no incentive to apply in one country more than once, and the fingerprinting system prevents them from receiving benefits twice. Individuals could, however, show up as a hit in two receiving countries if they applied for asylum in more than two countries.

### 3.4 Additional Data

For heterogeneity analyses we use additional origin-country variables: great-circle bilateral distance (CEPII), GDP per capita (PPP) from the World Bank WDI, battle-related deaths per capita from UCDP/PRIO, the 2010 stock of co-nationals in Germany from UN DESA International Migrant Stock, and 2011 internet penetration from the ITU. To benchmark asylum flows against general immigration we use Destatis's *Wanderungsstatistik* (Genesis table `12711-0011`), which records all registered foreign arrivals to Germany by month. These arrival counts are based on population-register (*Melderegister*) movements and therefore include asylum seekers, who are registered on arrival.[12] To isolate non-asylum foreign immigration we therefore subtract monthly first-time asylum applications from the total, the series plotted in Section 5.2. Data on asylum benefits over time across countries are collected from primary national legislation and secondary sources, predominantly AIDA country reports; the construction of the cross-country benefit panel is described in Appendix A.2, and the full list of sources by country is given in Appendix Table A2. Monetary amounts are expressed in constant 2012 euros: each country's nominal benefit is deflated by its own Harmonised Index of Consumer Prices (HICP, rebased to 2012=100), so that 2012 values equal their nominal amounts.

## 4 Model and Empirical Strategy

We start by presenting a model of asylum migrants' destination choice to discuss the structural and policy-relevant parameters we aim to identify, how to estimate them, and the main threats to identification.

### 4.1 Random-utility Model of Asylum Migration

Consider a potential asylum seeker $i$ from origin $o$ who decides, at time $t$, whether and where to apply for asylum. The choice set comprises the EU+ destinations $d \in \mathcal{D}$ and an outside option of staying home (or applying outside the EU), indexed by $o$. Following McFadden (1974, 1978), the utility of destination $d$ decomposes into a deterministic part $W_{odt}$ and an individual $i$-specific taste shock $\varepsilon_{iodt}$,

$$V_{iodt} = W_{odt} + \varepsilon_{iodt}, \tag{1}$$

$$W_{odt} = \alpha\, b_{dt} + \mu_{odt}. \tag{2}$$

Here $b_{dt}$ is the policy change in Germany: $b_{dt} = 1$ from the month after the constitutional court ruling onward for $d =$ Germany only and $b_{dt} = 0$ otherwise. $\alpha$ is the effect of the ruling on Germany-specific utility for a prospective asylum applicant. $\mu_{odt}$ captures all other deterministic elements of destination-specific utility that are constant across individuals. The outside option, staying in the home country, has deterministic value $W_{oot}$. Following Bertoli and Fernández-Huertas Moraga (2013), we assume that $\varepsilon_{iodt}$ is distributed differently for the destination countries and for the outside option of not applying for asylum in Europe: the utility shock is correlated across destinations, governed by

[12] Destatis migration statistics count registrations and deregistrations of cross-border moves recorded in the municipal population registers (*Melderegister*); asylum seekers assigned to a reception facility are registered and so enter these counts. For the methodology see Statistisches Bundesamt (Destatis), *Qualitätsbericht Wanderungsstatistik*, and GENESIS-Online table 12711-0011 (`https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Bevoelkerung/Wanderungen/_inhalt.html`).

the dissimilarity parameter $\tau$, which is 0 if utility shocks are perfectly correlated and 1 if utility shocks are independent. The model yields a closed form solution of migration probabilities.

**Choice probabilities.** Utility maximization of Equation (1) under the distributional assumption yields a convenient multiplicative expression for the probability that an origin-$o$ applicant chooses destination $d$ at time $t$ (hereafter: destination choice probability):

$$m_{odt} = \underbrace{\frac{\Sigma_{ot}^{\tau}}{\Sigma_{ot}^{\tau} + e^{W_{oot}}}}_{\text{migrate}} \cdot \underbrace{\frac{e^{W_{odt}/\tau}}{\Sigma_{ot}}}_{d \text{ conditional on migrating}}, \qquad \Sigma_{ot} \equiv \sum_{k \in \mathcal{D}} e^{W_{okt}/\tau}. \tag{3}$$

The choice probabilities only depend on the deterministic components of utility across countries, $W_{odt}$, and the parameter $\tau$. The first factor is the probability of migrating at all (versus staying) and the second factor is the conditional probability of choosing $d$ given migration.[13] The expected number of bilateral asylum applications is given by $\mathbb{E}[M_{odt}] = N_{ot}\, m_{odt}$, where $N_{ot}$ is the origin population size.

**The diversion share.** The migration probabilities show that an increase in the attractiveness of a destination country $d$ affects migration flows through two mechanisms. First, it affects the probability of migrating at all, through $\Sigma_{ot}^{\tau}$; second, it affects the conditional probability of migrating to country $d$. The relative strength of these two margins, which rely on structural parameters, determine whether the response mostly diverts asylum applicants across destination countries or attracts applicants that otherwise would not have migrated. In addition, because of the possibility to apply in more than one country, diverted applicants may or may not still first apply in another destination country. We can thus subdivide marginal applicants of the German reform into three categories:

- **A** (attraction): would not have applied for asylum in the EU+ absent the ruling; appears only in the asylum application statistics in Germany, with no matching record elsewhere.
- **D1** (observed diversion): also applied in another EU+ country, so the application in Germany generates a cross-border Eurodac hit.
- **D2** (unobserved diversion): would have applied in another EU+ country absent the ruling but, given the higher benefit, applies only in Germany without ever filing elsewhere; like type A, appears only in the asylum application statistics in Germany.

EU policy makers may be interested in the degree to which reforms divert asylum flows, that is, in the diversion share. This can be written as:[14]

$$\rho = \rho_{D1} + \rho_{D2} = \frac{D1 + D2}{A + D1 + D2} \in [0, 1]. \tag{4}$$

Sections 7.1 and 7.2 offer two complementary routes to bounding $\rho$: estimation of the lower bound $\rho_{D1}$ using cross-border Eurodac fingerprint matches, and corroborating evidence from changes in asylum applications from stable origin countries to the whole EU.

[13] When $\tau \to 1$, the expression simplifies to a multinomial logit model.

[14] Appendix D.2 shows that the diversion share can be linked to structural parameters of the model and the baseline migration probabilities.

### 4.2 Estimation

Taking the natural logarithm of $\mathbb{E}[M_{odt}]$, substituting (2) into (3), and rearranging yields the following gravity model equation (Anderson and van Wincoop, 2003; Ortega and Peri, 2013; Beine, Bertoli and Fernández-Huertas Moraga, 2016):

$$\ln M_{odt} \;=\; \beta\, b_{dt} \;+\; \gamma_{ot}, \qquad \beta \;\equiv\; \frac{\alpha}{\tau}, \tag{5}$$

where $b_{dt}$ is the indicator for the policy change. The coefficient on $b_{dt}$ is $\beta = \alpha/\tau$: we can only identify $\alpha$ up to the dissimilarity parameter $\tau$ (Beine, Bertoli and Fernández-Huertas Moraga, 2016). The estimate of $\beta$ thus provides an upper bound of $\alpha$, whereas point identification of $\alpha$ requires knowledge of $\tau$.[15]

$\gamma_{ot}$ captures the size of the origin country pool $N_{ot}$, as well as the *multilateral resistance* term (Bertoli and Fernández-Huertas Moraga, 2013). Those terms depend on the attractiveness of the alternative destinations and the value of staying—through the sum $\Sigma_{ot}$ and the denominator $\Sigma_{ot}^{\tau} + e^{W_{oot}}$, as seen in (3). Under the distributional assumptions for $\varepsilon_{iodt}$, those multilateral resistance terms only vary at the origin-time level: when Germany raises its benefits, the multilateral resistance changes proportionally across all destinations.

Estimating (5) by ordinary least squares on log counts faces two well-known problems. First, the bilateral panel has many zero cells and the log of zero is undefined; common patches such as omitting the zero cells or using the $\ln(M_{odt}+1)$ transformation do not enable estimation of a (semi-)elasticity and yield misleading inference (Chen and Roth, 2024; Fitzgerald et al., 2025). Second, even without zeros, $\hat{\beta}$ is biased under heteroskedasticity (Santos Silva and Tenreyro, 2006). We therefore estimate (5) by Poisson pseudo-maximum likelihood (PPML), which only requires the conditional mean to be correctly specified. We estimate the following model:

$$M_{odt} \;=\; \exp\!\Big(\beta\, \mathrm{DE}_d \times \mathbb{1}(t > t_0) \;+\; \phi_{ot} \;+\; \psi_{od}\big( \;+\; \phi_{odm}\big)\Big) \cdot \epsilon_{odt}, \tag{6}$$

where $b_{dt}$ is included through the interaction of $\mathrm{DE}_d$, a binary indicator for Germany, and $\mathbb{1}(t > t_0)$, a binary indicator for the period after the ruling date $t_0$ = July 2012. $\phi_{ot}$ denote origin × year-month fixed effects, which capture multilateral resistance and the origin country population, as well as any other origin-specific push shocks that drive asylum migration such as the onset of conflict or deteriorating economic conditions. $\psi_{od}$ is an origin × destination (pair) fixed effect, which captures all time-independent factors that drive the level of asylum migration between $o$ and $d$. In several specifications we also include destination-month or pair-month ($\phi_{odm}$) fixed effects, to account for seasonal patterns in asylum applications. $\epsilon_{odt}$ is a multiplicative error term with conditional mean one.

To assess pre-trends and study dynamic effects, we replace the post-indicator with event-time dummies relative to June 2012 (the month before the ruling) as the omitted reference,

$$M_{odt} \;=\; \exp\!\left(\sum_{k \neq \text{June 2012}} \beta_k\, DE_d \times \mathbb{1}[m_t = k] \;+\; \phi_{ot} \;+\; \psi_{od}\right) \cdot \epsilon_{odt}, \tag{7}$$

[15] Appendix D.2 shows that an estimate of the diversion share, together with the observable pre-ruling application shares, inverts to a point estimate of $\tau$ and hence of $\alpha = \beta\tau$; see equation (A7).

Pre-treatment coefficients $\beta_k$ for $k < 0$ test the parallel-trends assumption directly; post-treatment coefficients $\beta_k$ for $k \geq 0$ trace the dynamic response of asylum applications in Germany.

### 4.3 Identification of the policy-relevant parameter

Although the previous section shows how $\beta$, $\alpha$ and $\tau$ can be identified, those are not the only parameters of interest. For example, a German policy maker may want to know by how much the ruling changed asylum applications relative to a counterfactual without the ruling. We call this policy-relevant parameter $\pi$, with the relative effect of the reform equal to $(e^{\pi} - 1) \cdot 100\%$. Difference-in-differences estimation with multiplicative models relies on the assumption that treated and control units follow parallel trends in proportions, i.e. that growth rates are equal across both groups. Even if asylum applications in Germany and the control destinations were on parallel trends, the DiD estimate $\hat{\beta}$ is not an unbiased estimate of $\pi$: the control units are themselves affected by diversion of applications from the control group to Germany, so the no spillover (SUTVA) assumption is violated and $\hat{\beta}$ overestimates $\pi$.

To recover $\pi$ we need the diversion share $\rho_{D2} \in [0, 1]$, the share of marginal applicants in Germany who would otherwise have applied in another EU+ destination but did not because of the ruling (type D2 in Section 4.1), which is discussed in detail in the following section. This is best illustrated by a $2 \times 2$ Poisson DiD—Germany ($T$) versus the rest of the EU+ ($C$), pre versus post—which is representative of our high-dimensional estimator because all fixed effects are origin-nested, so $\hat{\beta}$ is a weighted average of origin-specific $2 \times 2$ contrasts.[16] The number of diverted applicants from the control group equals the number of diverted marginal applicants by construction, which allows us to derive an analytic expression of the policy-relevant treatment effect as a function of the estimate $\hat{\beta}$, the diversion share $\rho_{D2}$ and the pre-period ratio of asylum applications in the treated and control groups $\frac{M_0^T}{M_0^C}$ (see Appendix D.3 for a derivation):

$$e^{\hat{\pi}} - 1 \;=\; \frac{e^{\hat{\beta}} - 1}{1 + e^{\hat{\beta}} \dfrac{M_0^T}{M_0^C} \rho_{D2}}. \tag{8}$$

At $\rho_{D2} = 0$, $\hat{\pi} = \hat{\beta}$, whereas the correction is largest as $\rho_{D2} \to 1$ (pure diversion). Hence, to obtain a lower (upper) bound for the policy-relevant estimate we can estimate $\hat{\beta}$ and set $\rho_{D2} = 1$ ($\rho_{D2} = 0$). Because $\frac{M_0^T}{M_0^C}$ is not very large, the lower and upper bound are relatively close to each other.

### 4.4 Threats to Identification

Studying a policy change in a single country raises several identification concerns. We list them below and address each in Section 5.3.

*Origin-country factors and structural breaks.* The most immediate origin-country concern is Syria. The conflict escalated over the first half of 2012, and the summer of that year brought the first considerable increase in Syrian asylum applications in Europe, at a moment when European recognition rates for Syrian applicants were also rising steeply. The two developments coincide with the ruling in calendar time, which is why we treat Syria separately throughout: column (2) of Table 1 drops

[16] Empirically, collapsing the estimation sample to Germany versus the combined rest of the EU+ (summed over origins) yields a near-identical estimate to the main specification (Appendix Table A12).

Syrian origins altogether, and the recognition-rate heterogeneity of Section 5.1 places Syria among the high-recognition origins for which we find little response. More broadly, the Syrian refugee crisis could contaminate the analysis: refugees fleeing large-scale civil war may differ systematically from earlier arrivals, and the 2015 surge in applications from many origins at once may have caused a structural break in asylum migration—altering irregular migration routes and the feasibility of reaching Germany, especially once Germany suspended Dublin returns. To keep these forces from affecting our estimates, we restrict attention to a short symmetric 11-month window around July 2012 and, in robustness checks, drop origin countries with large idiosyncratic shocks and re-estimate leaving out each origin in turn (Appendix G). A related concern is the Dublin III Regulation, which added obligations on receiving states to verify that a transfer would not expose the applicant to harm; it entered into force in July 2013 and applies only to applications lodged from January 2014, so it lies outside the estimation window and cannot drive our results, though it might otherwise have affected destinations asymmetrically.

*Germany-specific factors.* The estimates would be biased if, within the same window, Germany enacted other policies affecting asylum migrants, or if Germany's relative or absolute attractiveness changed for unrelated reasons. On the first, no notable German asylum-policy change falls inside the analysis window. On the second, one might worry that, as the EU recovered from the euro crisis, Germany's attractiveness rose faster than that of the harder-hit southern economies under austerity—in particular, the euro-area sovereign-debt crisis made Germany differentially attractive on the labor-market dimension. We address this by comparing asylum applications with non-asylum foreign immigration to Germany over the same window, using immigration statistics from the German statistical office (Section 5.2).

*Other destination-country factors.* Other EU+ destinations may have changed their own reception or asylum policies over the window, which the gravity design does not absorb, and several did. Sweden moved early and in the same direction as Germany, though for one specific group: it stopped enforcing returns to Syria in January 2012 and treated the conflict as severe enough that protection became close to automatic, so the expected probability of a positive decision for Syrian applicants rose well before the better-known September 2013 switch to permanent residence permits (Andersson and Jutvik, 2023). Belgium moved the other way, and is the largest cumulative tightener among the European destinations between 2010 and 2013 in the DEMIG POLICY data, although it does so gradually across the whole window rather than in 2012 in particular (Appendix A.4). Luxembourg belongs to the group of destinations that designated origin countries as safe over these years, a lever whose cross-border externalities Guichard and Machado (2025) document directly, and it moved on the benefit margin in the same month as the German ruling: the grand-ducal regulation of 8 June 2012 cut the monthly cash allowance paid to applicants for international protection, by 23.5 percent for a single adult catering for himself and by more than three quarters for one accommodated with meals provided (Appendix C.3). Luxembourg thus moved in the opposite direction to Germany at almost exactly the same moment, which is why it turns up among the most extreme destinations in the placebo distribution of Appendix C.3. Romania illustrates a different mechanism altogether: its asylum applications over this period are dominated by transit movements that shifted with the routes out of North Africa after the Arab Spring, so what appears in the data as a sharp fall in a control destination reflects a change in routes rather than a policy response.

None of these changes are absorbed by the origin $\times$ time fixed effects, so we gauge their aggregate

importance against the DEMIG POLICY indicators for 2011–2013 (Appendix A.4) and probe them in the robustness analysis (Section 5.3). Most directly, column (7) of Table 1 drops Sweden, Belgium, Luxembourg, and Romania from the control group; the coefficient falls from 0.476 to 0.385 but remains large and statistically significant. The main specification uses all EU+ destinations as controls; restricting the control group to narrower sets of destinations, and leaving out each control destination in turn (Appendix G), shows that no single control destination has more than a minor effect on either the point estimate or the $t$-statistic.

*Residual multilateral resistance.* The model in Section 4.1 imposes a single dissimilarity parameter $\tau$, constant across origins and destination pairs, so that all EU+ destinations are equally substitutable within the migrate nest. This is a genuine limitation: empirically plausible substitution patterns—for example a tighter Germany–Austria–Switzerland cluster—would require additional structure on the joint distribution of taste shocks and introduce multilateral-resistance terms that vary at the origin–destination–time level, so that a German policy change could trigger stronger substitution out of Austria than out of, say, Spain. The same heterogeneity implies that the response itself can differ across origins, which we examine directly in the heterogeneity analysis (Section 5.1). Because $\alpha = \tau\,\beta$, the implied $\hat{\tau}$—and hence $\alpha$—should be read as a population average over the relevant pool of substitutes rather than a pair-specific object; the direction of any resulting bias in $\alpha$ is a priori ambiguous, depending on whether the marginal applicants substitute more or less strongly than the average pair. The policy-relevant treatment effect $\pi$, by contrast, does not rely on $\tau$: it is recovered from the Poisson estimate together with the directly estimated diversion share (Section 7), which already nets out these substitution patterns.

*Policy response.* The ruling anecdotally raised applications especially from Western Balkan countries (Section 7.3). In response, the BAMF introduced an administrative fast-track for Western Balkan claims in October 2012, shortening processing times from roughly three months to nine days and so cutting the period of benefit receipt for that group. Because this shortens the discounted benefit gain precisely where the response was strongest, it biases our estimates for low-recognition origins toward zero. The episode affects only a small set of origins—Western Balkan countries that had recently gained visa-free travel—and is documented in Section 5.1 and Appendix I.

*EU-wide attention effects.* A remaining risk is that a German reform salient enough to draw broader attention to the EU asylum channel could raise awareness of asylum-seeking in the EU+ generally, not just in Germany specifically: someone who first learns of the option because of the German ruling might end up applying elsewhere in the EU+ rather than in Germany itself, a channel our diversion/attraction accounting is not designed to detect. This is conceptually related to the option-value mechanism in Görlach and Motz (2021), where a policy change in one destination can reshape the perceived value of the EU asylum channel as a whole rather than only the payoff to applying in the reforming country. We cannot test this channel directly, but two pieces of evidence bound its likely size: the stable-origins analysis (Section 7.2) shows that, for origins with stable push factors, total EU+ applications did not rise after the ruling—if anything they fell slightly—suggesting that attraction to the EU+ at large is unlikely to be a first-order force; and the Frontex external-border detection counts (Appendix B.4) show no jump around the ruling, consistent with no broad increase in new arrivals to the EU+ as a whole. Neither piece of evidence directly tests this specific channel, however, and we regard it as an open threat to identification rather than one we can fully rule out.

### 4.5 Inference

Treatment is assigned at the destination × time level and Germany is the only treated destination, so the design has a single treated cluster and cluster-robust standard errors overstate precision (Conley and Taber, 2011; Ferman and Pinto, 2019). Two families of procedures have been developed for panel data with a single treated cluster, which differ in whether they assume that untreated potential outcomes are exchangeable across destinations or across periods (Almeida et al., 2025). The first motivates approaches that permute units, in which treatment status is permuted across control units and the observed $t$-statistic is compared to the distribution of $t$-statistics across permutations (MacKinnon and Webb, 2020). The second motivates approaches that permute time periods. An approach that additionally accounts for serial correlation is conformal inference, in which the treated unit's own residual series is rotated cyclically through every starting point and the statistic on the unrotated post-reform block is ranked against the rotations (Chernozhukov, Wüthrich and Zhu, 2021).[17]

However, both unit and time exchangeability are unlikely to fully hold in practice. Exchangeability across units fails because control destinations differ strongly in the scale of asylum applications. Small destination countries receiving few asylum seekers typically face larger relative changes than larger destinations such as Germany. Moreover, simultaneous policy changes in control countries strongly affect the variance and lead to large and significant placebo estimates.[18] Permutation-inference $p$-values are therefore overly conservative.

Two further features of this setting keep the placebo draws using control units from being clean realizations of the null, and both arise because the control destinations are affected by Germany's reform. The first is diversion. Germany's gain in applications came in part at the expense of the other destinations, and that loss was not spread evenly across them: destinations that are close substitutes for Germany among the origins that responded gave up proportionally more applications than destinations that are not. Every candidate for reassignment therefore absorbed some part of the treatment, and the ones that absorbed most are exactly those whose reassigned indicator picks up the largest negative movement, so the placebo distribution is drawn from units that are heterogeneously treated rather than untreated. The second concerns Germany itself, which stays in the panel as a control unit in every placebo draw. Germany is one of the largest destinations and its applications rise sharply after July 2012, so it pushes up the origin × year-month effects in the post-ruling months, which biases every reassigned coefficient downwards.

Exchangeability across periods fails because other shocks may affect the attractiveness of Germany at different times. In principle, the number of pre- and post-periods can be restricted so as to exclude periods containing other policy changes, which would otherwise lead to overrejection. However, since the approach of Chernozhukov, Wüthrich and Zhu (2021) controls size better with many pre-periods, we extend the pre-period by two more years. This approach may therefore also be conservative if other shocks occur during the pre-period.

Table A6 decomposes the variance of asylum application rates across destinations and time periods, following the diagnostic of Almeida et al. (2025). The diagnostic shows that the variance across destinations is considerably larger than across time periods, suggesting that unit exchangeability is

[17] Both are rank-based tests of the sharp null, so p-values are discrete and cannot fall below $1/K$, where $K$ is the number of permutations.

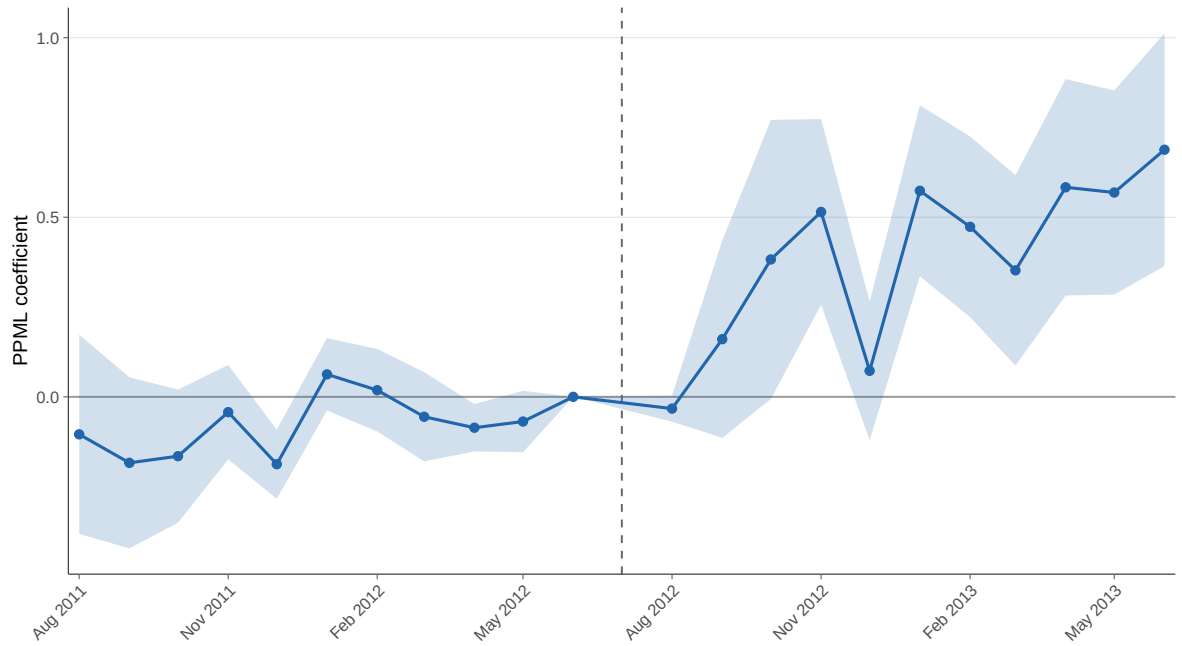
[18] Policy changes in control units also affect the main estimate for Germany. However, individual control countries are small and contribute little, and because policy changes across control countries go in opposite directions, their impact on the PPML estimate for Germany is limited.

unlikely to hold. Taken together with the two contaminations just described, this is why we do not report permutation inference in the main text. Our main regression tables report cluster-robust standard errors and $p$-values alongside conformal $p$-values; the permutation procedure is estimated throughout and reported in Appendix C.3, where its placebo distribution can be shown draw by draw and the destinations driving it named, which is what makes its $p$-values interpretable. Because of the limited flexibility of both permutation-based approaches we show only cluster-robust standard errors for the heterogeneity analyses and 95% confidence intervals in the event-study plots. For more details on the inferential procedures, see Appendix C.

## 5 Asylum Applications

Figure 5 plots the monthly event-study coefficients $\hat{\beta}_k$ from equation (7). The pre-treatment coefficients are close to zero with no systematic trend, consistent with parallel trends before the ruling. The coefficient for August 2012, the first month after the ruling, is still close to zero; from September onwards the coefficients rise steeply and remain elevated, in line with the pooled estimate of about 0.476, a 61 percent increase. Applications dip for a single month in December 2012 before rebounding immediately.

Figure 5: Event study: monthly PPML estimates




*Notes:* PPML event-study coefficients from equation (7), with 95 percent confidence intervals (clustered at origin and destination level). The omitted category is June 2012 (one month before the ruling); July 2012 is excluded from the sample. Sample: August 2011–June 2013, 25 destination countries and 168 non-EU origin countries. The dashed vertical line marks the ruling date (July 2012).

Table 1 reports the main DiD estimates over the same period for a range of samples. These estimates are not corrected for the diversion that depresses applications in the control group, and so exceed the policy-relevant treatment effect (Section 7). The full-sample coefficient is 0.476, implying a 61 percent increase in applications in Germany relative to the rest of the EU+. Given the 54 percent benefit increase, the uncorrected application-benefit elasticity is 1.1; under full diversion ($\rho_{D2} = 1$,

Appendix D.3) the policy-relevant effect falls to 45%, an elasticity of 0.8.[19]

Columns (2)–(4) test whether a few origins whose inclusion could be questioned drive the estimate. They drop, in turn, Syria; all Middle East and North Africa origins, whose applications over the window were shaped by acute conflict rather than by the benefit change; and the Western Balkan origins, all of which except Kosovo had obtained visa-free Schengen access shortly before the ruling, and for which we document a distinct attraction pattern in Section 7.3.[20] The coefficient is fairly stable across all three omissions, somewhat larger when Middle East and North Africa origins are dropped and somewhat smaller when the Western Balkans are dropped.

Columns (5) and (6) instead restrict the control destinations to more comparable groups: the old EU+ and the North-West. The coefficient is stable across these columns and rises slightly when the controls are confined to North-Western Europe. Column (7) drops the four control destinations whose asylum policies or migration routes changed during the window, namely Belgium, Luxembourg, Romania, and Sweden (Section 4.4). The coefficient falls to 0.385, a 47 percent increase, but remains large and statistically significant.

Table 1: Main difference-in-differences results

| | | Excluded origins | | | Control destinations | | |
|---|---|---|---|---|---|---|---|
| | Full | Syria | Excl. MENA | West-Balkan | Old EU | North-West | Excl. BE, LU, RO, SE |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| DE × $\mathbb{1}(t > t_0)$ | 0.476*** | 0.531*** | 0.613*** | 0.417*** | 0.512*** | 0.545*** | 0.385*** |
| | (0.102) | (0.105) | (0.119) | (0.096) | (0.129) | (0.159) | (0.089) |
| Destination × Origin FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Origin × Year-month FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 25,577 | 25,054 | 20,806 | 24,167 | 21,391 | 14,653 | 20,360 |
| Origins | 168 | 167 | 148 | 163 | 168 | 168 | 168 |
| Destinations | 25 | 25 | 25 | 25 | 16 | 11 | 21 |

*Notes:* PPML estimates of equation (6). The dependent variable is monthly first-time asylum applications from Eurostat. All specifications include destination × origin and origin × year-month fixed effects. Standard errors (in parentheses) are clustered at the origin and destination country level. The treatment variable equals one for Germany after July 2012. The pre-treatment period is August 2011–June 2012 and the post-treatment period is August 2012–June 2013 (July 2012 excluded). Column (2) drops Syria; column (3) drops all Middle East and North Africa origins (World Bank regional classification); column (4) drops the Western Balkan origins (Albania, Bosnia and Herzegovina, Serbia, Montenegro, Kosovo, and North Macedonia). Columns (5)–(7) instead change the control destinations: column (5) restricts them to the pre-2004 EU members plus Norway and Switzerland ("Old EU+"), column (6) to the North-Western destinations (Belgium, Denmark, Finland, Germany, Ireland, Luxembourg, Netherlands, Norway, Sweden, Switzerland, and the United Kingdom), and column (7) drops Belgium (BE), Luxembourg (LU), Romania (RO), and Sweden (SE), the destinations with contemporaneous policy or route changes discussed in Section 4.4.
*Inference.* The bracketed rows give $p$-values for the same coefficient under two procedures; Appendix C sets out the assumptions each requires and how they fail here. *Cluster-robust* is the $p$-value implied by the standard error above it. *Conformal* is the moving-block test of Chernozhukov, Wüthrich and Zhu (2021) applied to Germany's own monthly log gap between observed and counterfactual applications, which is rotated cyclically through every starting point and the unrotated post-ruling block ranked against the rest; it is run on a longer window carrying a 35-month pre-ruling period from August 2009, so $T = 46$ and the smallest attainable value is 0.022. Permutation inference across destinations is estimated for every column but reported in Appendix C.3 rather than here, for the reasons given in Section 4.5. Stars refer to the cluster-robust $p$-value: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

[19] Figure A14 plots the corrected percentage as a function of $\rho_{D2} \in [0, 1]$.

[20] Visa-free travel to the Schengen area was granted to Serbia, Montenegro, and North Macedonia in December 2009 (Council Regulation (EC) No 1244/2009) and to Albania and Bosnia and Herzegovina in November 2010 (Council Regulation (EU) No 1091/2010). Citizens of Kosovo remained subject to a visa requirement throughout the sample period.

## 5.1 Heterogeneity

The aggregate effect masks substantial heterogeneity. We look first at *who* responds—the demographic composition of the applicants—and then at *which* origins respond and why.

**Individual characteristics: age and sex.** The Eurostat age × sex breakdown is reported only at the destination level, with origins aggregated, so this disaggregation cannot be taken to the origin dimension: the binned series does not record the country of origin (Section 3.1). Within that limit, Table 2 re-estimates the ruling effect for each age × sex bin on the destination × month panel, alongside a baseline column for total applications on the same panel. The response is broadly homogeneous across the age × sex bins, if anything slightly larger for children, and falls off only for the elderly (65+), where it is small and statistically insignificant. This homogeneity is not an artifact of the benefit increase itself varying by household composition: as noted above, the percentage increase was similar for a single adult (54 percent) and for a family of two adults and two children (49 percent), so the reform did not differentially incentivize any particular age × sex group through the level of the increase.

Table 2: Heterogeneity by age and sex

| | Baseline | Age 0–13 | | Age 14–17 | | Age 18–34 | | Age 35–64 | | Age 65+ | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (all) | M | F | M | F | M | F | M | F | M | F |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
| DE × $\mathbb{1}(t > t_0)$ | 0.447*** | 0.448*** | 0.455*** | 0.384*** | 0.297*** | 0.441*** | 0.425*** | 0.406*** | 0.448*** | 0.168 | 0.114 |
| | (0.064) | (0.096) | (0.098) | (0.093) | (0.072) | (0.058) | (0.064) | (0.063) | (0.072) | (0.106) | (0.107) |
| *p-values* | | | | | | | | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [0.114] | [0.289] |
| Conformal | [0.022] | [0.022] | [0.022] | [0.174] | [0.065] | [0.022] | [0.022] | [0.022] | [0.022] | [0.022] | [0.304] |
| $N$ | 622 | 596 | 574 | 578 | 468 | 621 | 575 | 621 | 551 | 330 | 330 |
| Pre-reform total | 305,495 | 30,410 | 27,810 | 15,490 | 5,340 | 121,900 | 43,210 | 37,820 | 21,205 | 990 | 1,320 |

*Notes:* Each column is a separate PPML difference-in-differences on the destination × month panel (destination and year-month fixed effects; standard errors clustered on destination). The age × sex series are reported by Eurostat at the destination level only, so the regression has no origin dimension. The observation count varies across columns: Eurostat reports each age × sex bin for a different set of destination-months, and destination-months that are singletons under the fixed effects (or that report zero throughout) drop out of the corresponding regression, so the sparser bins—most visibly the small 65+ group—are estimated on fewer destination-months. "Pre-reform total" sums each bin's applications across all destination-months in the 11-month pre-reform window (Aug. 2011–Jun. 2012), to indicate group sizes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under two procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it, and the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point. The conformal counterfactual is the destination × month fit rather than the gravity one, estimated over the window extended back to August 2009, so $T = 46$ and the smallest attainable value is 0.022. Permutation inference across destinations is estimated for every column but reported in Appendix C.3 rather than here, for the reasons given in Section 4.5.

**Origin characteristics.** The data carry little information on individual applicants, but how the response varies across origin countries is itself informative about who the marginal applicants are. Table 3 interacts the post-ruling DiD indicator with one origin moderator at a time, and jointly in the final column. The moderators fall into two groups. The Germany-specific recognition rate $R_o$ (on its native $[0, 1]$ scale) and the processing-time proxy measure how long a typical applicant draws benefits. Log distance to Germany, the DEMIG VISA inward indicator (Germany, 2012), and conflict deaths per 100k measure the cost and nature of the move, while the bilateral migrant-network share, internet penetration, and log GDP per capita PPP measure how readily the change is learned about

and acted on and the broader economic pull. All continuous moderators are standardized. Under a welfare-magnet mechanism, the response should be largest where the discounted gain from the higher benefit is largest and the cost of acting on it lowest, a reading we develop when interpreting the estimates below.

Table 3: Heterogeneity by origin characteristics

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| DE × $\mathbb{1}(t > t_0)$ | 0.775*** | 0.311* | 0.358*** | 0.551*** | 0.163** | 0.523*** | 0.564*** | 0.352*** | 0.216 |
| | (0.134) | (0.181) | (0.095) | (0.102) | (0.075) | (0.108) | (0.076) | (0.085) | (0.172) |
| × Recognition rate | -1.15*** | | | | | | | | -0.440 |
| | (0.271) | | | | | | | | (0.272) |
| × Processing time (std.) | | 0.140 | | | | | | | 0.222*** |
| | | (0.169) | | | | | | | (0.071) |
| × Visa not needed | | | 0.830*** | | | | | | 0.758** |
| | | | (0.255) | | | | | | (0.361) |
| × Conflict deaths p.c. (std.) | | | | -0.046*** | | | | | -0.024*** |
| | | | | (0.007) | | | | | (0.009) |
| × Log distance (std.) | | | | | -0.331*** | | | | -0.204** |
| | | | | | (0.065) | | | | (0.088) |
| × Log GDP p.c. PPP (std.) | | | | | | -0.107* | | | -0.015 |
| | | | | | | (0.063) | | | (0.029) |
| × Internet users share (std.) | | | | | | | 0.370*** | | 0.089 |
| | | | | | | | (0.050) | | (0.089) |
| × Migrant network share (std.) | | | | | | | | 0.263*** | -0.001 |
| | | | | | | | | (0.055) | (0.082) |
| Destination × Origin FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Origin × Year-month FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 23,261 | 22,948 | 25,577 | 24,966 | 25,029 | 24,222 | 24,966 | 24,966 | 20,569 |

*Notes:* PPML estimates of equation (6) augmented with origin-characteristic interactions; one column per moderator, plus the joint specification in the last column. Each row reports the interaction of the post-ruling indicator (DE × $\mathbb{1}(t > t_0)$) with the labeled moderator; the top row is the non-interacted ruling effect. "Std." in a label denotes a moderator standardized to mean 0 and SD 1 across origins; the recognition rate $R_o$ enters on its native $[0, 1]$ scale. The moderators are: the Germany-specific recognition rate (2010–H1 2012); the processing-time proxy (Germany, June 2012); the DEMIG VISA inward indicator (Germany, 2012; baseline "visa needed", reported interaction with "visa not needed"); conflict deaths per capita (origin, 2011–2013); log distance origin-to-Germany; log GDP per capita PPP (origin, 2011–2013); the internet-user share (origin, 2012); and the bilateral migrant-network share (origin-to-Germany stock relative to origin population, 2010). All specifications include destination × origin and origin × year-month fixed effects. Standard errors (in parentheses) are clustered at the origin and destination country level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Columns (1)–(8) each interact the ruling indicator with a single origin moderator; column (9) enters them jointly. The bivariate results point consistently to one economic margin: the response is concentrated among applicants for whom the discounted value of the higher benefit is largest and the cost of acting on it is lowest.

*Recognition rate (column 1).* The dominant moderator is the origin's pre-ruling recognition rate. The interaction is large, negative, and highly significant: the proportional response falls from $e^{0.78}$ for an origin never recognized ($R_o = 0$) to $e^{0.78+-1.15}$—essentially zero—for an origin recognized with certainty ($R_o = 1$). This is exactly the welfare-magnet prediction: applicants whose claims are almost always rejected, but who cannot be deported, draw asylum seeker benefits for the longest (Section 2.2), so the present-discounted gain from the increase is largest for them. Figure 6 shows the split graphically: applications from origins with a recognition rate below 0.33 (roughly one positive decision in three) jump sharply at the ruling and stay elevated, while those above it are flat with no pre-trend. The gradient does not depend on how the recognition rate is measured or how it enters the regression: it survives quadratic and binned functional forms (Appendix G) and alternative definitions of $R_o$ (Appendix G).

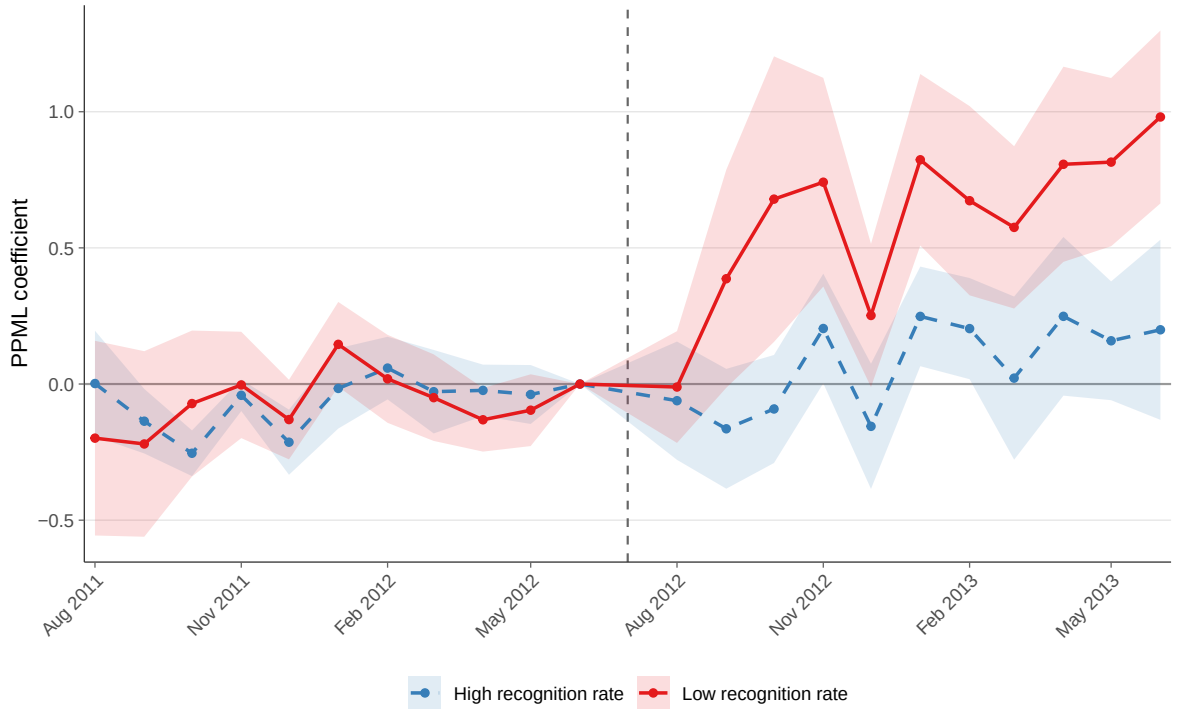
*Migration cost and claim strength (columns 2–6).* The remaining bivariate moderators trace the same margin from different angles, and all are signed as the mechanism predicts. The response is significantly stronger for origins closer to Germany (negative log-distance interaction, column 5) and less conflict-affected (negative conflict-deaths interaction, column 4)—origins whose nationals are plausible economic rather than protection migrants and who face low travel costs—and it is markedly larger for origins with visa-free Schengen access (column 3), the sharpest single proxy for a low-cost outside option. Because the EU sets a common visa list, an origin whose nationals can enter Germany without a visa can also enter every other EU member state except the United Kingdom and Ireland without one, including the transit countries Austria and Slovenia. This interaction therefore captures low-cost access to the EU as a whole, not an entry advantage specific to Germany. The log-GDP interaction is negative and significant at the ten-percent level (column 6), so poorer origins respond somewhat more. Longer expected procedures (column 2) carry the predicted positive sign—more months on benefits imply a larger discounted gain—but the bivariate estimate is not individually significant.

*Information frictions (columns 7–8).* Two moderators capture how readily the change is learned about and acted upon. Both the origin's internet-user share (column 7) and the size of the pre-existing co-national network in Germany (column 8) raise the response significantly, consistent with information about the German change diffusing faster, and migration being less costly, where connectivity and networks are already in place.

*Joint specification (column 9).* The moderators are strongly collinear—low-recognition origins are disproportionately visa-free, nearby, peaceful, and well-networked—so the joint column cannot cleanly separate them: the recognition-rate, GDP, internet, and network interactions all lose individual significance, and the standalone ruling term is no longer distinguishable from zero. The visa-free, processing-time, conflict-deaths, and distance channels nonetheless retain their bivariate signs and remain statistically significant even jointly, and the processing-time interaction in fact strengthens. We therefore read columns (1)–(8) as complementary proxies for a single underlying margin—low-cost migration with weak protection prospects—rather than as eight independent channels.

The recognition-rate result may itself be conditional. Low-recognition origins are disproportion-

Figure 6: Event study by recognition-rate group ($R_o \lessgtr 0.33$)




*Notes:* PPML event-study estimates from equation (7), estimated separately for origins with $R_o > 0.33$ and $R_o < 0.33$. Among the origins with a defined Germany-specific recognition rate (those with at least one first-instance decision in Germany over 2010–H1 2012), 70 fall below the threshold and 10 above; the remaining origins in the estimation sample have no German decisions on which to compute $R_o$. Confidence intervals are based on standard errors clustered at the origin and destination country level. Shaded areas are 95 percent confidence intervals. The dashed vertical line marks the ruling date (July 2012). A per-origin coefficient histogram and an alternative functional-form specification are in Appendix E.1.

ately origins whose nationals can reach the Schengen area without a visa, so the group that responds most is also the group facing the lowest cost of acting on the change. Splitting the low-recognition group by visa status (Appendix Figure A18) shows the visa-free origins responding several times as strongly at the post-ruling peak, and doing so immediately, while the visa-required origins build more slowly to a smaller effect that is nonetheless clearly present. Read as a second-order interaction, this suggests that weak protection prospects and low mobility costs are complements rather than substitutes: the ruling raises the value of applying in Germany, and the origins that can act on it fastest are those whose nationals do not first have to obtain a visa. We do not push that reading far, because the complementary check is unavailable. Of the 10 high-recognition origins only 1 is visa-free, and it contributes a negligible number of applications, so whether the same visa gradient appears where protection prospects are strong cannot be established from this variation.

Table 4 re-estimates the ruling effect separately for each origin region—one difference-in-differences per region, run on that region's countries alone rather than as interactions within a single model. We split the World Bank's Europe & Central Asia and Sub-Saharan Africa aggregates into finer subregions, since each pools origins with very different migration profiles, and pool where the reverse problem arises: Eastern and Southern Africa share a column because Southern Africa alone has only five origin countries, and South Asia joins East Asia & Pacific for the same reason, each being too sparse on its own for the destination × origin fixed-effect structure (a handful of all-zero cells drive the Poisson estimator to a degenerate fit). The Western Balkans and Russia show the largest effects within Europe & Central Asia, and West & Central Africa the largest within Sub-Saharan Africa,

while the Middle East and North Africa—the most conflict-affected region in this window—is small and indistinguishable from zero, consistent with the recognition-rate heterogeneity.

Table 4: Heterogeneity by origin region

| *Panel A* | Russia | Western Balkans | Rest of Europe & Central Asia | Middle East & North Africa |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| DE × $\mathbb{1}(t > t_0)$ | 1.397*** | 0.968*** | 0.544*** | 0.079 |
| | (0.335) | (0.276) | (0.148) | (0.085) |
| *p-values* | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [0.351] |
| Conformal | [0.087] | [0.022] | [0.022] | [0.130] |
| Pre-period applicants (DE) | 1,705 | 7,885 | 3,285 | 13,445 |
| *N* | 545 | 1,713 | 3,849 | 4,771 |
| *Panel B* | West & Central Africa | East & Southern Africa | Asia & Pacific | Latin America & Caribbean |
| | (5) | (6) | (7) | (8) |
| DE × $\mathbb{1}(t > t_0)$ | 0.801*** | 0.356*** | 0.131** | 0.025 |
| | (0.153) | (0.087) | (0.061) | (0.201) |
| *p-values* | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [0.032] | [0.901] |
| Conformal | [0.022] | [0.391] | [0.217] | [–] |
| Pre-period applicants (DE) | 2,710 | 2,255 | 12,350 | 110 |
| *N* | 5,937 | 3,187 | 3,827 | 1,419 |

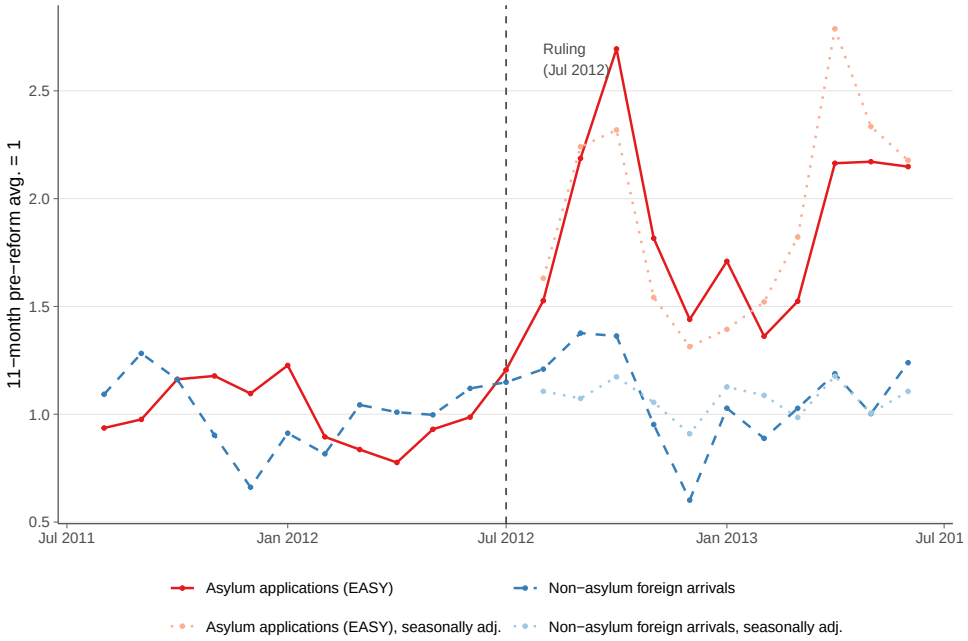
*Notes:* Each column is a separate PPML difference-in-differences estimate of equation (6), run on that origin region's countries only, with destination × origin and origin × year-month fixed effects. Standard errors (in parentheses) are clustered at the origin and destination country level, except for Russia, a single-origin column where clustering is on destination alone; *N* is the regression's number of observations. "Pre-period applicants (DE)" reports, for that region's origins, total first-time asylum applications to Germany over the same 11-month pre-ruling window as the main specification (August 2011–June 2012). Regions follow the World Bank classification, with Europe & Central Asia and Sub-Saharan Africa split further and Asia pooled: *Russia* is Russia alone; *Western Balkans* is Albania, Bosnia and Herzegovina, Kosovo, Montenegro, North Macedonia, and Serbia; *rest of Europe & Central Asia* is the remaining origins in the region (e.g. Ukraine, Turkey, the South Caucasus and Central Asian republics); *West & Central Africa* covers West Africa (e.g. Nigeria, Senegal, Mali) and Central Africa (e.g. Cameroon, Chad, DR Congo); *East & Southern Africa* pools East Africa and the Horn (e.g. Ethiopia, Somalia, Sudan) with Southern Africa (Botswana, Eswatini, Lesotho, Namibia, South Africa); *Asia & Pacific* pools World Bank South Asia with East Asia & Pacific, each of which is thin on its own for this fixed-effect structure, and excludes the Middle East and North Africa, which has its own column. *Inference.* The bracketed rows give *p*-values for the same coefficient under the two procedures of Appendix C, as in Table 1. *Cluster-robust* is the value implied by the standard error above it. *Conformal* is the moving-block test of Chernozhukov, Wüthrich and Zhu (2021) on Germany's monthly log gap for that region's origins, run on the window extended back to August 2009, so $T = 46$ and the smallest attainable value is 0.022. Permutation inference across destinations is estimated for every column but reported in Appendix C.3 rather than here, for the reasons given in Section 4.5. Stars refer to the cluster-robust *p*-value: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

## 5.2 Did Germany Get More Attractive for Other Migrants?

A natural concern is that the surge reflects a general improvement in Germany's attractiveness during the euro-area crisis rather than the ruling itself. Figure 7 plots asylum applications to Germany—from the EASY reception-center registration system, which tracks the true intake even when application processing lags behind (Appendix G)—against non-asylum foreign immigration of *Ausländer* from the Destatis *Wanderungsstatistik*. Both series are normalized to their 11-month pre-reform mean; the lighter, dotted lines add a seasonally-adjusted version of each series (the ratio to the same calendar month one year earlier), shown only for post-reform months, so that a generic seasonal pattern common

to both years cannot be mistaken for a ruling effect. Non-asylum immigration is broadly flat and dips marginally in late 2012, while asylum applications roughly double by late 2013—inconsistent with a generic attractiveness shock and consistent with the ruling raising the expected value of the asylum channel specifically. A sharper benchmark would compare non-asylum immigration to Germany and to other European destinations by origin and visa type, within the same gravity design. Such data are not available at the required frequency: Eurostat reports first residence permits by citizenship and reason only at the annual level, which cannot separate the months before and after a ruling issued in July.

Figure 7: Asylum applications and non-asylum foreign immigration to Germany




*Notes:* Monthly asylum applications to Germany (EASY registrations; red) and non-asylum foreign immigration of *Ausländer* (Destatis Wanderungsstatistik total minus Eurostat asylum applications; blue), both normalized to their 11-month pre-reform (Aug. 2011–Jun. 2012) average (solid/dashed lines). The lighter dotted lines add, for each series, the same normalization applied to its seasonally-adjusted value—its own value divided by the value in the same calendar month one year earlier—shown for post-reform months only. Asylum seekers are subject to Germany's residence registration requirement and are consequently already included, though not separately identified, in the Wanderungsstatistik total; subtracting Eurostat applications from that total isolates the non-asylum component (Statistisches Bundesamt , Destatis). The vertical dashed line marks July 2012.

### 5.3 Robustness

Appendix G collects a battery of robustness exercises that consistently confirm the main result; we summarize them here and point to the corresponding exhibits. Using total rather than first-time applications, which additionally lets us add Austria, leaves the effect unchanged (Table A10), as do more demanding fixed-effect structures and a destination × quarter unemployment control (Table A11). Collapsing the panel step by step to Germany versus the combined rest of the EU+ leaves the coefficient essentially unchanged (Table A12), the property that underlies the bias correction of Section 7. Replacing formal applications with the EASY reception-center registration counts, in a collapsed destination × month event study, leaves the jump essentially unchanged (Figure A21). Estimating applications per 10,000 inhabitants by linear DiD, synthetic control, and synthetic DiD yields effects of similar magnitude (Table A13), as does the population-weighted rates event study (Figure A23,

panel b). The estimate is robust to alternative pre/post windows of 6, 11, and 18 months (Table A16), to placebo rulings in other years and a comparison to the placebo years using July 2009–2011 cohorts (Appendix C.2), and to leaving out any single destination or origin (Figure A26).

The placebo ruling dates deserve a brief comment. Assigning a fake July ruling date to other years and re-running the event study (Appendix C.2) does yield several statistically significant effects. This is not surprising: in several of these windows Germany's attractiveness relative to other EU+ destinations shifted for reasons unrelated to the 2012 ruling—most notably the suspension of Dublin returns to Germany during the 2015 refugee crisis. Two features nonetheless set the true ruling apart. First, most of these placebo jumps do not line up with July of the respective year, whereas the 2012 response breaks sharply at the ruling date. Second, and more tellingly, the July 2012 effect is extreme in magnitude relative to the placebo distribution once applications are split by origin recognition rate. A generic shock to Germany's attractiveness—seasonal or crisis-driven, as in the placebo windows—raises applications from all origins roughly in proportion, whereas only a shock to asylum seeker benefits selectively pulls in applicants from low-recognition origins with the weakest protection prospects. No placebo date reproduces this strong low-recognition response, which is the signature of the ruling (Figure A11).

## 6 Applicant Composition

The strong heterogeneity over recognition implies that most marginal applicants come from origins whose claims are typically rejected, suggesting that the total number of refugees admitted should increase far less strongly. In this section, we study this directly using the first-instance asylum decisions data described in Section 3.2.

**Refugee decisions.** One way to translate the application response into refugee grants is to scale it by each origin's recognition rate while holding within-origin recognition rates fixed (the fixed-composition calculation of Appendix E.1). If marginal applicants from a given origin are weaker claimants than the inframarginal pool, however, this calculation overstates the refugee response. We can sidestep the assumption by estimating directly on Eurostat's quarterly first-instance decisions (`migr_asydcfstq`).

We estimate the same Poisson DiD design as in Section 5 on the decisions sample described in Section 3.2, with country-level quarterly decision counts as the outcome (Germany versus the EU controls, 2011 Q1–2013 Q4, ruling quarter excluded, reference 2012 Q2). Two features of these data warrant caution *ex ante*, before any pre-trend is inspected. First, decision counts are an administratively determined outcome: they reflect the agency's processing capacity, backlog, and case prioritisation as much as the underlying inflow, and the ruling-induced cohort reaches a first-instance decision only after a procedure lasting roughly seven to fifteen months. The decision margin should therefore respond with a lag and more noisily than the application margin. Second, decision counts are sensitive to batch-processing episodes. We accordingly read the decisions evidence as corroborating the heterogeneity result rather than as a primary estimate.

Two omissions follow directly from this sensitivity, and we make them before inspecting any estimates. We drop *Italy* from the control group: Italian first-instance decisions surged over 2011–2012 as the country cleared the backlog of the 2011 North-African (Arab Spring) arrivals, a processing spike unrelated to the German ruling that shows up as a large, isolated peak in the raw decision

counts. We drop *Syrian, stateless, and unknown* origins: German positive decisions for Syrians jumped in 2012 Q2—precisely the event-study reference quarter—when the Federal Office began granting protection to Syrian applicants en masse as the civil war escalated, which would mechanically distort the reference period. Both episodes are driven by administrative or conflict timing rather than by benefit levels. The decisions sample is therefore Germany versus the remaining EU24 destinations (Italy excluded), over origins excluding Syria, stateless, and unknown; the total-applications and positive-decisions columns of Table 5 use this same sample.

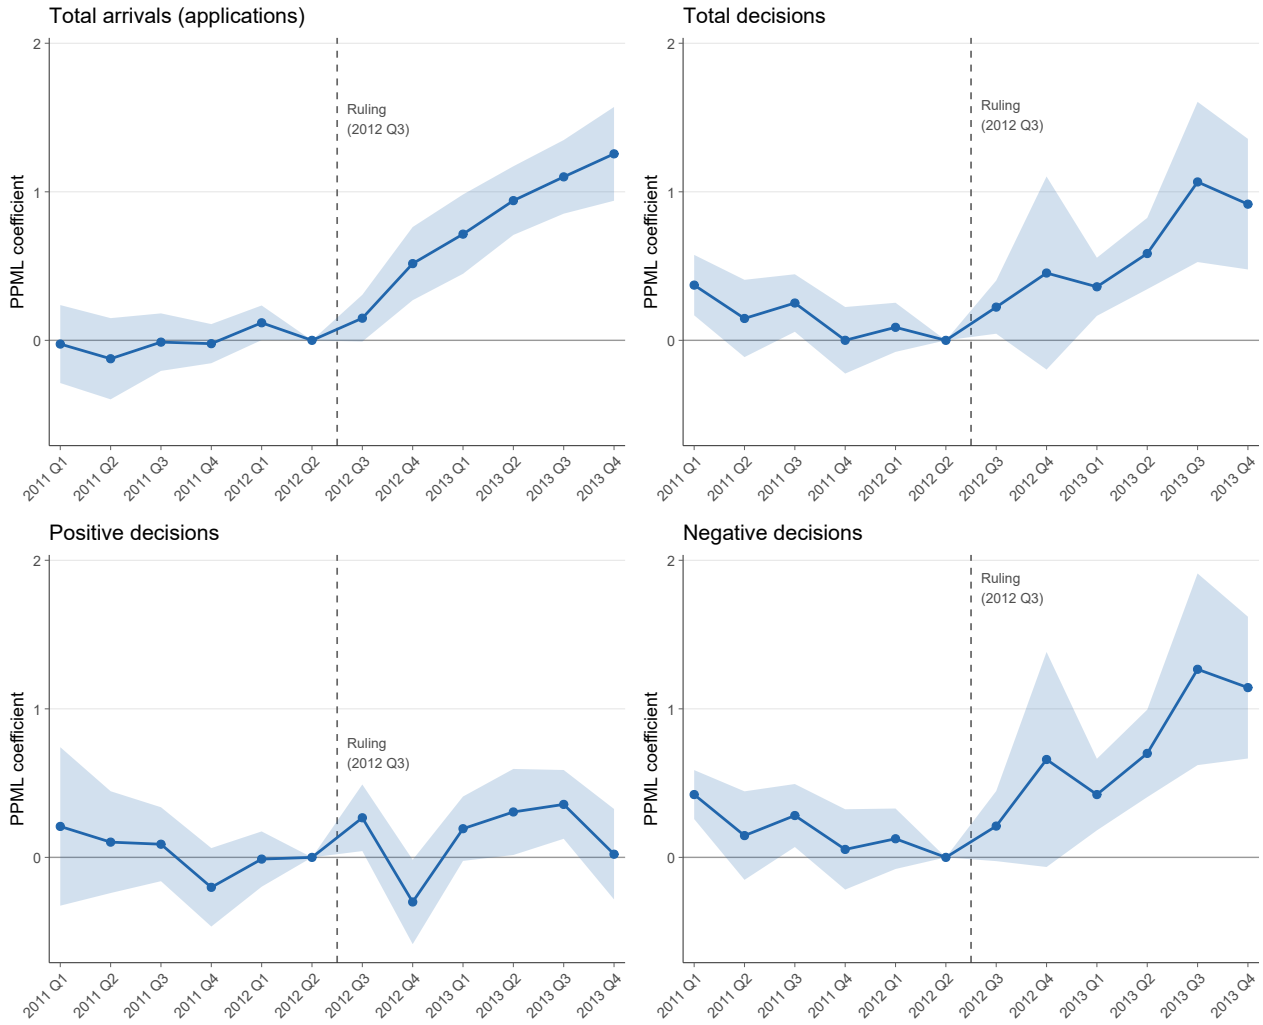
Figure 8 reports event-study coefficients on the origin × destination × quarter panel that is the quarterly counterpart of the main applications specification in equation (7), for four outcomes: total arrivals—i.e. first-time asylum *applications* (Eurostat `migr_asyappctzm`), the same series used throughout the rest of the paper, not EASY registrations—and total first-instance decisions in the top row, positive and negative first-instance decisions in the bottom row. Arrivals and total decisions jump together at the ruling, but the decision-level split shows why the recognized-refugee response is so much weaker: positive decisions, the quantity of primary interest, rise only modestly after the ruling—an increase of about 13 percent (Table 5), far below the 61 percent jump in applications—and show no visible pre-trend, while negative decisions rise by more, consistent with the marginal applicants skewing toward low-recognition origins. The implied refugee-benefit elasticity is therefore about 0.24, against 1.0–1.1for applications, consistent with the heterogeneity finding that marginal applicants came predominantly from low-recognition origins. The muted overall response is robust across estimation windows and aggregation levels, including the destination × time aggregation.

Table 5 places the two responses side by side, each as a single-coefficient quarterly DiD estimated at ±4 and ±6 quarters around the ruling. Applications rise sharply on every window, while positive decisions rise far less, compressing the event-study contrast into one number: applications jump at the ruling while the number of recognized refugees increases only modestly. All columns are estimated on the same origins and destinations, so the gap is not an artifact of differing coverage.

Two timing effects could nonetheless bias the decision response downward. Because a first-instance decision takes roughly seven to fifteen months, decisions in the first post-ruling quarters mostly concern applicants who arrived before the ruling, so a short window captures only part of the ruling-induced cohort. In addition, the BAMF fast-track for Western Balkan claims from October 2012 (Section 4.4) brought forward rejections for the group that responded most, so negative decisions on the marginal cohort fall inside the window while positive decisions on more slowly processed claims may fall outside it. The ±6-quarter window extends the post-period to the end of 2013, by which time a large part of the cohort that arrived in the first months after the ruling should have received a first-instance decision, which largely addresses this concern. On that window, too, positive decisions rise far less than applications.

The heterogeneity analysis is silent about how marginal applicants differ from inframarginal applicants *from the same origin*. A complementary calculation fills this gap by asking how much the number of recognized refugees would change if the composition of applicants within each origin were unchanged. Applying the estimated recognition-rate gradient of the application response to the application-weighted distribution of origin recognition rates, and holding within-origin recognition rates fixed, implies a predicted increase of 26 percent in recognized refugees against an aggregate application increase of 63 percent, an implied refugee-benefit elasticity of about 0.48 (Appendix E.1, equation A11 and Figure A16); the prediction is insensitive to the heterogeneity specification used.

Figure 8: Event study: arrivals and asylum decisions in Germany, origin × destination × time specification




*Notes:* Quarterly PPML event study on the origin × destination × quarter panel, with destination × origin and origin × year-quarter fixed effects and standard errors clustered at the origin and destination country level; this is the quarterly counterpart of the main applications specification in equation (7). Top-left: total arrivals, i.e. first-time asylum *applications* (Eurostat `migr_asyappctzm`)—the same application series used throughout the paper, not EASY registrations. Top-right: total first-instance decisions. Bottom-left: positive first-instance decisions. Bottom-right: negative first-instance decisions. Decisions are from Eurostat `migr_asydcfstq` and are published only at quarterly frequency. Germany is the treated unit; the controls are the EU24 destinations excluding Italy, and Syrian, stateless, and unknown origins are excluded (see Section 6). The reference quarter is 2012 Q2 and the dashed line marks the ruling quarter (2012 Q3); shaded bands are 95 percent confidence intervals; all four panels share the same $y$-axis scale.

The first-instance decisions in Table 5 show an even weaker grant response than this fixed-composition benchmark. The marginal applicants attracted by the ruling were therefore less likely to obtain a status than the inframarginal applicants who would have arrived in its absence, even those that originate from the same country.

Table 5: Applications versus positive decisions: difference-in-differences

| | Total applications | | Positive decisions | | Negative decisions | |
|---|---|---|---|---|---|---|
| | ±4 q. (1) | ±6 q. (2) | ±4 q. (3) | ±6 q. (4) | ±4 q. (5) | ±6 q. (6) |
| DE × $\mathbb{1}(t > t_0)$ | 0.570 | 0.821 | 0.170 | 0.119 | 0.428 | 0.619 |
| | (0.110) | (0.138) | (0.044) | (0.063) | (0.168) | (0.212) |
| *p-values* | | | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [0.057] | [0.011] | [0.004] |
| Conformal | [0.056] | [0.050] | [0.111] | [0.050] | [0.111] | [0.100] |
| *N* | 6,263 | 10,193 | 3,307 | 5,550 | 6,007 | 9,680 |
| Origins | 162 | 162 | 162 | 162 | 162 | 162 |
| Destinations | 19 | 19 | 19 | 19 | 19 | 19 |

*Notes:* Columns (1)–(6): PPML single-coefficient DiD (Germany × post) of equation (6) for total first-time applications, positive first-instance decisions, and negative first-instance decisions, at ±4 and ±6 quarters around the ruling, on the origin × destination × quarter panel with destination × origin and origin × quarter fixed effects and two-way clustering on destination and origin. The reform quarter (2012 Q3) is kept in the post period. All columns use the same sample (Section 6): EU24 destinations excluding Italy and the origins common to the application (`migr_asyappctzm`) and decision (`migr_asydcfstq`) data, with Syria, stateless, and unknown origins excluded. Small differences in the per-column *N* reflect singleton drops under the fixed effects. *Inference.* The bracketed rows give *p*-values for the reported coefficient under two procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it, and the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point. This panel is quarterly, so the conformal test rotates quarters and reaches back to 2009Q1 rather than to the column's own window, which would leave too few periods to rotate; $T$ is 18 for the ±4-quarter columns and 20 for the ±6-quarter columns, giving floors of 0.056 and 0.050. Permutation inference across destinations is estimated for every column but reported in Appendix C.3 rather than here, for the reasons given in Section 4.5.

# 7 Diversion or Attraction?

The DiD coefficient $\hat{\beta} = 0.476$ identifies the effect on applications in Germany *relative to* the rest of the EU+ (a 61% rise), but does not by itself distinguish two mechanisms with very different economic content. Under *attraction*, higher German benefits draw in wholly new EU entrants and shift the EU-wide entry decision, generating an externality at the EU's external frontier. Under *diversion*, applicants who would otherwise have filed in another EU country reroute to Germany, redistributing flows within the EU+ and imposing a fiscal externality on partner member states. The diversion share $\rho$ derived in Section 4.1 pins down the split.

This section delivers the empirical magnitudes. The Eurodac decomposition gives a lower bound $\hat{\rho}_{D1} = 0.56$ (Section 7.1), and the stable-origin comparison indicates that the true diversion share lies above that bound (Section 7.2), with two large attraction-prone groups, the Western Balkans and the Russian Chechens, qualifying the latter (Section 7.3). Plugging these into the bias correction (A9) over the plausible range $\rho_{D2} \in [0, 0.44]$ yields a counterfactual effect on applications in Germany of 53–61% (application–benefit elasticity 1.0–1.1): diversion is the dominant margin, but attraction is not zero, so the EU-wide total inflow effect is strictly positive, consistent with the flat external-frontier border-crossing detections after the ruling (Appendix B.4).

## 7.1 Onward Migration

The Eurodac system records the fingerprints of all asylum applicants in most EU+ countries, as discussed in Section 3.3. A Category 1–Category 1 cross-border hit occurs when an applicant whose prints were stored under country $s$ at some prior date appears in receiving country $j$'s query and

matches. Each first-time asylum application in Germany in $t$ decomposes as

$$M_{\mathrm{DE},t} \;=\; N_{\mathrm{DE},t} \;+\; \sum_{s\neq \mathrm{DE}} O^{s}_{\mathrm{DE},t}, \tag{9}$$

where $N_{\mathrm{DE},t}$ counts applicants who never previously applied for asylum elsewhere in the EU+ (residual against the Eurostat total) and $O^{s}_{\mathrm{DE},t}$ counts those whose fingerprint matched a stored asylum record in country $s$.

Marginal applicants of type D1 increase $\sum_{s\neq \mathrm{DE}} O^{s}_{\mathrm{DE},t}$; the relative increase in this sum can be estimated by a gravity model at the receiving country-time level with receiving country and time fixed effects, with treatment applying to Germany after July 2012. Types A and D2 are observationally indistinguishable. However, our main estimate $\hat{\beta}$ provides an estimate of the relative increase in the total number of asylum applications in Germany. Those estimates can be used to provide estimates of the total increase in $\widehat{\Delta O}_{\mathrm{DE}}$ and $\widehat{\Delta M}_{\mathrm{DE}}$, while $\widehat{\Delta N}_{\mathrm{DE}}$ is estimated directly by applying the same PPML specification to the never-applied subsample.[21] Those estimates enable estimation of the share of all marginal applicants that are of type D1 and provide a *lower bound* on the diversion share $\rho$:

$$\hat{\rho} \;\geq\; \frac{\widehat{\Delta O}_{\mathrm{DE}}}{\widehat{\Delta M}_{\mathrm{DE}}} \;\equiv\; \hat{\rho}_{D1}, \tag{10}$$

where $\widehat{\Delta M}_{\mathrm{DE}}$ is the increase in total applications implied by the headline DiD estimate $\hat{\beta}$ (equivalently $\widehat{\Delta N}_{\mathrm{DE}} + \widehat{\Delta O}_{\mathrm{DE}}$ by the decomposition (9)). This object is a lower bound rather than an unbiased estimate of $\rho$: cross-border Eurodac hits capture only diverted applicants who had *already* registered in another member state, so $\hat{\rho}_{D1}$ misses unobserved diversion (type D2) entirely and is further attenuated by any matches lost to fingerprint-coverage gaps or reporting lags. Both forces push in the same direction, so $\hat{\rho}_{D1}$ understates the true diversion share. We estimate the two components below.

We now report the underlying event-study evidence and decompose the inward-hit response by country of first application. Estimating the main PPML (6) separately on never-applied applicants and on inward Category 1–Category 1 hits yields $\hat{\beta}_N = 0.302$ and $\hat{\beta}_O = 0.938$, both highly significant; Table 6 collects these coefficients with the implied level changes and the lower-bound diversion share. Figure 9 (a) plots the corresponding event studies: both series rise after the ruling, but the proportional rise in inward hits is roughly three times larger. The headline applications DiD implies an increase over the 11-month post-ruling period of $\widehat{\Delta M}_{\mathrm{DE}} \approx 27,262$, of which the inward Category 1–Category 1 response accounts for $\widehat{\Delta O}_{\mathrm{DE}} \approx 15,180$, giving $\hat{\rho}_{D1} = \widehat{\Delta O}/\widehat{\Delta M} = 0.56$. Applicants who never applied elsewhere are estimated directly from the same PPML specification applied to that subsample, $\widehat{\Delta N}_{\mathrm{DE}} \approx 12,343$ ($\hat{\beta}_N = 0.302$). Because that group bundles genuine attraction $A$ with the unobserved diversion $D2$, $\hat{\rho}_{D1}$ is a lower bound on the total diversion share $\rho$, and equivalently $1-\hat{\rho}_{D1}$ is an upper bound on the attraction share.

**Diversion share over event time.** Figure A19 (Appendix F) plots the per-period diversion share $\hat{\rho}_k = \Delta O_k/(\Delta O_k + \Delta N_k)$. The pooled estimate $\hat{\rho} = 0.56$ is reached within two months of the ruling and stays broadly stable thereafter, consistent with quick learning by both applicants and intermediaries. The high and stable diversion share fits the design of the German system, in which applicants with

[21] Absolute increases can be calculated using the pre-ruling level and the PPML estimate: $\widehat{\Delta X}_{\mathrm{DE}} = X_{0,\mathrm{DE}}\left(e^{\hat{x}} - 1\right)$.

Table 6: Decomposing the increase in applications in Germany by Eurodac status

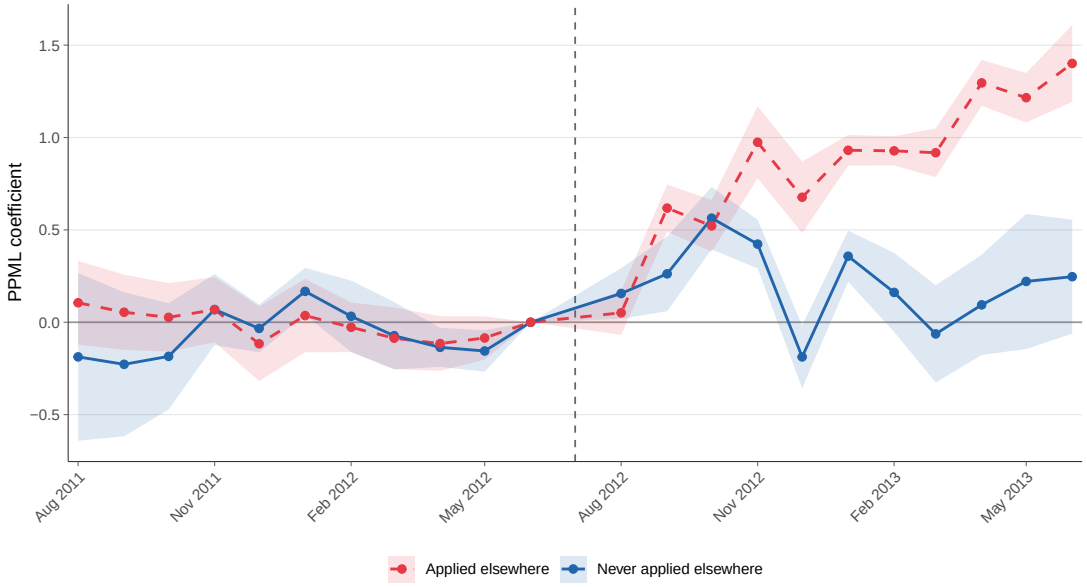
| | DiD coefficient $\hat{\beta}$ (s.e.) (1) | Pre-ruling base (11-month total) (2) | Implied increase $\widehat{\Delta}$ (over 11 months) (3) |
|---|---|---|---|
| Total first-time applications ($\Delta M$, baseline DiD) | 0.476 (0.102) | 44,720 | 27,262 |
| Inward Cat. 1–1 hits ($\Delta O = D1$) | 0.938 (0.064) | 9,763 | 15,180 |
| Implied diversion share $\hat{\rho}_{D1} = \widehat{\Delta O}/\widehat{\Delta M}$ | | | 0.56 (0.16) |
| *Components of $\widehat{\Delta O}$, by sender group* | | | |
| Arrival | 1.143 (0.085) | 3,537 | 7,550 |
| Transit | 0.772 (0.067) | 2,363 | 2,752 |
| Destination | 0.807 (0.054) | 3,863 | 4,793 |

*Notes:* PPML DiD (Germany × post) of equation (6) over the main DiD window. The first row is the headline total first-time applications response ($\widehat{\Delta M}$); the second row is the response of inward cross-border Category 1–Category 1 hits ($O = \sum_{s \neq \text{DE}} O^s_{\text{DE},t} = D1$, observed onward migration). Standard errors are clustered at the origin and destination country level. Implied increases use $\widehat{\Delta X} = X_0\,(e^{\hat{\beta}} - 1)$, with $X_0$ the total over the 11-month pre-ruling period, so that $\widehat{\Delta X}$ is the increase over the 11-month post-ruling period. The lower-bound diversion share is $\hat{\rho}_{D1} = \widehat{\Delta O}/\widehat{\Delta M}$ (its standard error is computed by the delta method); $\widehat{\Delta N} \approx 12,343$, the increase in applicants who never applied elsewhere, is estimated directly from the same PPML specification applied to the never-applied subsample (not the arithmetic residual $\widehat{\Delta M} - \widehat{\Delta O}$, which differs slightly because it combines two separately estimated coefficients) and bundles genuine attraction $A$ with unobserved diversion $D2$. The bottom panel decomposes $\widehat{\Delta O}$ by the sender group of the prior application (Arrival, Transit, Destination; see Figure 9 and its note for the exact country lists); the three rows sum to $\widehat{\Delta O}$ up to rounding.

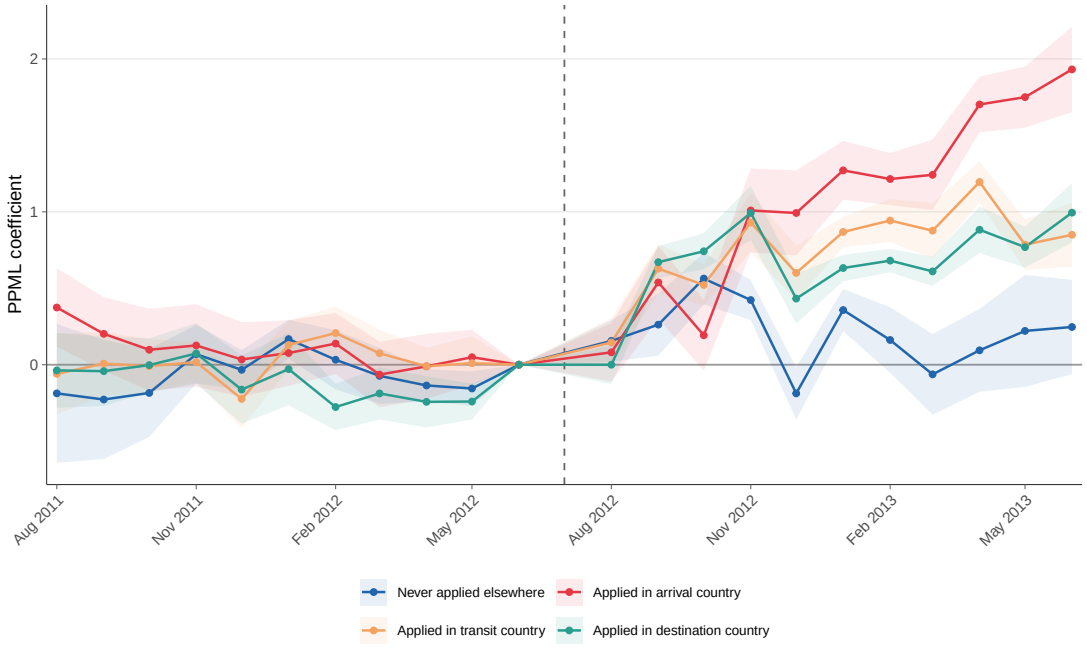
weak protection prospects can draw the raised benefits throughout a lengthy procedure irrespective of where they first applied, so that rerouting within the EU+ is privately attractive even for those who would otherwise have remained in another member state. The headline $\hat{\rho}$ is robust to alternative treatments of country-period cells where the residual $N_{\text{DE}}$ would be mechanically negative: $\hat{\rho} = 0.55$ (baseline), 0.55 (drop cells), 0.55 (drop countries); see Appendix F.

**Sender-group decomposition.** Figure 9 (b) splits inward hits by the country of first application into three economically distinct groups. *Arrival* comprises the EU's southern and southeastern external-border countries (Greece, Spain, Italy, Bulgaria, Romania, Hungary, Slovakia, Slovenia), through which Mediterranean and Western Balkan routes pass, together with Poland—which hosted a large 2008–13 Chechen first-application population that is the dominant feeder for German Russian-origin applications—and Cyprus and Malta, external-border island entry points; an inward hit from this group typically represents an applicant who first applied at the external border and subsequently moved north. *Transit* countries (Czech Republic, Austria, France, Switzerland) lie geographically between origins and Germany; an inward hit from this group represents an applicant who began the procedure in a transit country but moved on to Germany. *Destination* captures secondary movement from every remaining EU+ country, all of which are themselves plausible end-destinations (Sweden, Netherlands, UK, Belgium, and the rest of northern/western Europe). The Poisson DiD coefficients (Appendix Table A9) are positive and significant for all three groups, with the largest *level* contributions from the Arrival and Destination groups; Appendix Figure A20 further splits each group into finer sub-regions, isolating Poland's disproportionately large *proportional* increase within Arrival. The Polish sub-channel is consistent with the Chechen attraction pattern documented anecdotally below; the Arrival channel is consistent with stronger onward migration from Mediterranean entry points; the Destination channel suggests that the ruling also pulled previously settled-elsewhere asylum seekers from northern Europe.

Figure 9: Event study: monthly applications in Germany by Eurodac status of the applicant




(a) Never applied vs. inward Category 1–Category 1 hits




(b) Never applied vs. three sender groups

*Notes:* PPML event studies on monthly applications in Germany by Eurodac status. Specifications and clustering as in Section 5; the omitted period is June 2012. Shaded bands are 95 percent confidence intervals. Panel (b) sender groups: *Arrival* = Greece, Spain, Italy, Bulgaria, Romania, Hungary, Slovakia, Slovenia, Poland, Cyprus, Malta; *Transit* = Czech Republic, Austria, France, Switzerland; *Destination* = every other EU+ country in the sample. Appendix Figure A20 splits each of these three groups into finer sub-regions.

## 7.2 Stable Origins

The stable-origins comparison speaks to attraction and diversion at once rather than to diversion alone: it asks how much of the ruling-induced increase in Germany shows up as a decline in the rest of the EU+, and how much as a genuine addition to EU-wide applications, among origins whose push factors did not themselves change. It is therefore a second, independent line of evidence on $\rho$, complementing

the Eurodac lower bound of Section 7.1. We restrict attention to origins with stable push factors, meaning no large conflict, growth, or political shocks during 2010–2013, and compare the Germany-specific DiD to the total-EU DiD on a matched origin set. The identifying logic is simple: for origins whose push factors were stable, total EU asylum applications should have remained unchanged absent a financial pull from Germany. If total EU applications from these origins rose after the July 2012 ruling, the increase reflects newly induced EU entrants ($\rho < 1$, attraction); if they stayed flat while applications in Germany rose, the increase is pure diversion ($\rho = 1$).

We estimate the main specification (6) on the stable-origin subsample alongside a total-EU regression of applications net of Eurodac repeat filings,

$$M_{ot}^{\text{EU,net}} = \exp\big(\beta_{\text{EU}}\,\mathbb{1}(t > t_0) + \phi_o + \gamma_{\tau(t)}\big) \cdot \varepsilon_{ot}, \tag{11}$$

with origin and calendar-month-of-year fixed effects, restricted to Eurodac-covered destination-months; the Germany-specific regression is estimated on the same covered destinations, and $\hat{\rho}$ follows from comparing the two (Appendix D.4).

Table 7: Diversion versus attraction: stable-origin evidence

| | (1) All origins | (2) Conflict only | (3) + GDP | (4) + VA | (5) + CC | (6) All stable |
|---|---|---|---|---|---|---|
| *Panel A: conflict, GDP, and governance stability* | | | | | | |
| Total applications (net of repeats) | 0.170* | 0.222* | 0.058 | 0.026 | 0.006 | −0.034 |
| | (0.091) | (0.121) | (0.048) | (0.049) | (0.048) | (0.047) |
| Germany applications | 0.476*** | 0.964*** | 0.691*** | 0.582*** | 0.577*** | 0.474*** |
| | (0.102) | (0.111) | (0.145) | (0.049) | (0.059) | (0.064) |
| $N$ origins | 168 | 122 | 102 | 50 | 53 | 26 |

| | (1) All origins | (2) $\lvert\Delta\rvert < 5\%$ | (3) $\lvert\Delta\rvert < 10\%$ | (4) $\lvert\Delta\rvert < 20\%$ | (5) $\lvert\Delta\rvert < 25\%$ |
|---|---|---|---|---|---|
| *Panel B: OECD outmigration stability, non-European destinations* | | | | | |
| Total applications (net of repeats) | 0.170* | −0.051 | −0.016 | 0.012 | −0.001 |
| | (0.091) | (0.060) | (0.059) | (0.044) | (0.041) |
| Applications in Germany | 0.476*** | 0.495*** | 0.454*** | 0.327*** | 0.355*** |
| | (0.102) | (0.015) | (0.075) | (0.068) | (0.066) |
| $N$ origins | 168 | 29 | 53 | 93 | 104 |

*Notes:* PPML on a matched origin set, restricted throughout to destination-country-months covered by Eurodac reporting. Panel A, col. 1: full main-DiD set; cols. 2–5 cumulatively impose: max UCDP battle deaths $\leq 25$/yr over 2010–2013; PPP-GDP-pc growth in $(-2\%, +10\%)$; SD of WGI Voice & Accountability $< 0.10$; SD of WGI Control of Corruption $< 0.10$. Panel B classifies origins by the absolute change in their total inflow to non-European OECD destinations over 2011–2013 (OECD International Migration Database, *Inflows of foreign population by nationality*), with col. 1 the full main-DiD set and cols. 2–5 progressively widening the outmigration-change band. Non-European destinations only, so that intra-European flows that might themselves respond to the ruling do not enter the classification. "Total applications" nets Eurostat first-time applications against Eurodac Category 1–1 inward hits, destination-by-destination, before summing across destinations, so that the total responds to new entrants rather than double-counting repeat filers (Appendix D.4); as Eurodac hits are not observed by applicant nationality, the destination-month repeat-application share is applied uniformly across origins within that cell. "Germany applications" uses the main DiD spec, estimated on the identical set of Eurodac-covered destination countries used to build the corresponding total-EU row, so each panel rests on one common destination panel. Implied diversion shares are discussed qualitatively in the text (Section 7.2) rather than tabulated: the point estimates in the tighter subsamples are individually noisy, and the qualitative pattern—$\hat{\beta}_{EU}$ falling toward zero or negative as the stability criteria tighten, while $\hat{\beta}_{DE}$ stays large throughout—is what supports the diversion reading, not any single ratio. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Panel A of Table 7 applies this across five stability subsamples. The Germany-specific DiD exceeds the net-total-EU DiD in every one of them, but it is the cumulative pattern rather than any single column that carries the argument. As the stability requirements tighten, $\hat{\beta}_{EU}$ falls monotonically, from a positive and significant coefficient on the full origin set to a small and insignificant one, and eventually to a negative one on the strictest subsample, while $\hat{\beta}_{DE}$ stays large and significant throughout. The implied $\hat{\rho}$ therefore rises from below zero on the unrestricted set, where attraction among the Western Balkans and the Russian Chechens still inflates the net-total-EU response, through the interior of $[0, 1]$, and into the over-diversion region on the strictest set of 26 origins.

Panel B replaces those criteria with one that does not rest on conflict, growth, or governance data at all. Origins are classified by how much their total migration outflow to non-European OECD destinations moved over 2011–2013 (OECD International Migration Database), a measure dominated by general emigration pressure rather than by EU-specific conditions; non-European destinations are used so that intra-European flows, which might themselves react to the ruling, do not enter the classification. The picture is the same. Among the most outmigration-stable origins, net EU applications did not rise while applications in Germany rose sharply, and the pattern attenuates as the band widens to admit less stable origins, consistent with attraction operating where overall emigration is more elastic.

The exercise cannot pin down an exact diversion share. The standard errors are large, and because $\hat{\rho}$ is a ratio estimator in which the EU-wide flow being differenced is considerably larger than the flow to Germany, that imprecision passes straight into it. What the estimates do establish is that we detect no considerable attraction effect among the most stable origins. The ruling therefore most plausibly reduced asylum applications in the rest of the EU+ by a substantial margin.

### 7.3 Anecdotal Attraction: Western Balkans and Russian Chechens

The heterogeneity analysis identified two origin groups with particularly large effects for which attraction is plausible: the Western Balkans and Chechens from Russia. Both account for a disproportionate share of the headline response, and both display the diagnostic attraction pattern, a post-ruling rise in non-German EU applications where pure diversion would instead have produced a decline. Appendix E.1 documents the two cases in detail (Figure A15): for the Western Balkans, visa-free, near-zero-recognition economic migrants whose applications to the rest of the EU+ rose in parallel with those to Germany; for Russian Chechens, a rise in Germany that precedes a delayed rise elsewhere, combining attraction with onward migration through the Polish first-application channel.

These cases caution against reading the headline diversion bound as the whole story. Diversion is the dominant margin on average, but attraction is clearly present for some elastic-emigration origins, so the EU-wide total inflow effect is strictly positive. A large country raising benefits therefore does two things at once: it *redistributes* EU asylum flows toward itself, a fiscal externality on partner member states, and it *raises* EU-wide flows from elastic origins, an externality at the external frontier. Both argue for EU-level coordination of reception standards.

## 8 Discussion

We used Germany's 2012 constitutional court ruling—a large, court-mandated, plausibly exogenous increase in asylum seeker cash benefits after nearly two decades of nominal freeze—to study how cash

benefits shape the size, composition, and geographic allocation of asylum flows in Europe. The central finding of this paper is that asylum applications respond strongly to benefit levels, with a corrected application-benefit elasticity of 1.0–1.1, whereas the recognized-refugee response is far weaker, with an implied elasticity of about 0.24. These bounds follow from the diversion share bracketed in Section 7: the Eurodac lower bound $\hat{\rho}_{D1} = 0.56$ and the stable-origin evidence together place the unobserved diversion share at $\hat{\rho}_{D2} \leq 0.44$, so over the plausible range $\hat{\rho}_{D2} \in [0, 0.44]$ the policy-relevant effect on applications in Germany runs from 53% to 61%.

The Eurodac lower bound combines two estimates obtained from different designs. The increase in applications comes from the origin-nested gravity specification, whereas the increase in inward hits is estimated on a destination × month panel, because Eurodac does not record the applicant's nationality. If the two designs measured different objects, their ratio would misstate $\rho_{D1}$. Estimating the applications response on the same destination × month panel, without origin × time fixed effects, gives a nearly identical coefficient (Table A12, column 2), which suggests that the difference in design does not materially affect the bound.

The same range of diversion shares also pins down the dissimilarity parameter of the random-utility model in Section 4.1. Inverting equation (A7) at $\rho \in [0.56, 1]$ gives $\hat{\tau}$ between 0.00 (full diversion) and 0.39, far below the value of one at which tastes for different destinations would be independent. For the marginal applicant, EU+ destinations are therefore close substitutes, so a change in Germany's attractiveness mainly reallocates applicants within the EU+ rather than changing whether they apply in Europe at all. Table 8 places the elasticity in the broader literature, distinguishing asylum application elasticities (Panel A), total or net flow elasticities (Panel B), and internal location-choice elasticities (Panel C). Importantly, all other references concern social assistance paid to unemployed refugees, whereas this paper explicitly studied benefits that asylum seekers receive during their asylum procedure.

Table 8 also reports a naive PPML panel regression of monthly asylum applications on the log of the real benefit level across all EU+ destinations and years, with the same destination × origin and origin × time fixed effects as our main specification, identified from the cross-country and over-time variation in benefit levels rather than from a single reform (Appendix H). This naive elasticity is very large, 1.82, and statistically significant. It is, however, likely biased, and probably upward: benefit levels are not set in isolation but tend to move with a country's broader refugee-policy stance, so that more generous benefits coincide with more open recognition and reception policies generally. The cross-sectional elasticity therefore captures this whole policy bundle rather than the benefit level alone, which is precisely why the court-mandated, benefit-specific reform we exploit identifies a smaller effect. The analogous naive regression for recognized refugees yields a similarly large and significant elasticity (1.40), reinforcing that these correlations overstate the causal response.

Table 8: Migration elasticities w.r.t. benefits: summary of recent evidence

| Study | Country | Elasticity | Type | Benefit type | Method |
|---|---|---|---|---|---|
| *Panel A: Asylum applications* | | | | | |
| This paper | Germany | 1.0–1.1 | — | Asylum seeker benefits | DiD (court ruling) |
| Appendix H | EU | 1.82 | — | Asylum seeker benefits | Panel (non-causal) |
| Adema (2025) | Denmark | 0.8–1.1 | — | Social assistance | DiD (policy change) |
| *Panel B: Total (refugee) flows* | | | | | |
| This paper | Germany | $\approx 0.24$ | Inflow / recognized refugees | Asylum seeker benefits | DiD (court ruling) |
| Appendix H | EU | 1.40 | — | Asylum seeker benefits | Panel (non-causal) |
| Agersnap, Jensen and Kleven (2020) | Denmark | 1.3 | *Net* flows | Social assistance | DiD (policy change) |
| Adema (2025) | Denmark | $\leq 0.3$ | Stock and inflow | Social assistance | DiD (policy change) |
| Adema et al. (2025) | EU | 0.1–0.2 | Destination preference | Social assistance | Conjoint experiment |
| *Panel C: Internal location choice* | | | | | |
| Ferwerda, Marbach and Hangartner (2024) | Switzerland | small | Within-country location | Social assistance | DiD (cantonal) |
| Wett, Gallegos Torres and Steinmayr (2024) | Austria | small | Within-country location | Social assistance | DiD (states) |

*Notes:* Approximate elasticities of migration with respect to benefit levels. Elasticities are not directly comparable across studies due to differences in outcome (applications, net flows, stocks, location shares), population (asylum seekers, all migrants, refugees), and setting. The *Benefit type* column distinguishes asylum seeker benefits paid during the asylum procedure (this paper) from social assistance paid to recognized refugees or migrants (all other studies).

The table shows that asylum application elasticities in the case of Denmark and Germany are comparable in magnitude, although the Danish reforms were political decisions accompanied by other policy changes which likely caused an upward bias (Adema, 2025). This upward bias could have been counteracted by the response of asylum seekers with a low probability of receiving asylum, who likely anticipated that their chance of obtaining protection—and thus social assistance—would be small.

However, in both the German and Danish case application elasticities are larger than stock and gross inflow elasticities (Panel B). Estimates of internal mobility elasticities based on subnational and temporal variation in social assistance are also considerably smaller (Panel C). These patterns can be explained by the following. First, filing an asylum claim in another EU country is nearly costless in the presence of free travel and affordable bus or train tickets. The large extent of onward applications in the Eurodac data corroborates this. However, moving after several months or years of residence may come with considerable costs, explaining the limited internal mobility responses in the Swiss and Austrian cases. Second, our results show that the application response is concentrated among applicants with low recognition rates; their claims are overwhelmingly rejected, so the increase in applications translates into few additional *recognized* refugees—though, because deportation enforcement is only partial, the number of rejected-but-tolerated applicants who remain in Germany may still have risen (Section 6). This response is in line with the fact that procedure lengths are longer for those with ex-ante less plausible applications and the fact that a considerable fraction of rejected applicants continue to receive benefits.

## 9 Conclusion

These estimates are specific to the German system we study, and reading them onto other settings requires some care. Germany was comparatively generous to rejected-but-non-deportable applicants during our sample period: benefits continued throughout what could be a lengthy procedure, with no hard cutoff pushing claimants toward departure, and rejected applicants who could not be deported eventually moved onto the regular social-assistance (SGB II) rate after 48 months rather than facing a reduced or in-kind-only regime (Section 2.2). Countries whose reception systems more aggressively curtail support for rejected claimants—for example by shifting them onto departure-oriented, largely in-kind assistance, as in the Netherlands or Austria—offer a smaller discounted benefit gain to the same low-recognition-probability applicant, so our estimated elasticity should be read as informative about a system with the German institutional features specifically, plausibly an upper bound on the response in less generous ones. That an applicant deterred by an initial low-recognition-probability outlook reroutes to the destination with the most generous benefits, as our diversion evidence suggests happens within the EU+, further makes the aggregate effect of any one country's benefit level a function of the entire reception-system landscape, not that country's policy in isolation.

The policy implications are direct. First, asylum seeker benefit levels matter for the number of asylum claims received, but primarily from applicants who will ultimately be rejected—generating an externality on other EU+ destinations. This could lead to a situation where countries compete to lower benefits, which may not meaningfully reduce the total number of asylum applications in the EU+, as witnessed by the large estimated diversion share, while reducing living standards for asylum applicants and worsening the integration of refugees (Dustmann, Landersø and Andersen, 2024*b*). Yet the Common European Asylum System sets only minimum reception standards and leaves the

level and form of material support to member states, so large differences across countries persist.[22] Harmonization of the real value of benefits could reduce the extent to which asylum applicants apply in multiple EU+ countries, lowering the administrative and legal burden of asylum migration.

Second, targeted procedural responses such as the BAMF fast-track response to Western Balkan applicants and the safe-country designation (see Appendix I and Guichard and Machado (2025)) reduce the expected value of asylum seeker benefits for low-probability claimants without affecting the treatment of genuine refugees, and they appear to have been effective in the German case. However, this may not be a viable and effective strategy for applicants from countries without visa-free travel access. Although our study has contributed to the understanding of the drivers of asylum claims with low prospects of success, it does not seem that reducing asylum seeker benefits is a viable strategy for reducing the number of asylum applications from relatively safe countries, which has in general proven difficult (see also the work of Bah and Batista (2025)). Future research should study the role of alternative policies, including regular (labor) migration alternatives, resettlement of refugees and third country asylum processing.

[22] The Reception Conditions Directive (Directive 2013/33/EU during our sample period, recast as Directive (EU) 2024/1346) binds member states as to the result to be achieved but leaves the form and level of support to them. Under the recast directive, receiving countries should provide housing, food, clothing, personal hygiene products and a daily expenses allowance to enable applicants to enjoy a minimum degree of autonomy in their daily life. However, all of these can be provided in kind or as vouchers.

## References


**Abadie, Alberto, Alexis Diamond, and Jens Hainmueller.** 2010. "Synthetic control methods for comparative case studies: estimating the effect of California's tobacco control program." *Journal of the American Statistical Association*, 105(490): 493–505.

**Abdelaaty, Lamis.** 2023. "Do Rights Violations Deter Refugees?" *Journal of Ethnic and Migration Studies*, 49(1): 110–133.

**Adema, Joop.** 2025. "The Welfare Magnet Hypothesis: Evidence from an Immigrant Welfare Scheme in Denmark: Comment." Forthcoming, *American Economic Review: Insights*.

**Adema, Joop, Lasha Chargaziia, Yvonne Giesing, and Panu Poutvaara.** 2025. "Refugees from Ukraine value job opportunities over welfare." *Proceedings of the National Academy of Sciences*, 122(32): e2502420122.

**Agersnap, Ole, Amalie Jensen, and Henrik Kleven.** 2020. "The welfare magnet hypothesis: evidence from an immigrant welfare scheme in Denmark." *American Economic Review: Insights*, 2(4): 527–542.

**Allen, William L., Mariá Fernández-Reino, and Isabel Ruiz.** 2025. "Immigration and the welfare state." *Oxford Review of Economic Policy*, 41(1): 64–86.

**Almeida, Alexander, Susan Athey, Guido Imbens, Eva Lestant, and Alexia Olaizola.** 2025. "Estimating Variances for Causal Panel Data Estimators." *arXiv preprint arXiv:2510.11841*, Version 2.

**Alscher, Stefan, Johannes Obergfell, and Stefanie Ricarda Roos.** 2015. "Migrationsprofil Westbalkan: Ursachen, Herausforderungen und Lösungsansätze." Bundesamt für Migration und Flüchtlinge Working Paper 63, Nürnberg.

**Anderson, James E., and Eric van Wincoop.** 2003. "Gravity with gravitas: A solution to the border puzzle." *American Economic Review*, 93(1): 170–192.

**Andersson, Henrik, and Kristoffer Jutvik.** 2023. "Do Asylum-Seekers Respond to Policy Changes? Evidence from the Swedish–Syrian Case." *Scandinavian Journal of Economics*, 125(1): 3–31.

**Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager.** 2021. "Synthetic difference-in-differences." *American Economic Review*, 111(12): 4088–4118.

**Artuc, Erhan, and Caglar Ozden.** 2018. "Transit migration: All roads lead to America." *The Economic Journal*, 128(612): F306–F334.

**Avdan, Nazli, Andrew S. Rosenberg, and Christopher F. Gelpi.** 2025. "Where There's a Will, There's a Way: Border Walls and Refugees." *Journal of Peace Research*, 62(2): 375–389.

**Bah, Tijan L., and Catia Batista.** 2025. "Why do People Migrate Irregularly? Evidence from a Lab in the Field Experiment in West Africa." NOVAFRICA, Nova School of Business and Economics Working Paper.

**BAMF.** 2017. "Aktuelle Zahlen zu Asyl, Ausgabe: Dezember 2017. Tabellen, Diagramme, Erläuterungen." Federal Office for Migration and Refugees.

**Beine, Michel, Simone Bertoli, and Jesús Fernández-Huertas Moraga.** 2016. "A practitioners' guide to gravity models of international migration." *The World Economy*, 39(4): 496–512.

**Bertoli, Simone, and Jesús Fernández-Huertas Moraga.** 2013. "Multilateral resistance to migration." *Journal of development economics*, 102: 79–100.

**Bertoli, Simone, Herbert Brücker, and Jesús Fernández-Huertas Moraga.** 2016. "The European crisis and migration to Germany." *Regional Science and Urban Economics*, 60: 61–72.

**Bertoli, Simone, Herbert Brücker, and Jesús Fernández-Huertas Moraga.** 2022. "Do Applications Respond to Changes in Asylum Policies in European Countries?" *Regional Science and Urban Economics*, 93: 103771.

**Borjas, George J.** 1999. "Immigration and welfare magnets." *Journal of Labor Economics*, 17(4): 607–637.

**Bratu, Cristina, Matz Dahlberg, Mattias Engdahl, and Till Nikolka.** 2020. "Spillover Effects of Stricter Immigration Policies." *Journal of Public Economics*, 190: 104239.

**Brekke, Jan-Paul, Marianne Røed, and Pål Schøne.** 2017. "Reduction or Deflection? The Effect of Asylum Policy on Interconnected Asylum Flows." *Migration Studies*, 5(1): 65–96.

**Brekke, Jan-Paul, Marianne Røed, and Pål Schøne.** 2023. "Family Matters: The Impact of National Policies on Asylum Destinations." *International Migration Review*.

**Bundesrepublik Deutschland.** 2014. "Gesetz zur Einstufung weiterer Staaten als sichere Herkunftsstaaten und zur Erleichterung des Arbeitsmarktzugangs für Asylbewerber und geduldete Ausländer vom 31. Oktober 2014." *Bundesgesetzblatt Teil I Nr. 49, S. 1649, 5. November 2014*.

**Bundesrepublik Deutschland.** 2015. "Verordnung zum Asylverfahrensbeschleunigungsgesetz vom 24. Oktober 2015 (§26 Abs. 2 Beschäftigungsverordnung)." *Bundesgesetzblatt Teil I Nr. 41, S. 1789, 27. Oktober 2015*.

**Bundesverfassungsgericht.** 2012*a*. "Pressemitteilung Nr. 35/2012: Mündliche Verhandlung in Sachen "Asylbewerberleistungsgesetz/Grundleistungen" am 20. Juni 2012." *https://www.bundesverfassungsgericht.de/SharedDocs/Pressemitteilungen/DE/2012/bvg12-035.html*, Press release of May 30, 2012.

**Bundesverfassungsgericht.** 2012*b*. "Urteil des Ersten Senats vom 18. Juli 2012." BVerfGE 132, 134; Az. 1 BvL 10/10 und 1 BvL 2/11.

**Chen, Jiafeng, and Jonathan Roth.** 2024. "Logs with Zeros? Some Problems and Solutions." *The Quarterly Journal of Economics*, 139(2): 891–936.

**Chernozhukov, Victor, Kaspar Wüthrich, and Yinchu Zhu.** 2021. "An exact and robust conformal inference method for counterfactual and synthetic controls." *Journal of the American Statistical Association*, 116(536): 1849–1864.

**Conley, Timothy G., and Christopher R. Taber.** 2011. "Inference with "Difference in Differences" with a Small Number of Policy Changes." *The Review of Economics and Statistics*, 93(1): 113–125.

**Council of the European Union.** 2009. "Council Regulation (EC) No 1244/2009 of 30 November 2009 amending Regulation (EC) No 539/2001 listing the third countries whose nationals must be in possession of visas when crossing the external borders and those whose nationals are exempt from that requirement." *Official Journal of the European Union, L 336, 18 December 2009, pp. 1–3.*

**Czaika, Mathias, and Hein de Haas.** 2017. "The Effect of Visas on Migration Processes." *International Migration Review*, 51(4): 893–926.

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

**Di Iasio, Valentina, and Jackline Wahba.** 2024. "The Determinants of Refugees' Destinations: Where do refugees locate within the EU?" *World Development*, 177: 106533.

**Docquier, Frédéric, and Chrysovalantis Vasilakis.** 2024. "Migrants' self-selection and the vicious circle of right-wing populism." IZA Discussion Paper 17402.

**Dustmann, Christian, Rasmus Landersø, and Lars Højsgaard Andersen.** 2024*a*. "Refugee Benefit Cuts." *American Economic Journal: Economic Policy*, 16(2): 406–441.

**Dustmann, Christian, Rasmus Landersø, and Lars Højsgaard Andersen.** 2024*b*. "Unintended Consequences of Welfare Cuts on Children and Adolescents." *American Economic Journal: Applied Economics*, 16(4): 161–185.

**eu-LISA.** 2015. "Eurodac—Annual Statistics." European Union Agency for the Operational Management of Large-Scale IT Systems in the Area of Freedom, Security and Justice (eu-LISA), Tallinn.

**European Commission.** 2025. "Reports on Partner Countries' Compliance with Visa-Free Travel Requirements." *Migration and Home Affairs, European Commission.*

**European Stability Initiative.** 2013. "Statistical Update: Asylum Seekers from the Western Balkans." *https://www.esiweb.org/pdf/ESI%20Statistical%20Update%20-%20Balkan%20asylum%20seekers%20-%2015%20Feb%202013.pdf*, 15 February 2013.

**Ferman, Bruno, and Cristine Pinto.** 2019. "Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity." *The Review of Economics and Statistics*, 101(3): 452–467.

**Ferwerda, Jeremy, Moritz Marbach, and Dominik Hangartner.** 2024. "Do immigrants move to welfare? Subnational evidence from Switzerland." *American Journal of Political Science*, 68(3): 874–890.

**Fitzgerald, Jack, Joop Adema, Lenka Fiala, Essi Kujansuu, and David Valenta.** 2025. "Non-Robustness in Log-Like Specifications." *MetaArXiv preprint*, Version 1.

**Flüchtlingsrat Berlin.** 2012. "Tabelle: Grundleistungen nach §3 AsylbLG ab November 2012 (BVerfGE 132, 134)." *https://fluechtlingsrat-berlin.de/wp-content/uploads/Tabelle_AsylbLG_BVerfG_2013.pdf*.

**Fratzke, Susan.** 2015. "Not Adding Up: The Fading Promise of Europe's Dublin System." Migration Policy Institute Europe, Brussels.

**Giulietti, Corrado.** 2014. "The welfare magnet hypothesis and the welfare take-up of migrants." *IZA World of Labor*, , (37).

**Görlach, Joseph-Simon, and Nicolas Motz.** 2021. "Spillovers and Strategic Interaction in Immigration Policies." *Journal of Economic Geography*, 21(2): 287–315.

**Guichard, Lucas.** 2020. "Self-Selection of Asylum Seekers: Evidence From France." *Demography*, 57(3): 1089–1116.

**Guichard, Lucas, and Joël Machado.** 2025. "The Externalities of Immigration Policies on Migration Flows: The Case of an Asylum Policy." *Journal of Economic Geography*, 25(1): 59–74.

**Hatton, Timothy J.** 2004. "Seeking asylum in Europe." *Economic Policy*, 19(38): 6–62.

**Hatton, Timothy J.** 2009*a*. "The Rise and Fall of Asylum: What Happened and Why?" *The Economic Journal*, 119(535): F183–F213.

**Hatton, Timothy J.** 2009*b*. "The rise and fall of asylum: what happened and why?" *The Economic Journal*, 119(535): F183–F213.

**Hatton, Timothy J.** 2016. "Refugees, Asylum Seekers, and Policy in OECD Countries." *American Economic Review*, 106(5): 441–445.

**Hatton, Timothy J.** 2023. "Asylum Recognition Rates in Europe: Policies and Performance." *European Journal of Political Economy*, 76: 102267.

**Kaushal, Neeraj.** 2005. "New immigrants' location choices: magnets without welfare." *Journal of Labor Economics*, 23(1): 59–80.

**Kosyakova, Yuliya, and Herbert Brücker.** 2020. "Seeking asylum in Germany: Do human and social capital determine the outcome of asylum procedures?" *European Sociological Review*, 36(5): 663–683.

**Landessozialgericht Nordrhein-Westfalen.** 2010*a*. "Beschluss vom 22. November 2010." Az. L 20 AY 1/09; Vorlage an das Bundesverfassungsgericht (BVerfG-Az. 1 BvL 2/11).

**Landessozialgericht Nordrhein-Westfalen.** 2010*b*. "Beschluss vom 26. Juli 2010." Az. L 20 AY 13/09; Vorlage an das Bundesverfassungsgericht (BVerfG-Az. 1 BvL 10/10).

**LoPalo, Melissa.** 2019. "The Effects of Cash Assistance on Refugee Outcomes." *Journal of Public Economics*, 170: 27–52.

**MacKinnon, James G, and Matthew D Webb.** 2020. "Randomization inference for difference-in-differences with few treated clusters." *Journal of Econometrics*, 218(2): 435–450.

**McFadden, Daniel.** 1974. "Conditional logit analysis of qualitative choice behavior." In *Frontiers in Econometrics*. , ed. Paul Zarembka, 105–142. New York:Academic Press.

**McFadden, Daniel.** 1978. "Modelling the choice of residential location." In *Spatial Interaction Theory and Planning Models.* , ed. Anders Karlqvist, Lars Lundqvist, Folke Snickars and Jörgen W. Weibull, 75–96. Amsterdam:North-Holland.

**Mossaad, Nadwa, Jeremy Ferwerda, Duncan Lawrence, Jeremy Weinstein, and Jens Hainmueller.** 2020. "In search of opportunity and community: Internal migration of refugees in the United States." *Science advances*, 6(32): eabb0295.

**Müller, Andreas.** 2013. "Die Organisation der Aufnahme und Unterbringung von Asylbewerbern in Deutschland." Bundesamt für Migration und Flüchtlinge (BAMF) Working Paper 55, Nürnberg. Working Paper der Forschungsgruppe des Bundesamtes; European Migration Network (EMN).

**Müller, Ruth, Alexandra Reinig, and Reiner Klingholz.** 2017. "Beschäftigung und Migration in der Region Westbalkan: Übersicht der Wanderungsbewegungen und Arbeitsmärkte in der Westbalkanregion." Berlin-Institut für Bevölkerung und Entwicklung, Berlin.

**Neumayer, Eric.** 2004. "Asylum destination choice. What makes some West European countries more attractive than others?" *European Union Politics*, 5(2): 155–180.

**Ortega, Francesc, and Giovanni Peri.** 2013. "The role of income and immigration policies in attracting international migrants." *Migration Studies*, 1(1): 47–74.

**Razin, Assaf, and Jackline Wahba.** 2015. "Welfare magnet hypothesis, fiscal burden, and immigration skill selectivity." *The Scandinavian Journal of Economics*, 117(2): 369–402.

**Santos Silva, João M. C., and Silvana Tenreyro.** 2006. "The log of gravity." *The Review of Economics and Statistics*, 88(4): 641–658.

**Schreier, S., L. Skrabal, and M. Czaika.** 2023. "DEMIG-QuantMig migration policy database." *https://quantmig.eu/data_and_estimates/policy_database/*, Accessed: 17 January 2025.

**Statistisches Bundesamt (Destatis).** 2019. "Qualitätsbericht: Wanderungen (Binnenwanderung, Außenwanderung, Gesamtwanderung)." Statistisches Bundesamt, Wiesbaden.

**Wett, Valentin, Katia Gallegos Torres, and Andreas Steinmayr.** 2024. "Opportunities or Benefits: Local Conditions and Refugee Labor Market Integration." Working paper.

# A Asylum seeker benefits

## A.1 Benefit levels in Germany

Table A1 documents monthly cash benefit levels under the AsylbLG for all six need levels (*Regelbedarfsstufen*, RS 1–6).

Table A1: AsylbLG monthly benefit levels by need level, 1993–2026 (EUR)

| Year | RS 1 | RS 2 | RS 3 | RS 4 | RS 5 | RS 6 | Family‡ |
|---|---|---|---|---|---|---|---|
| 1993–2011 | 225 | 199 | 199* | 199 | 179 | 133 | 736 |
| 2012† | 346 | 311 | 277 | 271 | 238 | 205 | 1,100 |
| 2013 | 354 | 318 | 283 | 274 | 242 | 210 | 1,124 |
| 2014 | 362 | 326 | 290 | 280 | 247 | 215 | 1,150 |
| 2015 | 359 | 323* | 287* | 278* | 245* | 213* | 1,140* |
| 2016 | 354 | 318 | 283* | 274* | 241* | 210* | 1,123* |
| 2017 | 354 | 318 | 283* | 274* | 241* | 210* | 1,123* |
| 2018 | 354 | 318 | 283* | 274* | 241* | 210* | 1,123* |
| 2019 | 344 | 310 | 275 | 275 | 268 | 214 | 1,136 |
| 2020 | 351 | 316 | 280 | 280 | 273 | 218 | 1,158 |
| 2021 | 364 | 328 | 292 | 323 | 282 | 247 | 1,221 |
| 2022 | 367 | 330 | 294 | 326 | 283 | 249 | 1,229 |
| 2023 | 410 | 369 | 328 | 364 | 304 | 278 | 1,361 |
| 2024 | 460 | 413 | 368 | 408 | 341 | 312 | 1,526 |
| 2025 | 441 | 397 | 353 | 391 | 327 | 299 | 1,464 |
| 2026 | 455 | 409 | 365 | 405 | 337 | 309 | 1,510 |

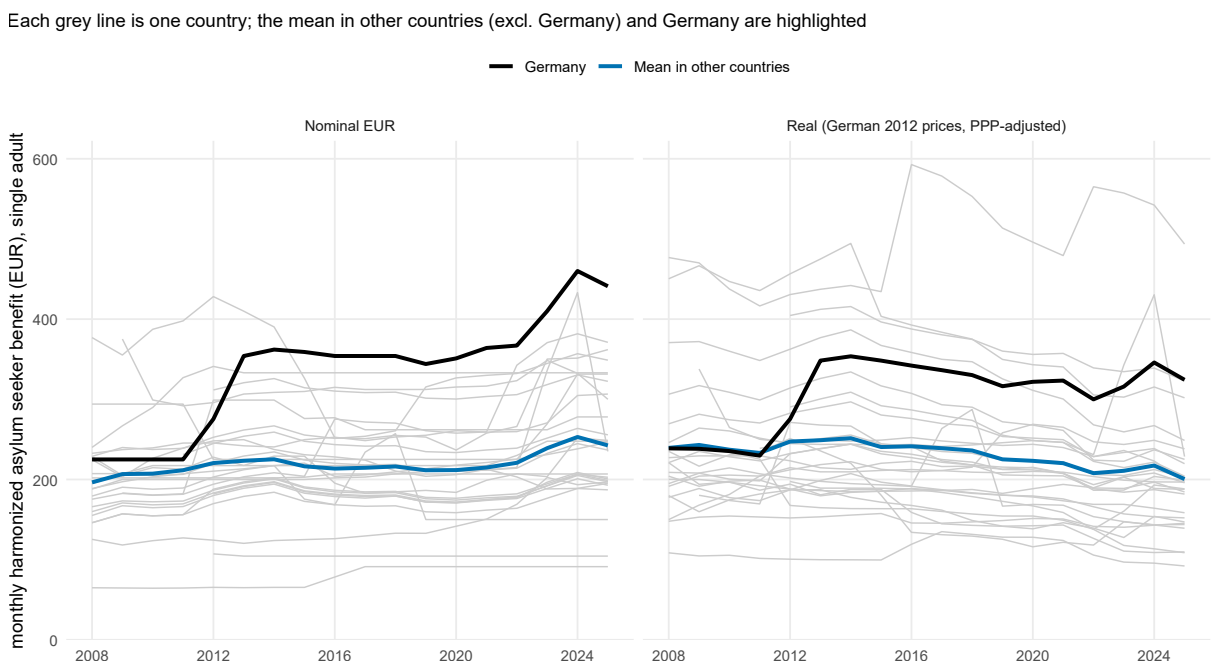
*Notes:* Monthly cash entitlements under § 3 AsylbLG (1993–2018) and § 3a AsylbLG (2019 onwards) for private accommodation. The need levels (*Regelbedarfsstufen*) follow the SGB categories: RS 1 is a single adult or single parent; RS 2 each adult partner in a couple sharing a household; RS 3 an adult without an own household (e.g. living in a facility or another person's household); RS 4 a youth aged 14–17; RS 5 a child aged 6–13; and RS 6 a child aged 0–5. A single adult applicant (RS 1) is the rate used throughout the paper. † Values effective from July 2012 following the Federal Constitutional Court ruling (*BVerfGE* 132, 134). ‡ Combined monthly benefit for a reference family of two adults (RS 1, RS 2), a child aged 5 (RS 6), and a child aged 10 (RS 5).

## A.2 Cross-country asylum seeker benefits

For all EU-27 + UK, NO, CH over 2000–2025, we coded the monthly asylum seeker benefit for a single adult applicant housed in collective quarters under two concepts: *full-in-kind* (housing and food in kind, cash is pocket money) and *housing-only* (only housing in kind). We combine the two into a single *harmonized asylum seeker benefit* (personal allowance plus a cash-or-imputed food allowance), the measure used throughout the paper. Coding rules: rates while the claim is being processed (not pre-registration or appeal-stage rates); largest city when subnational variation matters; basic health care typically in kind. Local currencies converted at Eurostat annual-average exchange rates and deflated by each country's HICP (rebased to 2012=100). The panel is constructed from primary legal sources rather than from secondary indices.

The 2012 harmonized-benefit cross-section is shown in the main text (Figure 3): Germany sits near the bottom of the cash-benefit distribution pre-ruling and near the top post-ruling. Figure A1 traces nominal and real levels over time; the German jump in July 2012 is by far the largest discrete change in the panel.

Figure A1: Harmonized asylum seeker benefits over time, nominal vs. real (German 2012 prices)




*Notes:* Each gray line is one country's monthly harmonized asylum seeker benefit for a single adult; the mean in other countries (blue, unweighted, excluding Germany) and Germany (black) are highlighted separately. Left panel: nominal euros. Right panel: real in German 2012 prices — deflated to 2012 by each country's own HICP, then multiplied by the ratio of Germany's 2012 food price level to the country's own. Germany 2012 is the numeraire (untransformed).

## A.3 Public attention to the AsylbLG (Google Trends)

Figure A2 plots the Google Trends search-interest index for the term *Asylbewerberleistungsgesetz* in Germany, aggregated to 4-week bins, from the beginning of 2010 to the end of 2014. Before the ruling, attention is low and sporadic, with occasional brief spikes but no sustained level; the ruling itself produces by far the single largest spike in the entire 2010–2014 window, and the subsequent baseline settles persistently higher than before. This supports the reading that the reform—rather than some earlier, gradually building policy debate—is what put asylum-seeker benefits into sustained public view (Section 2.3). A modest, short-lived rise in search interest in the four weeks before the ruling coincides with the Bundesverfassungsgericht's oral hearing on the case, held June 20, 2012 (marked by the dotted vertical line), and its accompanying press coverage; this pre-ruling attention did not carry through to a sustained elevated baseline the way the ruling itself did. Such anticipation would, if anything, bias against finding an effect of the ruling on applications, but is worth flagging explicitly so the pre-ruling bump is not mistaken for a threat to identification.

## A.4 Broader migration-policy restrictiveness (DEMIG)

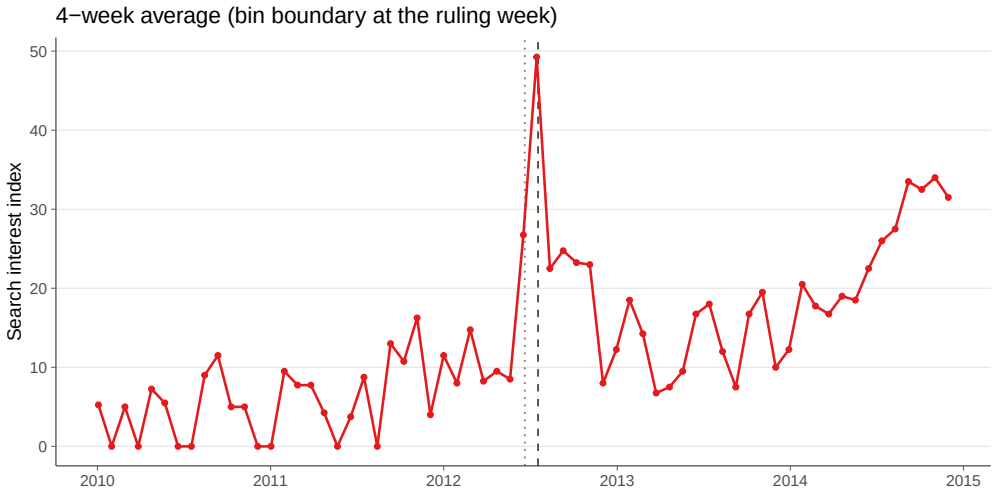
To gauge whether other European countries were adjusting their own migration or asylum policies over the same period, Figure A3 plots the cumulative change in refugee-targeted policy restrictiveness for 31 European countries between 2010 and 2013, using the DEMIG POLICY database (Schreier, Skrabal and Czaika, 2023). For each country and year, policy changes classified as more restrictive score +1 (scaled to +1.5 for changes coded as a "major change"), less restrictive changes score −1, and scores are cumulated from a 2010 baseline of zero. Movement is modest relative to the German shift and accumulates gradually across the whole window rather than clustering around the July 2012 ruling:

Table A2: Sources for the cross-country reception-benefit panel, by country

| Country | Sources |
|---|---|
| Austria | AIDA (2013, 2016–2024) |
| Belgium | AIDA (2016–2024, current) |
| Bulgaria | AIDA (2018, current) |
| Cyprus | AIDA (current country report) |
| Czech Republic | Act 325/1999 on Asylum |
| Denmark | Danish Aliens Act § 42b |
| Estonia | Act on Granting International Protection |
| Finland | Reception Act 746/2011 § 19 |
| France | AIDA (2016–2024, current) |
| Germany | asyl.net; BVerfG 1 BvL 10/10; Drucksache 12/4451; einwanderer.net; Fluechtlingsrat Berlin schedule; Niedersaechsischer Fluechtlingsrat; AIDA (2013) |
| Greece | AIDA (2016–2024, current) |
| Hungary | Govt Decree 301/2007 (XI.9.) |
| Ireland | SI 422/2000 Direct Provision; AIDA (2016–2024) |
| Italy | AIDA (2016–2024, current) |
| Latvia | Cabinet Regulation No. 595 |
| Lithuania | Lithuanian Refugee Council |
| Luxembourg | EMN Luxembourg 2013 Reception |
| Malta | AIDA (current country report) |
| Netherlands | AIDA (2016–2024, current) |
| Norway | UDI rates |
| Poland | Act on Granting Protection 13 June 2003; AIDA (2018, 2024) |
| Portugal | Lei 27/2008 Asilo |
| Romania | AIDA (current country report) |
| Slovakia | AIDA (current country report) |
| Slovenia | ZMZ Article 78 |
| Spain | Orden TIN/3469/2011; AIDA (2017–2018, 2023–2024) |
| Sweden | Forordning 1994:361; AIDA (2016–2024) |
| Switzerland | Asylverordnung 2 (AsylV 2); AIDA (2018) |
| United Kingdom | Asylum Support Regulations 2000; AIDA (2016–2024) |

*Notes:* Sources underlying the cross-country reception-benefit panel (Appendix A.2). "AIDA (years)" denotes the Asylum Information Database country reports for the listed years (plus the current online country page where indicated); AIDA is maintained by ECRE and the reports are available at https://asylumineurope.org (country pages and annual PDF updates). Remaining entries are national legislation, administrative rate schedules, EMN studies, and refugee-council documentation.

Figure A2: Google Trends search interest: "Asylbewerberleistungsgesetz" (Germany)

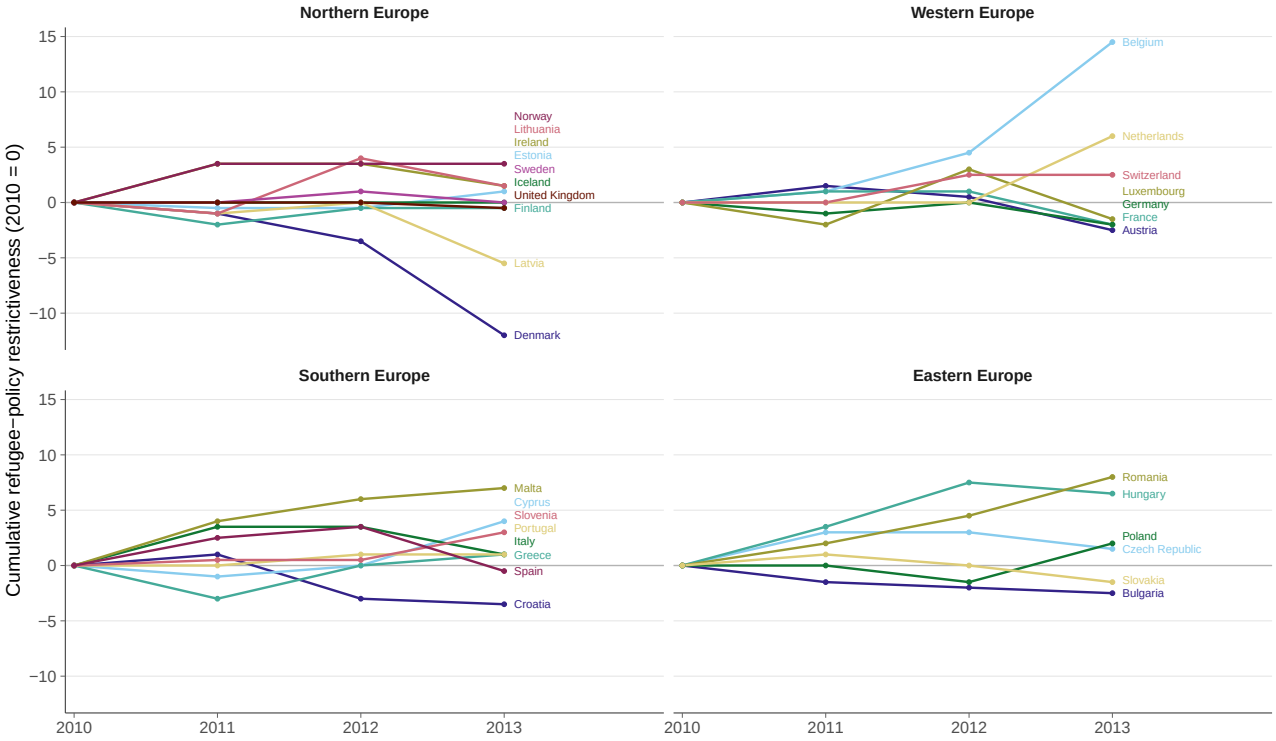



*Notes:* Google Trends relative search-interest index (0–100, Germany, all categories) for "Asylbewerberleistungsgesetz", aggregated to 4-week bins. The bins are anchored so that a bin boundary falls exactly at the week of the July 18 2012 ruling, separating a pre-reform from a post-reform bin rather than splitting the ruling week across bins. The dashed vertical line marks the ruling date; the dotted vertical line marks the Bundesverfassungsgericht's oral hearing on the case (June 20, 2012).

most countries drift within a narrow band around zero throughout, and even the countries with the largest cumulative shifts by 2013—Belgium (more restrictive) and Denmark (less restrictive)—move steadily across all four years rather than jumping in 2012 specifically. This is consistent with the

identifying assumption in Section 4.4 that contemporaneous policy changes elsewhere in the EU+ are not driving the German response.

Figure A3: Cumulative refugee-policy restrictiveness by country, 2010–2013




*Notes:* Cumulative refugee-targeted policy-restrictiveness score by country (2010 = 0), DEMIG POLICY database (Schreier, Skrabal and Czaika, 2023), faceted by UN geoscheme region (panels share a common $y$-axis). Restrictiveness-classified policy changes, 1990–2017, excluding nationality- or citizen-specific policies, are scored +1 (more restrictive) or −1 (less restrictive), scaled by 1.5 for changes coded as a "major change", and summed within country-year for policies tagged as targeting refugees (including irregular migrants) or all migrants.

# B Asylum application data

## B.1 First-time asylum applications

The main outcome is the monthly number of first-time asylum applications by origin × destination from Eurostat (`migr_asyappctzm`); we use first-time rather than total applications to capture recently arrived applicants (Section 3.1). The estimation sample covers the 25 EU+ destination countries listed in Table A3 and 168 origin citizenships (Appendix B.2).

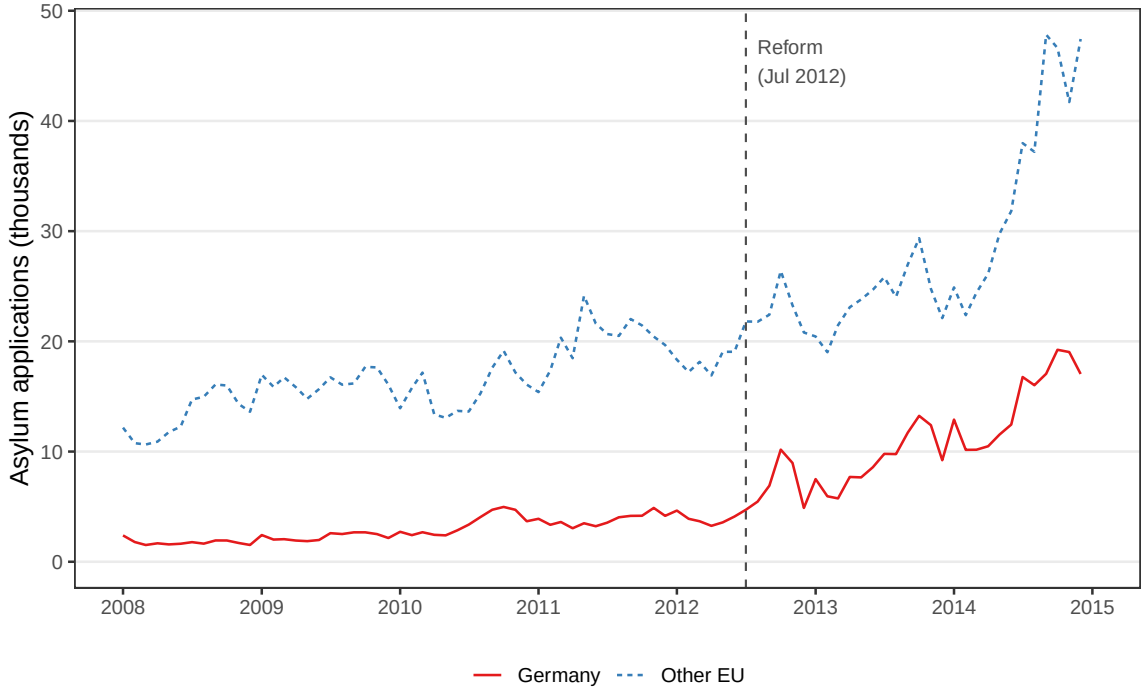
Table A3: Destination countries in the estimation sample

| | | |
|---|---|---|
| Belgium | Ireland | Romania |
| Bulgaria | Italy | Slovakia |
| Czech Republic | Latvia | Slovenia |
| Denmark | Lithuania | Spain |
| Estonia | Luxembourg | Sweden |
| Finland | Netherlands | Switzerland |
| France | Norway | United Kingdom |
| Germany | Poland | |
| Greece | Portugal | |

*Notes:* The 25 EU+ destination countries in the main difference-in-differences sample (2012 EU members excluding Austria, Croatia, Cyprus, Hungary, and Malta, plus Norway and Switzerland), sorted alphabetically.

Figure A4 shows monthly asylum applications to Germany and to the rest of the EU+ over 2008–2014.

Figure A4: Monthly asylum applications to Germany and other EU+ countries, 2008–2014




*Notes:* Monthly first-time asylum applications, Germany against the sum of the other EU+ destinations, 2008–2014. The series are plotted on separate vertical scales because the levels differ by roughly an order of magnitude. The vertical line marks July 2012, the month of the ruling. Source: Eurostat, `migr_asyappctzm`.

Table A4 reports the top-5 origin countries by application volume in Germany and in the rest of the EU+ separately, over the 11+11-month window.

Table A4: Top-5 origin countries by sub-window (Aug 2011–Jun 2013)

| | **Pre (Aug 2011–Jun 2012)** | | | | **Post (Aug 2012–Jun 2013)** | | |
|---|---|---|---|---|---|---|---|
| # | Country | Apps | Share | # | Country | Apps | Share |
| *Panel A: Germany* | | | | | | | |
| 1 | Afghanistan | 6,755 | 15.4% | 1 | Russia | 12,085 | 15.4% |
| 2 | Iraq | 5,030 | 11.5% | 2 | Serbia and Montenegro (former) | 8,960 | 11.4% |
| 3 | Serbia and Montenegro (former) | 4,690 | 10.7% | 3 | Syrian Arab Republic | 8,395 | 10.7% |
| 4 | Iran, Islamic Republic of | 3,530 | 8.1% | 4 | Afghanistan | 6,785 | 8.7% |
| 5 | Pakistan | 3,220 | 7.4% | 5 | Macedonia, the former Yugoslav Republic of | 5,045 | 6.4% |
| *Panel B: Rest of EU* | | | | | | | |
| 1 | Afghanistan | 13,005 | 7.2% | 1 | Syrian Arab Republic | 21,060 | 9.9% |
| 2 | Pakistan | 11,860 | 6.6% | 2 | Russia | 19,735 | 9.3% |
| 3 | Russia | 11,085 | 6.2% | 3 | Pakistan | 13,835 | 6.5% |
| 4 | Somalia | 9,335 | 5.2% | 4 | Afghanistan | 11,740 | 5.5% |
| 5 | Nigeria | 7,535 | 4.2% | 5 | Kosovo | 11,735 | 5.5% |

## B.2 Origin countries in the estimation sample

The main difference-in-differences sample comprises 168 origin citizenships, listed alphabetically in Table A5.

Table A5: Origin countries in the main estimation sample

| | | | |
|---|---|---|---|
| Afghanistan | Dominica | Liberia | San Marino |
| Albania | Dominican Republic | Libya | Sao Tome and Principe |
| Algeria | Ecuador | Macedonia, the former Yugoslav Republic of | Saudi Arabia |
| Andorra | Egypt | Madagascar | Senegal |
| Angola | El Salvador | Malawi | Serbia and Montenegro (former) |
| Antigua and Barbuda | Equatorial Guinea | Malaysia | Seychelles |
| Argentina | Eritrea | Maldives | Sierra Leone |
| Armenia | Eswatini | Mali | Singapore |
| Australia | Ethiopia | Marshall Islands | Solomon Islands |
| Azerbaijan | Federated States of Micronesia | Mauritania | Somalia |
| Bahamas | Fiji | Mauritius | South Africa |
| Bahrain | Gabon | Mexico | South Korea |
| Bangladesh | Gambia | Moldova, Republic of | South Sudan |
| Barbados | Georgia | Monaco | Sri Lanka |
| Belarus | Ghana | Mongolia | Sudan |
| Belize | Grenada | Montenegro | Suriname |
| Benin | Guatemala | Morocco | Syrian Arab Republic |
| Bhutan | Guinea | Mozambique | Taiwan |
| Bolivia | Guinea-Bissau | Myanmar/Burma | Tajikistan |
| Bosnia and Herzegovina | Guyana | Namibia | Tanzania, United Republic of |
| Botswana | Haiti | Nauru | Thailand |
| Brazil | Holy See (Vatican City State) | Nepal | Timor-Leste |
| British overseas countries and territories | Honduras | New Zealand | Togo |
| Brunei | India | Nicaragua | Tonga |
| Burkina Faso | Indonesia | Niger | Trinidad and Tobago |
| Burundi | Iran, Islamic Republic of | Nigeria | Tunisia |
| Cambodia | Iraq | Oman | Turkey |
| Cameroon | Israel | Pakistan | Turkmenistan |
| Canada | Ivory Coast | Palau | Tuvalu |
| Cape Verde | Jamaica | Palestinian Territory, Occupied | Uganda |
| Central African Republic | Japan | Panama | Ukraine |
| Chad | Jordan | Papua New Guinea | United Arab Emirates |
| Chile | Kazakhstan | Paraguay | United States |
| China | Kenya | Peru | Uruguay |
| Colombia | Kiribati | Philippines | Uzbekistan |
| Comoros | Korea, Democratic People's Republic of | Qatar | Vanuatu |
| Congo | Kosovo | Russia | Venezuela |
| Congo, the Democratic Republic of the | Kuwait | Rwanda | Vietnam |
| Cook Islands | Kyrgyzstan | Saint Kitts and Nevis | Western Sahara |
| Costa Rica | Lao People's Democratic Republic | Saint Lucia | Yemen |
| Cuba | Lebanon | Saint Vincent and the Grenadines | Zambia |
| Djibouti | Lesotho | Samoa | Zimbabwe |

*Notes:* All 168 origin citizenships (Eurostat citizenship categories) in the main difference-in-differences sample, sorted alphabetically down columns. A few entries are Eurostat aggregate or residual categories (e.g. "British overseas countries and territories", "Recognised non-citizens") rather than sovereign states.

## B.3 Dublin transfers

The Dublin Regulation—Dublin II over our sample window—determines which member state is responsible for examining an asylum application; in principle the country of first registration is responsible. Aggregate acceptance rates were on the order of 65–75 percent in 2010–2013 and effective transfer rates about 25 percent, so Eurodac matches (Section 3.3) capture a lower bound on the diversion margin rather than a counterfactual stock of inevitable transfers. Figure A5 traces the widening gap between requests, acceptances, and effective transfers over the period.

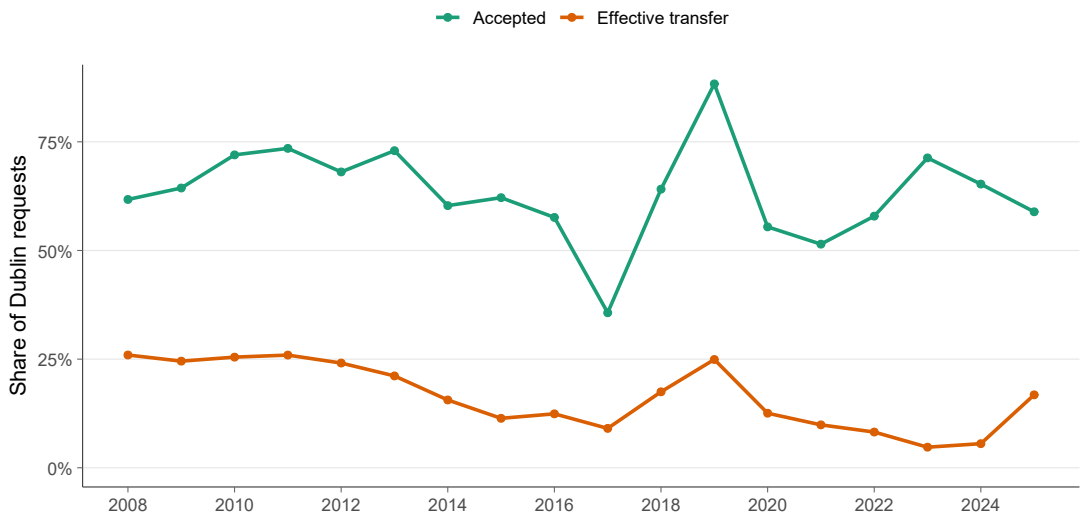
## B.4 Frontex irregular border-crossing detections

Total detections of irregular border crossings at the EU external frontier did not rise after the ruling relative to pre-ruling levels (Figure A6). This is what the diversion interpretation of Section 7 predicts: the ruling reshuffled applicants within the EU+ rather than pulling in new arrivals from outside it.

## B.5 Eurodac cross-border hits: descriptives

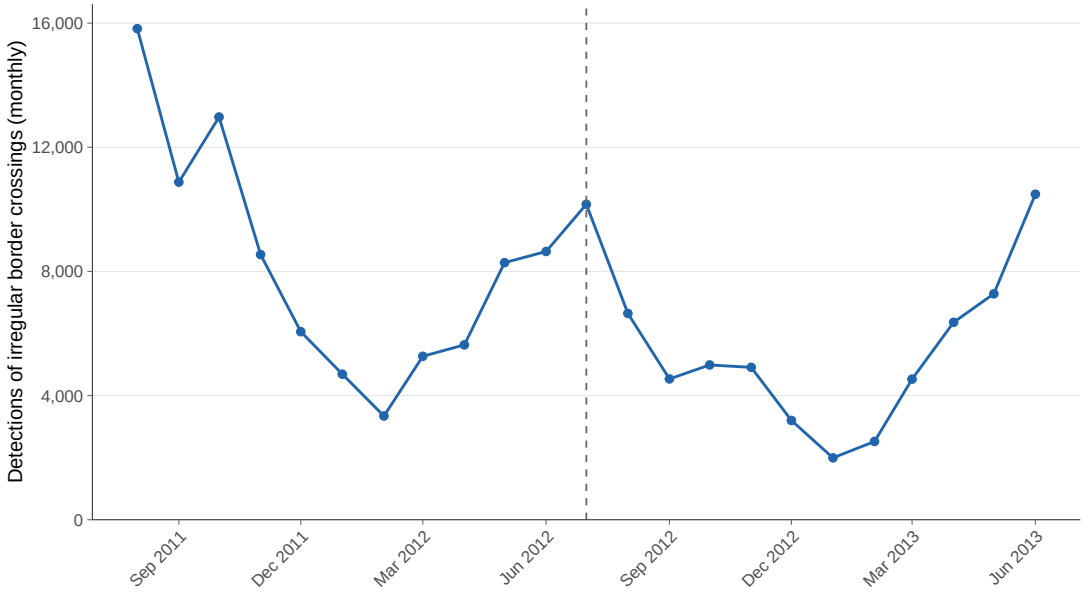
This subsection reports descriptive statistics on the Eurodac cross-border hits used in Section 7: inward and outward hits per first-time applicant by country (Figures A7 and A8), and the cross-border hit share for Germany versus the rest of the EU+ (Figure A9).

Figure A5: Dublin requests, acceptances, and effective transfers, EU+



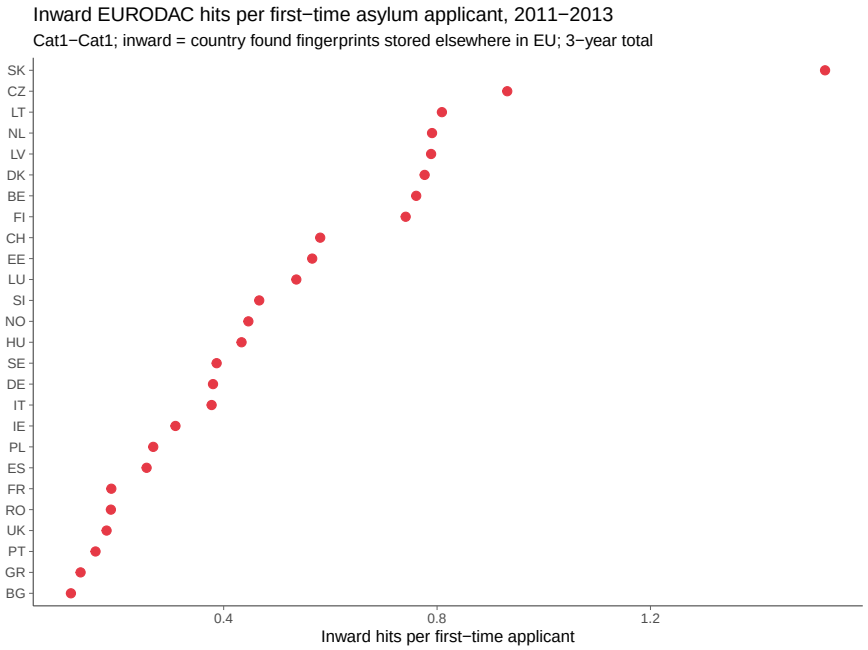

*Notes:* Aggregate outgoing Dublin requests, acceptances, and effective transfers across the EU+ (Eurostat `migr_dubro`/`migr_dubri`). Effective transfers remain a small fraction of detected onward movements, so a cross-border Eurodac match rarely results in an actual transfer.

Figure A6: Frontex IBC detections, August 2011–June 2013



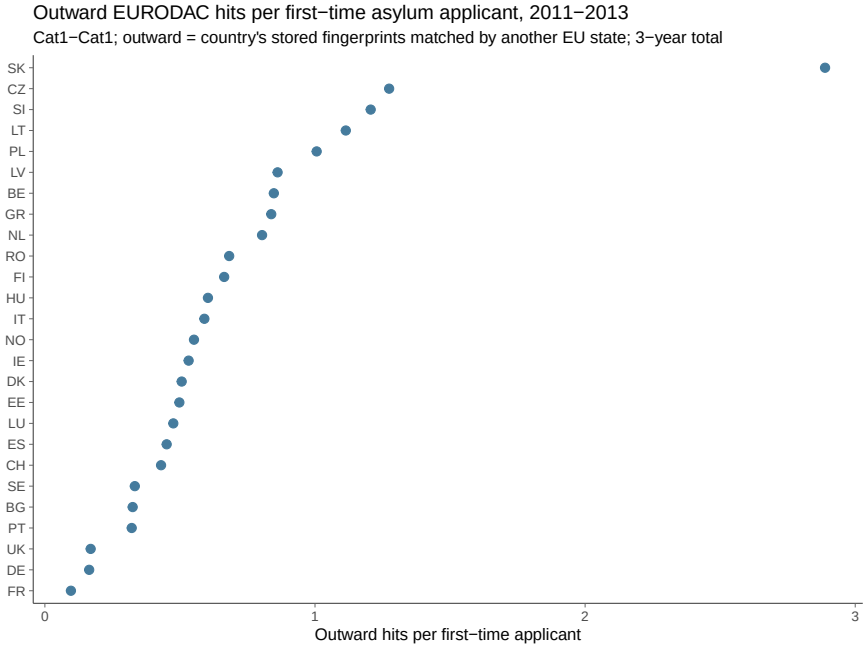

*Notes:* Monthly Frontex detections of irregular border crossings at the external EU border over the estimation window, August 2011–June 2013, by route. Detections are events rather than persons, so an individual crossing two sections is counted twice. The vertical line marks July 2012. Source: Frontex Detections of Illegal Border Crossings statistics.

Figure A7: Inward Eurodac hits per first-time applicant, by country



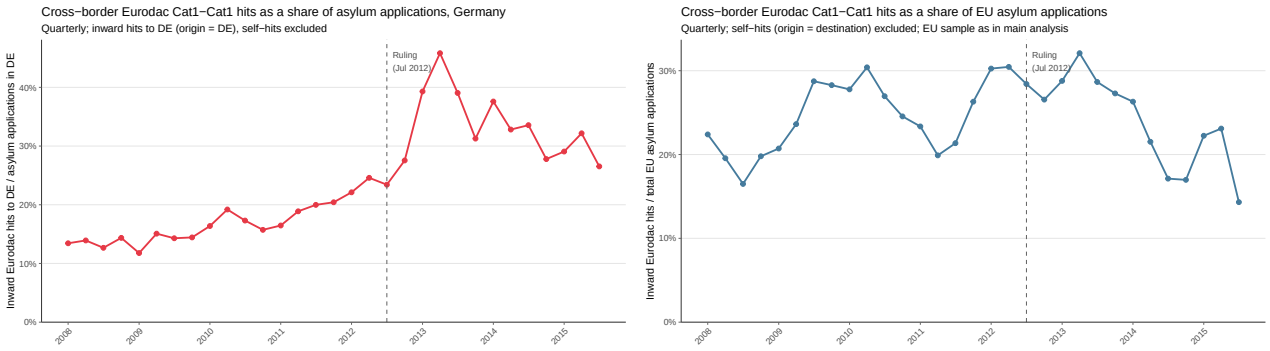

*Notes:* Inward Category 1–Category 1 Eurodac hits per first-time asylum application, by destination country, summed over 2011–2013. An inward hit occurs when a new applicant's fingerprint, taken on application in this country, matches a record stored earlier under another EU+ country (Section 3.3)—that is, the applicant had already applied for asylum elsewhere before arriving here. Germany sits among the higher-ratio destinations, consistent with substantial onward migration from southern entry countries toward Germany (compare Figure A8).

Figure A8: Outward Eurodac hits per first-time applicant, by country




*Notes:* Outward Category 1–Category 1 Eurodac hits per first-time asylum application, by country, same sample and construction as Figure A7 (summed over 2011–2013). An outward hit occurs when a country's own previously-stored fingerprint is later matched by another EU+ country's query—that is, an applicant who first applied here subsequently applied again elsewhere. Southern entry countries (Italy, Greece, Spain) show by far the highest outward ratios, while Germany sits in the middle of the distribution; together with Figure A7, this indicates the bulk of onward asylum migration runs from southern entry points toward Germany rather than the reverse.

Figure A9: Cross-border Eurodac hits as a share of asylum applications




(a) Germany (b) Rest of EU (aggregate)

*Notes:* Quarterly cross-border Eurodac hits as a share of asylum applications. Germany's share rises sharply after the July 2012 ruling, while the rest-of-EU aggregate moves only modestly.

# C Inference

Germany is the only treated destination and treatment is assigned at the destination × time level, so the design has a single treated cluster and cluster-robust standard errors overstate precision (Conley and Taber, 2011; Ferman and Pinto, 2019; MacKinnon and Webb, 2020). The procedures developed for this case replace the cluster-robust variance with a permutation distribution, and what they require in place of many treated clusters is an assumption of exchangeability: a statement that, absent the ruling, some set of relabelings of the data would have been equally likely to arise as the one observed. Which relabelings those are is what separates the two families of procedure.

Conformal inference permutes along time. It assumes that Germany's untreated outcomes are exchangeable across periods, so that absent the ruling the post-ruling window would be one draw among many from Germany's own path and the treatment date could be moved along the series without changing the distribution of what follows it (Chernozhukov, Wüthrich and Zhu, 2021). Permutation inference permutes across units. It assumes that untreated outcomes are exchangeable across destinations, so that absent the ruling Germany would be one candidate for treatment among the controls and the treatment could be reassigned to each of them in turn (MacKinnon and Webb, 2020).

Both assumptions are harder to state here than in the two-dimensional panels for which these procedures were written, because our data are an origin × destination × month panel and treatment varies on only one of its three margins. Each destination × month cell is itself a vector of origin-specific counts, so what has to be exchangeable is not a scalar outcome but the entire origin profile of a destination or of a month. Destinations differ in which origins they receive and months differ in which origins are on the move, so exchangeability on either margin asks for more here than the two-way phrasing suggests. The two procedures deal with the extra dimension differently: the conformal test collapses it, aggregating over origins into a single monthly series before permuting, whereas permutation inference leaves the panel intact and re-estimates the full gravity specification under each reassignment.

Neither assumption holds exactly here. Section C.1 asks which of the two margins is the more homogeneous, and so which of the two is the more plausible candidate; Section C.2 then sets out the conformal procedure and Section C.3 the permutation one.

## C.1 Which margin does the error variance vary over?

The choice between the two inference approaches is not purely a matter of taste, because the two margins are not equally homogeneous. Table A6 asks which of them carries more of the error variance, using the diagnostic of Almeida et al. (2025).

Two features of how we implement it should be stated first, because both depart from the specification the rest of the paper uses. The diagnostic is written for a linear model and does not apply directly to Poisson pseudo-maximum likelihood, where the Pearson residual is tied mechanically to the fitted mean and takes the value $-\sqrt{\hat{\mu}}$ at every zero of the outcome. We therefore run it on a log-transformed difference-in-differences at the destination × month level: applications are aggregated over origins, the outcome is the log of that total, and the fixed effects are destination and year-month. A *cell* in what follows is one such destination × month observation, and the *cell variance* is the variance of the regression error in it. The diagnostic is also fitted on untreated cells only, so Germany's post-ruling months are held out and the effect of the ruling cannot enter the variance model being estimated.

Writing the variance of the error in cell $(d, t)$ as $\sigma^2_{dt} = \exp(\kappa + \nu_d + \xi_t)$, the exponentiated fixed effects are multiplicative variance factors averaging one within their own margin, and the table reports how far each set of factors is spread around that mean.

The two margins differ considerably. A spread of 1.19 across destinations means that a typical destination has a cell variance more than one whole multiple above or below the panel average, and the destinations at either end differ by two orders of magnitude: the United Kingdom sits at 0.02 times the average variance and Bulgaria at 4.81. Across months the spread is 0.49, less than half as large, so months resemble one another much more closely than destinations do. Almeida et al. (2025) take a value above one as the marker that a margin matters for the choice of variance estimator, and only the destination margin clears it.

Germany is among the least volatile destinations in the panel, with a variance factor of 0.09, fourth-lowest of twenty-five. That ordering is what a model in relative changes implies, since a destination receiving few applications has noisier proportional movements than a large one, and it bears directly on how the permutation-inference $p$-values should be read: because the destinations making up the placebo pool are systematically more volatile than the treated unit, the placebo distribution is wider than the null it is meant to represent, and the resulting $p$-values are conservative on that account.

The corresponding concern on the time margin does not arise. If the post-ruling months were more volatile than the pre-ruling ones, a test that permutes the treatment date would compare a noisy window against quiet ones. They are not: the mean variance factor is 0.94 in the post-ruling months against 1.06 before.

One qualification belongs with this reading. What the diagnostic examines on the time margin is stricter than what the conformal procedure actually requires, because the moving-block version of Chernozhukov, Wüthrich and Zhu (2021) does not ask for exchangeability across months in the literal sense. It asks only that the residual process be stationary and strongly mixing, which admits the persistence that a monthly application series displays, at the cost of exact finite-sample size. The next subsection sets out what that relaxation buys and what it costs.

Table A6: Spread of the error variance across destinations and across months

| | |
|---|---|
| *Spread of the error variance* | |
| Across destinations | 1.19 |
| Across months | 0.49 |
| *Germany among the destinations* | |
| Variance factor | 0.09 |
| Rank, least to most variable | 4 of 25 |
| *Mean variance factor by period* | |
| Post-ruling months | 0.94 |
| Pre-ruling months | 1.06 |
| Mean of dependent variable | 5.53 |
| Destinations | 25 |
| Observations | 533 |

*Notes:* Variance diagnostic of Almeida et al. (2025), implemented as described above. The first block reports the standard deviation of each set of variance factors, a value above one indicating a margin across which the error variance is materially heterogeneous. The second block gives Germany's own factor and its rank among the twenty-five destinations, counting from the least to the most variable. The third block compares the mean factor of the eleven post-ruling months with that of the eleven pre-ruling months. *Sample.* All destinations, August 2011–June 2013, excluding July 2012. Hungary is excluded because its series is missing for sixteen of the twenty-two months, and twelve destination–months recording no applications are dropped because the outcome is logarithmic. Source: Eurostat, `migr_asyappctzm`.

## C.2 Conformal inference

The first of the two procedures for a single treated cluster permutes along time. We follow Chernozhukov, Wüthrich and Zhu (2021), whose conformal test treats the post-ruling window as one block among the many the series could have produced and asks how extreme the observed one is.

**Implementation.** Under the null that the ruling had no effect, Germany's post-ruling months are untreated months like any other, so the whole series can be treated as one draw from a stationary process and the post-ruling block asked whether it looks like the rest of that draw. The counterfactual is the PPML specification of equation (6) itself rather than a separate synthetic control, and it is estimated with Germany's post-ruling cells held out rather than with the null imposed on them, for reasons set out below. The test then runs in four steps.

1. *Estimate the counterfactual.* Equation (6) is fitted with Germany's post-ruling cells held out. The origin × year-month effect for a post-ruling month is then identified from the control destinations, and the Germany–origin pair effect from Germany's pre-ruling months alone, which is the count-data counterpart of choosing synthetic control weights on the pre-ruling period only.

2. *Form the gap series.* Applications in Germany are predicted for every month and aggregated over origins, leaving one gap per month, $u_t = \log(\sum_o Y_{ot}) - \log(\sum_o \hat{\mu}_{ot})$. Because the counterfactual comes from a Poisson pseudo-maximum-likelihood fit, it is worth being explicit that this is not a Pearson or deviance residual but a proportional deviation: $u_t = 0.10$ means applications in Germany ran about eleven percent above what the fixed effects predict. Only origins with a counterfactual in every month enter, so movement in $u_t$ is not driven by a changing set of cells.

3. *Compute the statistic.* The statistic is the mean absolute gap over the post-ruling block.

4. *Rank it among the rotations.* The gap series is rotated cyclically through every starting point, the statistic is recomputed on each rotation, and the unrotated value is ranked among the 46 of them. The $p$-value is that rank divided by 46.

**Why the null is not imposed at the estimation stage.** Chernozhukov, Wüthrich and Zhu (2021) fit the counterfactual on all $T$ periods with the null imposed, on the reasoning that post-treatment outcomes are untreated outcomes under $H_0$ and so carry information about the counterfactual like any other period. That is what makes their estimator a symmetric function of time and the residual series rotatable, and it costs nothing when the treated unit contributes little to the fit, as in a synthetic control whose weights fall on the donor pool. Our counterfactual is not of that kind. The origin × year-month effects of equation (6) are estimated from all destinations jointly, and Germany is the largest of them, accounting for a substantial share of applications in the typical origin–month cell. Leaving its post-ruling cells in the estimation would let the post-ruling increase pass into the very fixed effects it is measured against, along two channels. The origin × year-month effect for a post-ruling month would rise with Germany's own applications in that month, shrinking the measured gap; and the Germany–origin pair effects, estimated over the whole window rather than the pre-ruling part of it, would carry the post-ruling level back into the fitted values for the pre-ruling months, inflating the gaps there. The statistic on the true block would fall while the rotations it is ranked against would rise, so the test would lose power on both sides at once, and the loss would reflect Germany's weight

in the panel rather than anything about the ruling. We therefore hold those cells out, as step 1 above describes. The price is that the post-ruling residuals are out-of-sample while the pre-ruling ones are in-sample, so the estimator is no longer symmetric in time and the rotations are no longer exactly comparable. Whether that asymmetry inflates the post-ruling statistic is testable, and the placebo tests that follow are the test, since they hold a pseudo-treated block out in exactly the same way.

Two further choices in that construction deserve comment. The gap is taken in logs because equation (6) is multiplicative, so the proportional deviation is the scale-free residual; a gap in levels would inherit the trend in the size of the series and could not plausibly be treated as stationary. And we use moving block rather than unrestricted permutations. Both are valid under exchangeability, but only the moving blocks remain valid under serial dependence, because the realized statistic is computed on a block of consecutive months whose average carries more variance than an average over months drawn at random from a persistent series.[23] The price is resolution: the cyclic shifts number only $T$, so the smallest attainable $p$-value is $1/T$. This is why the test is run on a window with a 35-month pre-ruling period reaching back to August 2009 rather than on the symmetric 11-month window of the main specification. On the latter the floor would be $1/22 = 0.045$, which leaves a five percent test a single attainable rejection and no ability to discriminate among columns; on the former it is 0.022.

**Placebo specification tests.** Chernozhukov, Wüthrich and Zhu (2021) recommend testing the exchangeability assumption directly rather than asserting it, and Table A7 does so. The procedure discards the post-ruling months entirely, refits the counterfactual on the pre-ruling data alone, and hands the last $\tau$ pre-ruling months to the test as though they were the treated block. Nothing happened in those months, so the null of no effect is true by construction and a rejection is evidence against the maintained assumptions rather than evidence of an effect. With 35 pre-ruling months the floor is 0.029, leaving room to reject at conventional levels.

The specification passes. Nothing rejects at five percent in any column at any $\tau$, the smallest value in the table being 0.057, and the great majority lie well away from the boundary. This is also the test on the margin that matters, since the moving blocks are the reference set the data permit: the first-order autocorrelation of the pre-ruling gap is 0.77, so the unrestricted permutation group, which would deliver exact size under exchangeability, is not available here.

**Pointwise confidence intervals.** Because the null is sharp, specifying the effect for every month under test, the same machinery delivers confidence intervals and not only a $p$-value for zero: one inverts the test over a grid of candidate effects and retains those the data do not reject. Figure A10 does this month by month across the whole window. Steps 2 to 4 above are unchanged; only step 1 differs, because test inversion has to impose the candidate effect at the estimation stage. For a given month the data under the null are the pre-ruling months together with that month alone, the candidate

[23] Exchangeability of the residuals gives exact finite-sample size control under any permutation group, but Chernozhukov, Wüthrich and Zhu (2021) do not rely on it for the moving block case. They require instead that the residual process be stationary and strongly mixing with summable mixing coefficients, together with consistent estimation of the counterfactual, which delivers approximate validity with a size distortion that vanishes as the pre-ruling period lengthens. Mixing asks that the dependence between two stretches of the series fall away as the gap between them widens, and summability that it fall away fast enough for a long pre-ruling period to contain many effectively independent blocks, so the persistence of an autoregressive process is admissible while a unit root or long memory is not; see Chernozhukov, Wüthrich and Zhu (2021) for the size bound and for an alternative route that replaces consistency by a stability condition on the estimator.

Table A7: Placebo specification tests for the conformal procedure

| | | Excluded origins | | | Control destinations | | |
|---|---|---|---|---|---|---|---|
| | Full | Syria | Excl. MENA | West-Balkan | Old EU | North-West | Excl. BE, LU, RO, SE |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| $\tau = 1$ | 0.371 | 0.200 | 0.457 | 0.343 | 0.229 | 0.371 | 0.257 |
| $\tau = 2$ | 0.457 | 0.457 | 0.771 | 0.429 | 0.371 | 0.571 | 0.371 |
| $\tau = 3$ | 0.600 | 0.571 | 0.857 | 0.543 | 0.514 | 0.686 | 0.400 |
| $\tau = 4$ | 0.600 | 0.486 | 0.771 | 0.629 | 0.543 | 0.657 | 0.457 |
| $\tau = 6$ | 0.229 | 0.143 | 0.371 | 0.229 | 0.229 | 0.257 | 0.286 |
| $\tau = 8$ | 0.200 | 0.057 | 0.286 | 0.229 | 0.143 | 0.429 | 0.286 |
| $\tau = 11$ | 0.400 | 0.200 | 0.543 | 0.343 | 0.286 | 0.457 | 0.200 |

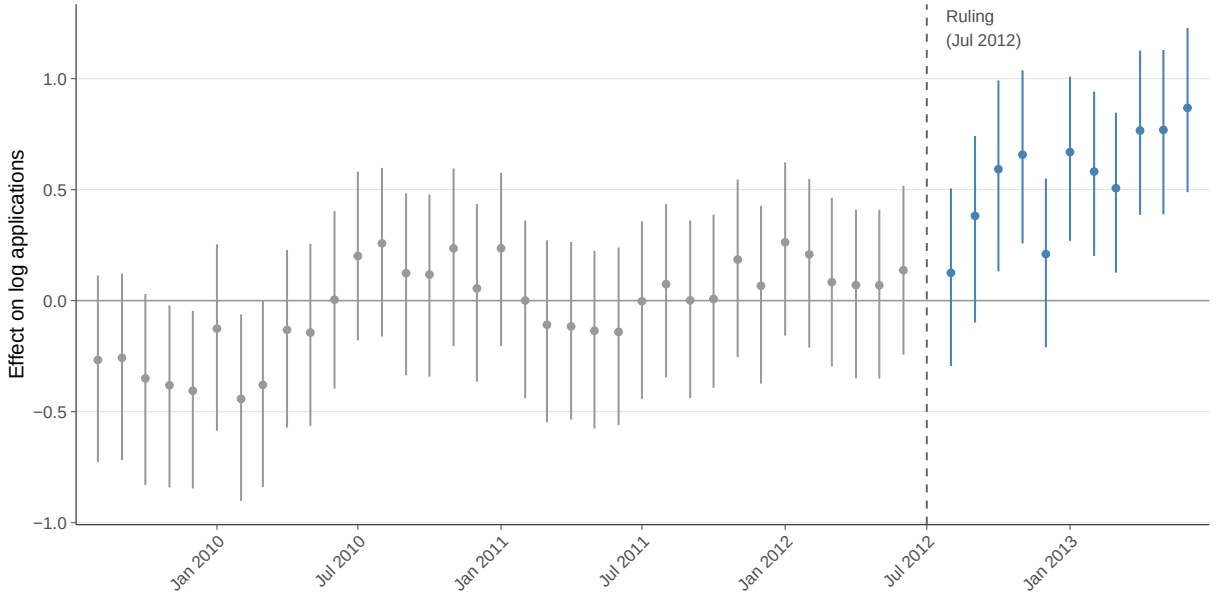
*Notes:* Placebo specification test of Chernozhukov, Wüthrich and Zhu (2021), their Section 4.3. Only the 35 pre-ruling months of the long window (August 2009–June 2012, excluding July 2012) are used. For each $\tau$ the null $H_0 : \alpha_{T_0-\tau+1} = \cdots = \alpha_{T_0} = 0$ is tested by treating the last $\tau$ pre-ruling months as the treated block, so under correct specification the null is true and the entries should behave like uniform draws. The counterfactual is built as in Table 1: Germany's pseudo-treated cells are held out of the PPML estimation, the origin $\times$ month effects for those months come from the control destinations, and only origins with a counterfactual in every month enter the gap series. The statistic is the mean absolute log gap over the pseudo-treated block, and the gap series is rotated through all 35 cyclic shifts, so the smallest attainable value is 0.029. Samples are identical to the corresponding columns of Table 1.

effect $a$ is imposed by scaling Germany's cells in that month by $\exp(-a)$, which is the count-scale counterpart of subtracting a level in a linear model, and the counterfactual is re-estimated on the null-imposed sample. That is what makes the estimator symmetric in time and the residuals rotatable. Absorption is at work here too, but with a milder consequence than it would have for the block test: at the candidate value equal to the true effect the imposed cells are untreated cells, so nothing is absorbed and the interval is centered where it belongs, while candidates away from the truth are partly absorbed and are therefore rejected less readily. The intervals are accordingly wider than a counterfactual Germany could not influence would deliver, rather than displaced.

Running the exercise over the pre-ruling months as well as the post-ruling ones puts the graphical placebo check and the estimate on one axis. The pre-ruling intervals are the check, since the null holds in those months by construction: 3 of the 35 exclude zero, or 9 percent, against the tenth a correctly sized ten percent procedure would leave. They also show a slow drift in the level, from below zero through 2009 and 2010 to above it by the first half of 2012, which the placebo test of Table A7 cannot register because a statistic built on the mean absolute gap treats a large negative early value and a large positive late value alike. Part of that drift is identifiable, the trough coinciding with the suspension of Dublin transfers to Greece and the Western Balkan visa liberalisation discussed below. The assumption of a stable dispersion is therefore defensible while the assumption of a stable level is weaker, and the estimates are measured against a counterfactual that was already drifting upwards before the ruling. The post-ruling intervals themselves trace a response that builds rather than jumps, from 0.13 in the first post-ruling month, with an interval spanning zero, to 0.87 by the end of the window, matching the event-study path of Figure 5.

**Placebo ruling dates.** A third exercise sits between the two procedures and is easier to read than either. Instead of permuting a residual series or reassigning the treated unit, we simply move the ruling date. Figure A11 re-estimates equation (6) at every calendar month from August 2009 onward, holding the $\pm 11$-month event window fixed throughout, and plots the resulting $t$-statistic against the date of the fake ruling. It rests on the same idea as the conformal test, that a real effect should

Figure A10: Pointwise ninety percent conformal confidence intervals, by month




*Notes:* Confidence intervals for the effect on log applications in Germany in each month, obtained by inverting the conformal test of Chernozhukov, Wüthrich and Zhu (2021) over a grid of candidate effects in steps of 0.02 and retaining those with $p > 0.10$. For each month the data under the null are the 35 pre-ruling months from August 2009, together with that month itself when it is post-ruling; the null is imposed by scaling Germany's cells in the month under test by $\exp(-a)$, the counterfactual is re-estimated on the null-imposed sample, and the block under test is the single month, so the cyclic rotations reduce to ranking its absolute log gap among those of the window. Gray markers are pre-ruling months, where the null of no effect is true by construction and the intervals serve as a graphical placebo check; blue markers are post-ruling months, in the color used for estimates elsewhere in the paper. The dashed vertical line marks July 2012, the month of the ruling, which is excluded throughout. The sample is that of column (1) of Table 1.

be large relative to what the same specification produces at dates when nothing happened, but it is computed on the estimating equation itself rather than on a residual series, so every point is directly comparable to the headline estimate.

Rolling the date this way also reaches far beyond the conformal window, and in particular it uses post-reform months. Those are legitimate comparisons rather than contaminated ones. Whatever the ruling did to the level of applications in Germany, there is no reason for it to have left *differential trends* behind years afterwards, so a window centered on 2018 or 2022 compares the German and control series on much the same terms as any other. What the exercise delivers is therefore a placebo distribution running over nearly two hundred months, against which July 2012 can be ranked.

Read that way, the ruling is extreme but not uniquely so. Two episodes exceed it in the pooled series, and the refugee-crisis years produce larger values still once the origins are split, which is worth confronting directly rather than leaving to a summary statistic. We take the three in turn, in chronological order.

- *2009–2011: Dublin transfers to Greece, and Western Balkan visa liberalisation.* The pooled series peaks at May 2010 ($t = 6.27$), and no single origin accounts for much of it, since removing the largest contributor moves the $t$-statistic by less than one. Two distinct things are happening, on opposite sides of the recognition-rate split.

  On the low-recognition side, Serbia, Montenegro and North Macedonia obtained visa-free Schengen travel in December 2009 and Albania and Bosnia and Herzegovina in December 2010, and applications from those origins rose sharply across the destinations that received them, Germany

among the largest.

The high-recognition side matters more for reading the figure, and it peaks separately at August 2010 ($t = 6.46$). It has a single identifiable cause: over precisely this window, returning asylum seekers to Greece under the Dublin Regulation ceased to be possible, so applicants who entered the EU through the eastern Mediterranean, and whose claims should have been examined in Greece, lodged them further north instead. The change came through a sequence of decisions rather than a single act, which is why the effect builds gradually rather than jumping at one date. The Bundesverfassungsgericht blocked an individual transfer to Greece by interim order on 8 September 2009 (2 BvQ 56/09) and again on 8 December 2009 (2 BvR 2780/09); administrative courts across Germany followed; the Second Senate heard the underlying constitutional challenge on 28 October 2010; and in January 2011 the Federal Ministry of the Interior directed the Federal Office to stop transferring to Greece altogether and to examine the claims itself, initially through 12 January 2012. In the same month the European Court of Human Rights decided *M.S.S. v Belgium and Greece*, which concerned an Afghan applicant transferred to Greece.

The applications data carry the signature this mechanism should leave. Over the event window around May 2010, applications in Germany from the origins that traveled the eastern Mediterranean route (Afghanistan, Iraq, Iran, Somalia, Syria, Eritrea and Pakistan) rose by 47 percent, while applications from those same origins to all other destinations *fell* by 7 percent, taking Germany's share of them from 19.7 to 27.8 percent. A general increase in arrivals would have raised both; only a reallocation of where a claim can be lodged raises one while lowering the other. The same reading explains the individual origins the leave-one-out exercise picks out: Iranian applications across Europe did rise after the crackdown that followed the June 2009 election, so Germany took a growing share of a growing total, whereas Somali applications in Germany rose while falling elsewhere in Europe, which is redistribution and not inflow. One mechanical caveat applies to this whole window. Fewer destinations reported to Eurostat's monthly series before 2011, between 18 and 24 against 27 from 2012 on, so the comparison group behind the earliest estimates is both smaller and less representative than later in the series.

- *2014–2017: the refugee crisis.* These months are excluded from the placebo band of Figure A11, and it is worth saying why rather than leaving the exclusion to the figure note. They are also the most extreme months in the whole series: the pooled panel reaches roughly $-12$ in 2016, and among the high-recognition origins a single crisis month reaches about $+15$, several times the ruling's own statistic in that panel. What moved was not a benefit level. Over these years applications reallocated across EU+ destinations on a scale no national policy change produces, and Germany's position within that reallocation moved repeatedly and for reasons ranging from the suspension of Dublin returns for Syrian applicants in 2015 to the closure of the Western Balkan route and the EU–Turkey statement in 2016. A window drawn from this period is not a draw from the same distribution as a quiet one, in either direction, so including these months would widen the reference distribution without saying anything about how large a spurious effect can plausibly be.

- *2022: the Western Balkan route after the pandemic.* This is the clearer-cut of the two exceeding episodes. The pooled series peaks at July 2022 ($t = 8.20$), and dropping origins one at a time

shows it resting on two: Afghanistan, whose applications in Germany rose by two thirds while the rest of the EU+ received slightly fewer, and Syria, whose applications in Germany rose by four fifths against a quarter elsewhere. Both traveled the Western Balkan route, which in 2022 recorded its heaviest traffic since 2015 once the pandemic-era closures lifted and which fed Germany disproportionately; Serbia's visa-free arrangements with several non-European states, suspended under EU pressure late that year, added a smaller set of origins to the same corridor.

The origin group is the one that drove 2009–2010, but the mixture is different. Applications in Germany from the eastern Mediterranean origins rose by 58 percent over this window, and applications from them to other destinations also rose, by 10 percent, taking Germany's share from 38.0 to 46.7 percent. Unlike 2009–2010, then, 2022 is a genuine increase in arrivals on top of a shift toward Germany rather than a reallocation alone. A more mundane factor works in the same direction. The Eurostat series counts formal applications, which in Germany are lodged some weeks after arrival is first registered, and that gap widens whenever the registration system is congested, so a surge in arrivals shows up in the applications data as a sharper rise and a sharper subsequent fall than the underlying arrivals warrant. That is part of why the 2022 swing is as large as it is in both directions.

None of the three is a threat to the design. Each has a visible cause that has nothing to do with a benefits ruling, each shows up in a series constructed precisely to make such episodes visible, and the ruling's own estimate is distinguished not by being the largest but by being the one for which the timing, the recognition-rate gradient, and the mechanism all line up.

### C.3 Permutation inference

The second of the two procedures for a single treated cluster permutes across units. We follow MacKinnon and Webb (2020): the treatment is reassigned to each control destination in turn, the baseline specification is re-estimated on every reassignment, and Germany's own statistic is ranked against the placebo distribution that results. This is the procedure Section 4.5 sets aside for the main tables, and the exhibits here are the reason: what makes its $p$-values hard to read is visible only draw by draw, with the destinations behind them named. The exercise runs in four steps.

1. *Choose the placebo pool.* The pool is the 15 non-German destinations with at least 1,000 asylum applications over the eleven pre-ruling months. The screen is discussed immediately below.

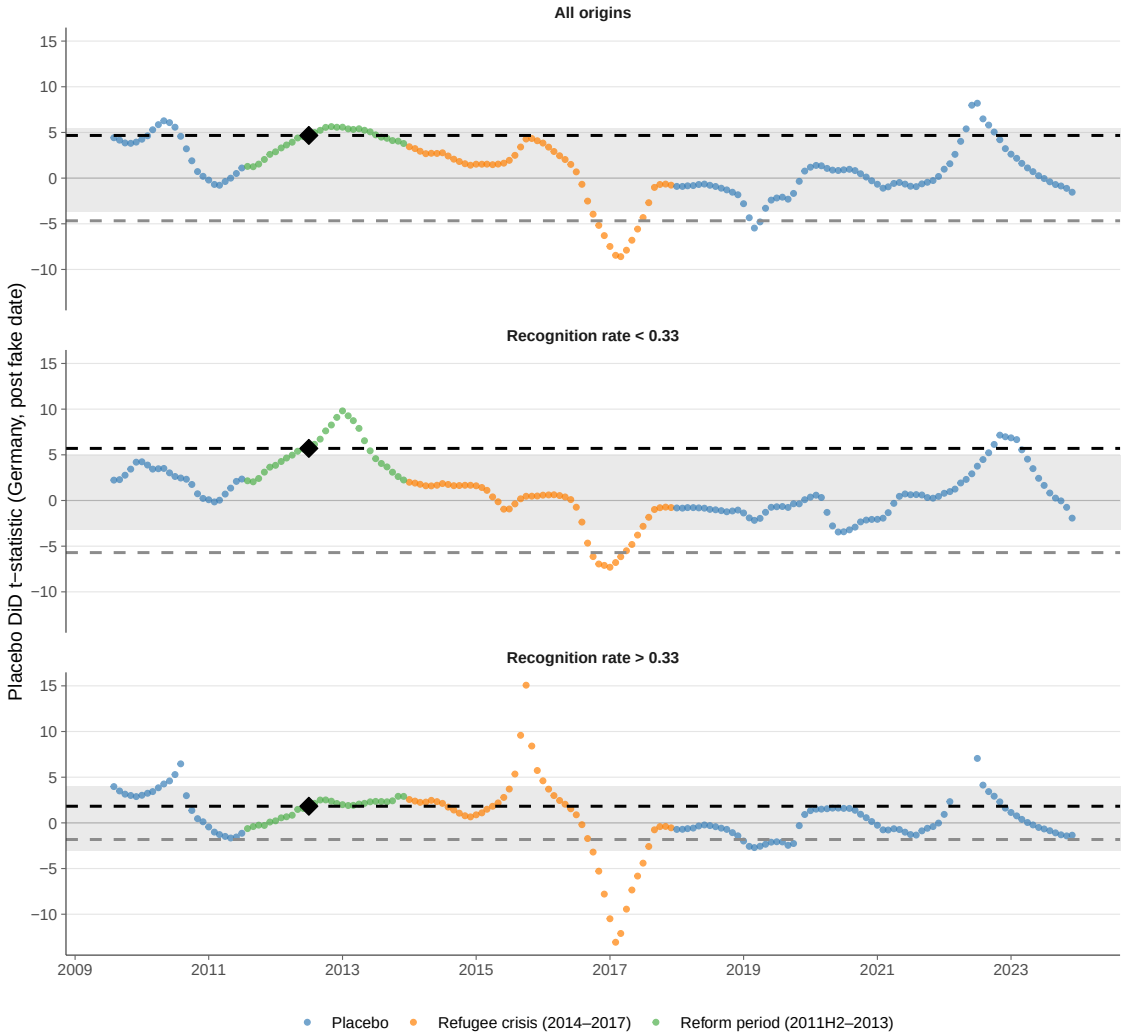
2. *Reassign the treatment.* For each destination in the pool, the July 2012 indicator is moved to it. Germany's own observations remain in the panel as an ordinary control destination in every draw.

3. *Re-estimate.* Equation (6) is refitted under each reassignment and the $t$-ratio on the reassigned indicator is recorded.

4. *Rank.* Germany's $t$-statistic is ranked among the $J+1$ draws. The two-sided $p$-value is the share of draws whose $|t|$ is at least Germany's, so the smallest attainable value is $1/(J+1)$.

Figure A12 shows that Germany's actual $t$-statistic lies in the tail of the placebo distribution. The most extreme placebo destinations are Belgium ($t = -8.37$), Romania ($t = -6.20$), Luxembourg ($t = -5.10$).

Figure A11: Placebo ruling dates at monthly frequency




*Notes:* Each point is the $t$-statistic on $\text{DE} \times \mathbb{1}(t > t_0)$ from equation (6) re-estimated with the ruling moved to that calendar month, keeping the same $\pm 11$-month event window; months whose full post-window would run past the end of the data are dropped rather than the window shortened. Panel (a) pools all origins. Panels (b) and (c) split them at $R_o \lessgtr 0.33$, using the fixed 2010Q1–2012Q2 classification of the main heterogeneity analysis, so the groups are the same set of countries at every date. Points are colored by period: the reform and its immediate aftermath in green, the 2014–2017 refugee crisis in orange, and every other candidate month in blue. The black diamond marks the actual July 2012 ruling at its true date; the black dashed line carries its value across the panel and the gray dashed line mirrors it at $-t$, so that a placebo month more extreme than the ruling in the opposite direction is as visible as one more extreme in the same direction. The shaded band spans the mean placebo $t$ in that panel plus and minus 1.645 standard deviations, computed over the blue months only: it is a 90 percent interval for the placebo draws themselves, not a confidence interval for their mean, since the question is whether a single estimate is unusual against the placebo distribution. The reform and crisis months are excluded from it because the first is contaminated by the ruling and the second by a reallocation of applications across destinations far larger than any national policy change produces. A handful of 2022 candidate months are dropped from panel (c), whose ten-origin sample is thin enough that the two-way clustered variance matrix becomes near-singular. All three panels share a vertical scale.

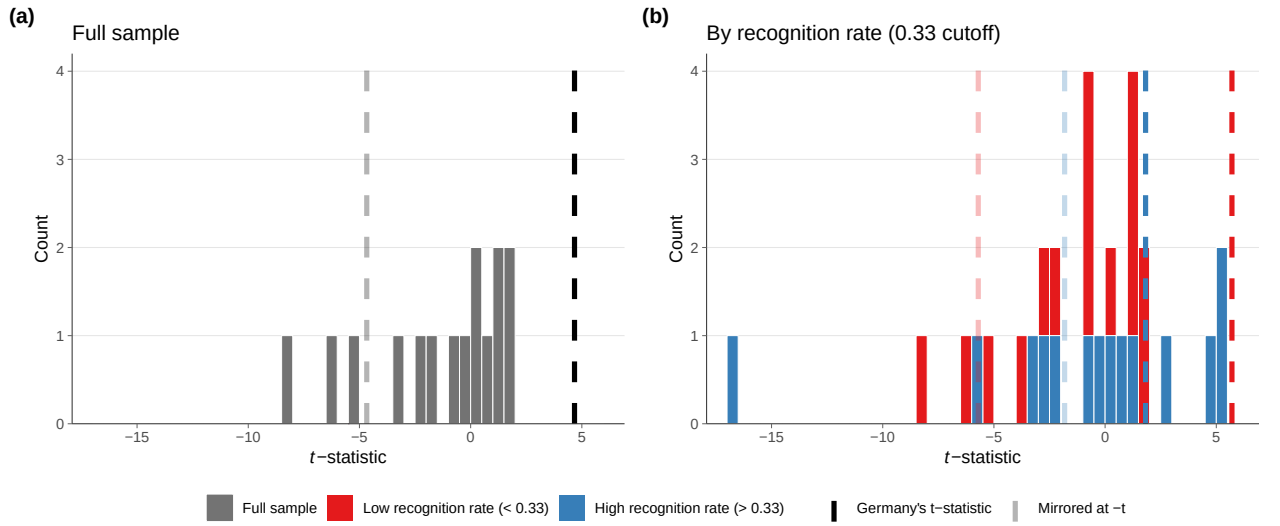
The minimum-application requirement restricts the pool because a placebo draw is informative only if it is a plausible realization of the null for a destination of Germany's size, and the smallest destinations are not. Several receive a few hundred first-time applications over the whole pre-period and the smallest fewer than a hundred, so destination $\times$ origin cells are frequently empty and the two-way clustered Poisson variance is estimated off almost no mass. The $t$-statistics they generate are numerical artifacts an order of magnitude larger than Germany's, not small-sample noise around zero,

and in some origin subsamples the clustered variance matrix becomes near-singular and the statistic diverges outright. The threshold is nonetheless a judgment call and it is not costless: restricting the pool to 15 destinations coarsens the test, since the smallest attainable $p$-value is the reciprocal of the number of draws.

The recognition-rate split in panel (b) of Figure A12 is the more informative half. Germany's statistic separates from the placebo draws among the low-recognition-rate origins that drive the heterogeneity result, but not among the high-recognition ones, which is the pattern the mechanism implies. The few draws that exceed Germany in absolute value are small destinations that changed policy over the same window, so the placebo distribution is wider than the null it stands for and the $p$-values are conservative, exactly as the variance diagnostic of Appendix C.1 anticipates.

Two further features keep the placebo draws from being clean realizations of the null, both set out in Section 4.5: diversion leaves the control destinations heterogeneously treated rather than untreated, and keeping Germany in the panel as an ordinary control biases every reassigned coefficient downwards. Both are visible in the exhibits below, and together they are why we read these $p$-values as descriptive of the placebo distribution rather than as the size of a test.

Figure A12: Permutation inference: placebo distribution of $t$-statistics




*Notes:* Each draw reassigns the July 2012 treatment indicator to one non-German destination and re-estimates equation (6); the statistic recorded is the $t$-ratio on the reassigned indicator, following MacKinnon and Webb (2020). Germany's own observations stay in the panel as part of the control group in every draw. The pool is the 15 non-German destinations with at least 1,000 applications over the eleven pre-ruling months; smaller destinations are omitted because the two-way clustered Poisson variance is estimated off too little mass to be informative, for the reason set out above. The darker dashed line marks Germany's own $t$-statistic and the faded line its mirror image at $-t$, so draws more extreme in either direction can be read off directly, and two-sided $p$-values are the share of draws exceeding Germany's in absolute value. Panel (a) uses the full origin sample ($t = 4.67$, $p = 0.200$). Panel (b) splits origins at a recognition rate in Germany of 0.33, holding the pool fixed: $t = 5.70$, $p = 0.133$ among low-recognition origins and $t = 1.82$, $p = 0.643$ among high-recognition ones, with one high-recognition draw dropped for a degenerate clustered variance matrix. Both panels share a horizontal axis, bin grid and vertical scale.

The screen is a judgment call, so it is worth showing what it removes rather than asserting that it removes only noise. Figure A13 repeats the exercise over all 24 non-German destinations, and Table A8 lists every destination whose placebo $|t|$ exceeds Germany's, which is exactly the set the two-sided $p$-value counts.

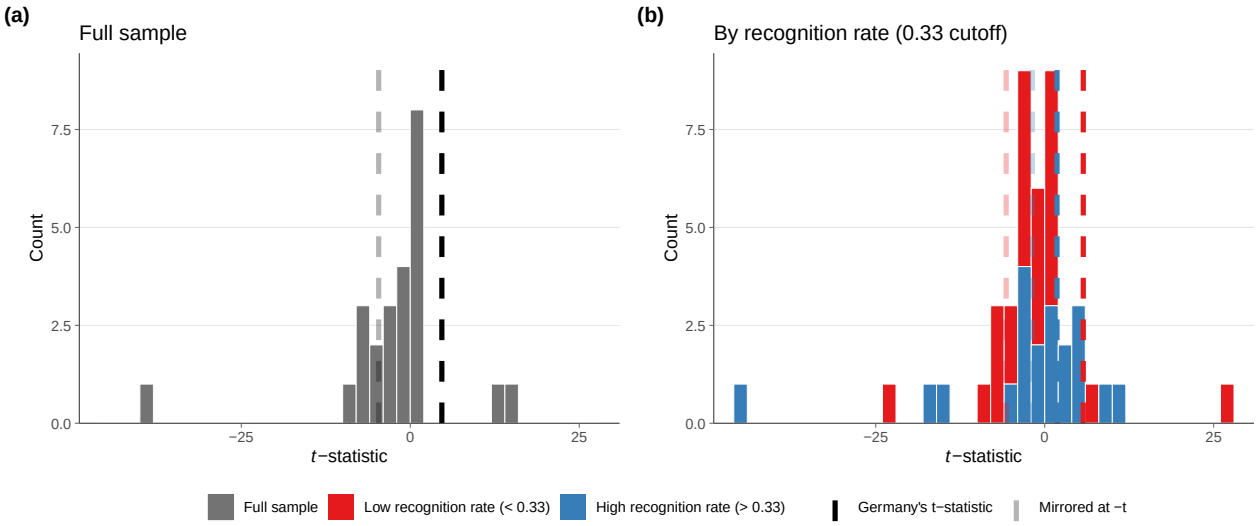
Two things are visible at once. Lifting the screen moves the $p$-value from 0.200 to 0.333 on the full sample and from 0.133 to 0.292 among the low-recognition-rate origins, so the additional draws are not

marginal. But they are also not draws one would want to condition on. Of the eight destinations that overtake Germany in Panel A, six received fewer than 2,500 applications over the entire eleven-month pre-ruling window, against Germany's 43,800, and the two most extreme, Latvia at $t = -39.4$ on 250 applications and Bulgaria at $t = +14.4$ on 495, sit an order of magnitude beyond anything the screened pool produces. Their $t$-statistics measure the thinness of the panel rather than the size of a plausible spurious effect; four further draws in the high-recognition half degenerate outright and are dropped.

The two destinations that clear the volume threshold and still overtake Germany are the ones Section 4.4 already identifies, and neither is a null draw either. Belgium bundled deterrence measures over 2011–2012. Luxembourg is the sharper case, because it moved on the same margin as Germany and in the same month: the grand-ducal regulation of 8 June 2012 replaced the schedule that had been in force since 2006 and cut the monthly cash allowance paid to applicants for international protection.[24] For applicants housed without meals provided, the allowance fell from EUR 294.00 to EUR 225.00 for a single adult, a cut of 23.5 percent, and from EUR 534.15 to EUR 300.00 for a two-person household, a cut of 43.8 percent. Where meals are provided, which is the arrangement the regulation treats as the rule, an adult's allowance fell from EUR 107.90 to EUR 25.00, a cut of 76.8 percent. Romania, just below the threshold, is unwinding the one-off Arab Spring pulse of 2011. These are genuine policy responses in the control group, which is a violation of the null rather than evidence for it.

[24] *Règlement grand-ducal du 8 juin 2012 fixant les conditions et les modalités d'octroi d'une aide sociale aux demandeurs de protection internationale*, Mémorial A no. 123 of 20 June 2012, p. 1586, whose Article 11 repeals the *Règlement grand-ducal du 1er septembre 2006* of the same title, Mémorial A no. 171 of 22 September 2006, p. 3109. Both schedules are compared at their published values, which is how the Luxembourg national contact point of the European Migration Network reports the reform in its 2012 policy report, describing the allowance for an adult accommodated with meals as falling "de 107,9 à 25 euros".

Figure A13: Permutation inference without the volume screen on the placebo pool




*Notes:* As Figure A12, except that the placebo pool is every non-German destination in the estimation sample rather than only those with at least 1,000 applications in the eleven-month pre-ruling window. Each draw reassigns the July 2012 treatment indicator to one control destination and re-estimates equation (6); the statistic recorded is the $t$-ratio on the reassigned indicator, following MacKinnon and Webb (2020). Germany's own observations stay in the panel as part of the control group in every draw. The darker dashed line marks Germany's own $t$-statistic and the faded line its mirror image at $-t$, so draws more extreme in either direction can be read off directly. Panel (a) uses the full origin sample; panel (b) splits origins at a recognition rate in Germany of 0.33, holding the pool fixed across the two halves. Draws whose two-way clustered variance matrix is near-singular produce four- and five-digit $t$-statistics that are numerical artifacts rather than null draws; those with $|t| > 50$ are dropped, which removes four draws from the high-recognition half and none elsewhere. Both panels share a horizontal axis, a bin grid and a vertical scale.

Table A8: Destinations whose placebo estimate is more extreme than Germany's

| | Coefficient | $t$-statistic | Pre-ruling applications | Reason |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| *Panel A. All origins* | | | | |
| Germany (treated) | 0.476 | 4.67 | 43,800 | — |
| Latvia | -1.374 | -39.37 | 250 | Thin panel |
| Bulgaria | 1.188 | 14.40 | 495 | Thin panel |
| Slovenia | 0.415 | 12.85 | 130 | Thin panel |
| Belgium | -0.695 | -8.37 | 20,795 | Deterrence package |
| Estonia | -0.452 | -7.04 | 70 | Thin panel |
| Lithuania | -0.639 | -6.85 | 500 | Thin panel |
| Romania | -0.836 | -6.20 | 2,230 | Arab Spring pulse unwinds |
| Luxembourg | -1.179 | -5.10 | 1,985 | Reception benefits cut |
| *Panel B. Low-recognition-rate origins ($R_o < 0.33$)* | | | | |
| Germany (treated) | 0.695 | 5.70 | 23,295 | — |
| Bulgaria | 2.392 | 27.50 | 65 | Thin panel |
| Latvia | -1.465 | -22.59 | 225 | Thin panel |
| Romania | -1.221 | -8.36 | 1,765 | Arab Spring pulse unwinds |
| Slovenia | 0.474 | 7.24 | 70 | Thin panel |
| Lithuania | -0.682 | -7.06 | 425 | Thin panel |
| Estonia | -0.490 | -6.67 | 60 | Thin panel |
| Belgium | -0.715 | -6.37 | 15,115 | Deterrence package |

*Notes:* Every destination whose placebo $|t|$ is at least as large as Germany's, from the unscreened pool of Figure A13, ordered by $|t|$; these are exactly the draws the two-sided $p$-value counts. Germany's own estimate heads each panel for comparison and is not a placebo draw. Coefficients are PPML estimates of equation (6) with the treatment reassigned to that destination, and applications are counted over the eleven-month pre-ruling window on that panel's own origin sample, so Panel B reports low-recognition-rate applications rather than each destination's overall total. *Reason* classifies why the draw is extreme. "Thin panel" marks destinations with fewer than 1,000 pre-ruling applications, where the two-way clustered Poisson variance is estimated off too little mass for the $t$-statistic to be informative. The three remaining cases are documented policy or route changes over the estimation window, discussed in Section 4.4: Belgium combined dissuasion campaigns, a 2011–2012 safe-country-of-origin list that fast-tracked the Western Balkan caseload, and a reception-capacity crisis that curtailed material support; Luxembourg cut the monthly cash allowance paid to applicants for international protection by 23.5 percent for a single adult and 43.8 percent for a two-person household in June 2012, in the same month as the German ruling and in the opposite direction; and Romania's decline reflects the unwinding of a one-off 2011 pulse of North African applications on the Serbia–Romania corridor rather than any policy change.

# D Derivations

This appendix derives (i) the diversion-bias correction used to bound the application elasticity, and (ii) the closed-form expression for the implied diversion share $\hat{\rho}$ used in the stable-origins corroboration (Section 7.2).

## D.1 Random Utility Model

Recall from Section 4.1 that the utility applicant $i$ from origin $o$ derives from destination $d$ at time $t$ is $V_{iodt} = W_{odt} + \varepsilon_{iodt}$, with deterministic component $W_{odt} = \alpha\, b_{dt} + \mu_{odt}$ and the outside option of staying valued at $W_{oot}$ (equations (1) and (2)). The simplest assumption on $\varepsilon_{iodt}$ is that the draws are i.i.d. across alternatives $d$, which is unrealistic in our setting. EU+ destinations have comparable economic conditions, share a common asylum framework, broadly comparable reception conditions, and free travel after recognition among Schengen countries; hence, an individual considering applying for asylum in Germany typically treats several other EU+ destinations as close substitutes much more readily than the stay option.[25] In other words, an applicant deterred from applying in one EU+ country is plausibly more likely to apply in another EU+ destination.

We allow for this by letting $\varepsilon_{iodt}$ follow the two-level generalised extreme value distribution of McFadden (1978), introduced to migration destination choice by Bertoli and Fernández-Huertas Moraga (2013): all EU+ destinations sit in a single nest with dissimilarity parameter $\tau \in (0, 1]$, while the stay option occupies a degenerate singleton nest. Within the migrate nest, taste shocks are correlated according to

$$\operatorname{Corr}(\varepsilon_{iodt}, \varepsilon_{iod't}) \;=\; 1 - \tau^2, \qquad d, d' \in \mathcal{D},\; d \neq d'. \tag{A1}$$

Small $\tau$ describes a setting in which EU+ destinations are close substitutes (an applicant deterred from Germany tends to apply elsewhere in the EU+ rather than stay home), and $\tau \to 1$ collapses the model to the i.i.d. (multinomial logit) case.

## D.2 Marginal-substitution patterns

A small benefit increase in destination $d'$ pulls probability mass off the other EU+ destinations and the stay margin and into $d'$. Suppressing origin and time subscripts and writing $S \equiv \sum_{k \in \mathcal{D}} m_k$ for the migrate share and $m_0 \equiv 1 - S$ for the stay share, differentiation of (3) gives the semi-elasticities

$$\frac{\partial \ln m_d}{\partial b_{d'}} \;=\; \alpha \left[ \frac{\mathbb{1}[d = d']}{\tau} \;-\; \frac{1 - \tau}{\tau}\, \frac{m_{d'}}{S} \;-\; m_{d'} \right], \qquad d \in \mathcal{D}, \tag{A2}$$

$$\frac{\partial \ln m_0}{\partial b_{d'}} \;=\; -\,\alpha\, m_{d'}. \tag{A3}$$

The bracketed expression in (A2) decomposes into three terms:

(i) the *direct response* $1/\tau$, active only for the affected destination ($d = d'$); pass-through is stronger when destinations are closer substitutes (smaller $\tau$);

[25] As we only observe asylum migration at monthly frequency in the EU+, the outside option includes all other destinations outside the EU+.

(ii) the *within-nest substitution* $-(1-\tau)/\tau \cdot m_{d'}/S$, the loss in conditional share suffered by every destination in the migrate nest as the nest's inclusive sum rises; it vanishes at $\tau = 1$;

(iii) the *outer-nest rebalancing* $-m_{d'}$, the proportional squeeze on every alternative—including the stay option, as the same coefficient also governs (A3)—as the migrate-vs-stay split tightens in favor of migration.

For a competing EU destination ($d \neq d'$) the direct term drops out, and the two negative terms send $m_d$ down; for the stay margin only the outer-nest rebalancing applies, but it likewise sends $m_0$ down. Both losses accrue to $d'$.

The structural diversion share is the fraction of marginal applicants drawn into $d'$ that comes from other EU+ destinations rather than from the stay margin,

$$\rho \;\equiv\; \frac{\sum_{d\neq d',\, d\in\mathcal{D}} \big(-\partial m_d/\partial b_{d'}\big)}{\partial m_{d'}/\partial b_{d'}}. \tag{A4}$$

The utility coefficient $\alpha$ enters every derivative symmetrically and cancels. Substituting (A2)–(A3) (converted to level derivatives via $\partial m_d/\partial b_{d'} = m_d\,\partial \ln m_d/\partial b_{d'}$), applying adding-up, and dividing through by $m_{d'}$ delivers the closed form in terms of the pre-ruling application ratio

$$\kappa \;\equiv\; \frac{\sum_{k\neq d'} m_k}{m_{d'}} \;\approx\; 4.8, \tag{A5}$$

calibrated to Germany's pre-ruling share of EU asylum applications ($M_0^T/M_0^C \approx 0.21$, so $\kappa = 1/0.21 \approx 4.8$). The closed form is

$$\rho \;=\; \frac{\kappa\,\big[\,1-\tau(1-S)\,\big]}{\kappa \;+\; \tau\,(1-S)}. \tag{A6}$$

Two features of (A6) matter for what follows. First, $\rho$ depends only on $\tau$ and the observable shares $\kappa$ and $S$: relating an empirical $\hat{\rho}$ to the structural $\tau$ requires no estimate of $\alpha$. Second, only a small share of any origin population applies for asylum in the EU, so $1 - S \approx 1$ and (A6) reduces to $\rho \approx \kappa(1-\tau)/(\kappa+\tau)$. At $\kappa = 4.8$ this gives $\rho \to 1$ as $\tau \to 0$ (full diversion: marginal applicants come entirely from other EU+ destinations), $\rho \approx 0.45$ at $\tau = 0.5$, and $\rho = 0$ at $\tau = 1$ (no diversion: marginal applicants come entirely from the stay margin, the multinomial-logit limit). Inverting (A6) writes the dissimilarity parameter as a function of the diversion share and the observable shares,

$$\tau \;=\; \frac{\kappa\,(1-\rho)}{(1-S)\,(\kappa+\rho)}, \tag{A7}$$

which reduces to $\tau \approx \kappa(1-\rho)/(\kappa+\rho)$ for $1-S \approx 1$. Given an estimate of $\rho$ and the observable shares, this point-identifies $\tau$ (and hence $\alpha = \beta\tau$).

### D.3 Characterizing the policy-relevant treatment effect under diversion

Consider two periods and two units (Germany $T$ and the rest of the EU+ $C$). Let $M_{0/1}^{T/C}$ denote applications, $g$ the common secular growth rate, $\delta$ the proportional treatment effect for Germany, and $\rho_{D2} \in [0,1]$ the share of marginal applications in Germany being diverted away from the control group countries. Then

$$M_1^T = M_0^T\,(1+g+\delta), \quad M_1^C = M_0^C(1+g) - (M_0^T)\rho_{D2}\delta,$$

and the Poisson DiD estimator identifies

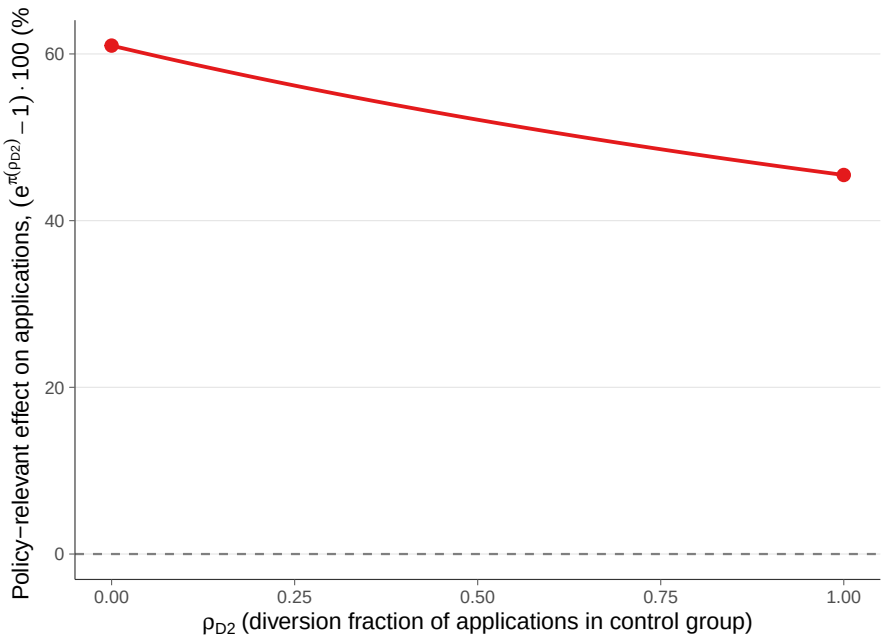
$$e^{\hat{\beta}} \;=\; \frac{1+g+\delta}{1+g-\left(M_0^T/M_0^C\right)\rho_{D2}\delta}. \tag{A8}$$

Here $\delta$ is the structural per-period treatment effect and the policy-relevant parameter is $\pi$, defined by $e^{\pi} \equiv 1+\delta/(1+g)$. Importantly, if secular growth rates are positive, the increase relative to the counterfactual $(e^{\pi}-1)$ is smaller than $\delta$ as the latter signifies the growth rate with respect to the smaller baseline level. Absent diversion from control countries ($\rho_{D2}=0$) the estimator yields $e^{\hat{\beta}} = e^{\pi}$, whereas $\rho_{D2}>0$ shrinks the denominator of (A8) and inflates $e^{\hat{\beta}}$ above $e^{\pi}$. To isolate $\pi$, cross-multiply (A8): $e^{\hat{\beta}}\big(1+g-(M_0^T/M_0^C)\rho_{D2}\delta\big) = 1+g+\delta$. Expanding and collecting the terms in $\delta$ gives $(1+g)(e^{\hat{\beta}}-1) = \delta(1+e^{\hat{\beta}}(M_0^T/M_0^C)\rho_{D2})$, hence $\delta/(1+g) = (e^{\hat{\beta}}-1)/(1+e^{\hat{\beta}}(M_0^T/M_0^C)\rho_{D2})$. Adding one and using $e^{\pi} \equiv 1+\delta/(1+g)$ gives

$$e^{\hat{\pi}} \;=\; 1 + \frac{e^{\hat{\beta}}-1}{1+e^{\hat{\beta}}\,\dfrac{M_0^T}{M_0^C}\,\rho_{D2}}. \tag{A9}$$

$g$ drops out: the corrected effect is identified from $\hat{\beta}$, $M_0^T/M_0^C$, and $\rho_{D2}$. Figure A14 plots (A9) at $e^{0.476}$ and $M_0^T/M_0^C \approx 0.21$.

Figure A14: Policy-relevant treatment effect as a function of $\rho_{D2}$




*Notes:* Policy-relevant treatment effect $e^{\hat{\pi}}-1$ from equation (8) as a function of the diversion share $\rho_{D2} \in [0,1]$, evaluated at $e^{0.476}$ and $M_0^T/M_0^C \approx 0.21$.

## D.4 Estimating the diversion share

The ruling-induced increase in applications in Germany is $\Delta M^T = A + D1 + D2$ (raw, uncorrected: every marginal applicant of any type files in Germany). Generalizing the onward-migration decom-

position (9) to an arbitrary destination $d$ gives $M_{d,t} = N_{d,t} + \sum_{s \neq d} O^s_{d,t}$, where $N_{d,t}$ counts applicants with no prior EU+ registration. Summing across destinations and subtracting the inward-hit term destination-by-destination recovers $M^{EU,\text{net}}_t = \sum_d N_{d,t}$, the count of unique new entrants to the EU+ asylum system that month. A type-D1 applicant's repeat filing is, by construction, an inward hit at whichever destination it lands, so it is removed wherever it appears and does not enter $\Delta M^{EU,\text{net}}$. A type-D2 applicant substitutes one $N$-event for another—lost at the control country they would otherwise have filed in, gained in Germany—and so nets to zero across the sum. Only type-A applicants, who would not have entered the EU+ asylum system at all absent the reform, add a new $N$-event. Hence

$$\Delta M^{EU,\text{net}} = \Delta A, \qquad \Delta M^{T} = \Delta A + \Delta D1 + \Delta D2, \qquad \Longrightarrow \qquad \hat{\rho} = \frac{\Delta D1 + \Delta D2}{\Delta M^{T}} = 1 - \frac{\Delta M^{EU,\text{net}}}{\Delta M^{T}}.$$

The absolute increase in net EU+ applications is $\Delta M^{EU,\text{net}} = M_0^{EU,\text{net}}(e^{\hat{\beta}_{EU}} - 1)$, and the absolute increase in applications in Germany is the raw PPML effect, $\Delta M^{T} = M_0^{T}(e^{\hat{\beta}_{T}} - 1)$. This gives the closed form used in the text:

$$\hat{\rho} = 1 - \frac{M_0^{EU,\text{net}}\left(e^{\hat{\beta}_{EU}} - 1\right)}{M_0^{T}\left(e^{\hat{\beta}_{T}} - 1\right)}, \qquad \hat{\rho}_{D2} = \hat{\rho} - \hat{\rho}_{D1}. \tag{A10}$$

$\hat{\rho}$ is identified from $(M_0^{EU,\text{net}}, M_0^{T}, \hat{\beta}_{EU}, \hat{\beta}_{T})$ alone; $\hat{\rho}_{D1}$, the Eurodac lower bound of Section 7.1, enters only to split $\hat{\rho}$ into its D1 and D2 components, not to identify $\hat{\rho}$ itself.

Constructing $M^{EU,\text{net}}_{d,t} = M_{d,t} - \sum_{s \neq d} O^s_{d,t}$ requires inward-hit counts by destination and month, but the Eurodac source data records hits only by receiving and sending EU+ country, not by applicant nationality. We therefore compute the inward-hit share $\left(\sum_s O^s_{d,t}\right)/M_{d,t}$ at the destination-month level, pooling all nationalities, and apply it uniformly to every origin's applications in that destination-month—an approximation forced by the coarser granularity of the Eurodac data relative to the Eurostat data it corrects. A destination-month whose inward-hit count exceeds its Eurostat total (which arises for a minority of country-months, concentrated in a handful of smaller destinations) is treated as fully repeat ($N_{d,t} = 0$) rather than allowed to go negative, mirroring the floor used in the Eurodac decomposition of Section 7.1. Destination-months absent from the Eurodac file altogether are dropped from both the net EU+ sum and, for comparability, from the Germany-specific control group, so that $\hat{\beta}_T$ and $\hat{\beta}_{EU}$ are estimated on one common destination panel.

The two coefficients are estimated on the respective stable-origin sample: $\hat{\beta}_T$ is the main difference-in-differences coefficient of (6) (Germany versus the EU+ controls, with destination × origin and origin × year-month fixed effects), while $\hat{\beta}_{EU}$ is the post-July-2012 coefficient from a Poisson regression of *net* EU+ applications from those origins (summed across Eurodac-covered destinations) on the post indicator, with origin and calendar-month-of-year fixed effects.

# E Additional Results

## E.1 Anecdotal attraction: Western Balkans and Russian Chechens

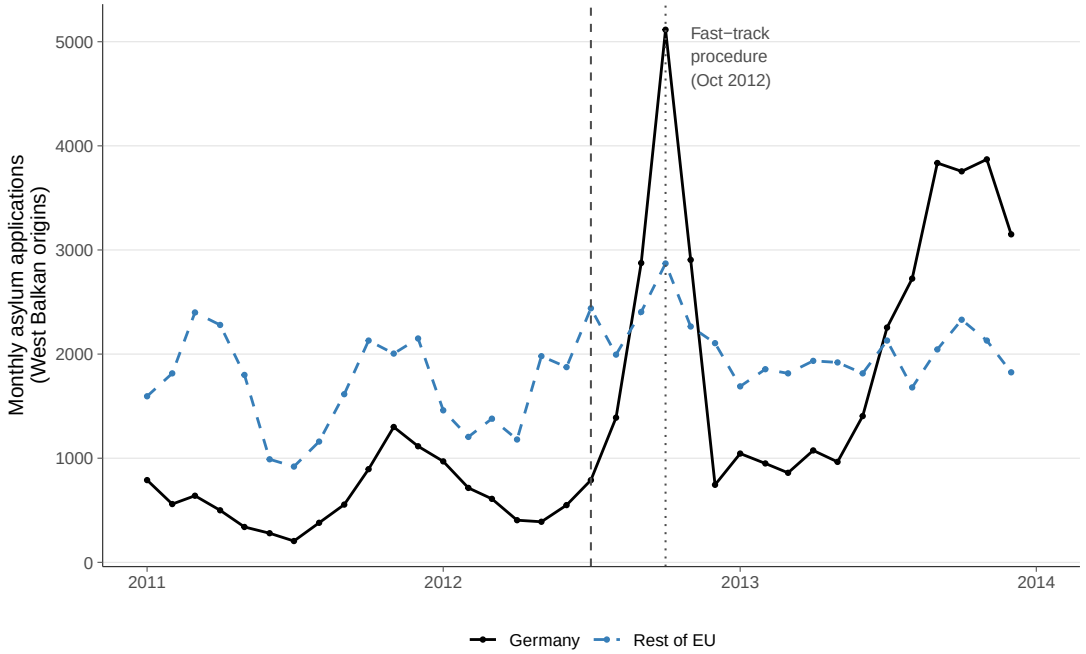
This appendix documents in detail the two origin groups, summarized in Section 7.3, for which the data show attraction rather than diversion.

**Western Balkans.** For Western Balkan origins (Serbia, Kosovo, Bosnia–Herzegovina, Albania, Montenegro, North Macedonia), Figure A15a shows that applications to the rest of the EU+ rise in parallel with applications to Germany after the July 2012 ruling. Under pure diversion, applications to the rest of the EU+ would have *declined* as Western Balkan applicants reallocated to Germany; the parallel rise instead indicates that the ruling raised the expected value of EU asylum migration above the reservation utility of staying home for previously infra-marginal Western Balkan economic migrants. This is exactly the population for whom emigration is a viable outside option—they faced no enforcement obstacles to travel and had near-zero protection prospects. Visa-free short stays in the Schengen area applied from 19 December 2009 for citizens of North Macedonia, Montenegro, and Serbia and from 15 December 2010 for citizens of Albania and Bosnia–Herzegovina, while citizens of Kosovo remained subject to a visa requirement throughout the sample period (Council of the European Union, 2009; Alscher, Obergfell and Roos, 2015). Because the exemption was set in EU law, it applied on the same dates in every EU member state except the United Kingdom and Ireland, which kept their own visa requirements, and therefore also in the transit countries Slovenia and Austria. Germany thus offered no easier entry than its neighbors, yet applications concentrated there: in 2013 Germany received more than 60 percent more Western Balkan applications than France, the second-largest recipient (Alscher, Obergfell and Roos, 2015). Most applicants belonged to the Roma minority (Alscher, Obergfell and Roos, 2015), which throughout the region has poor access to education, employment, and housing and faces a high risk of poverty; among decided applications in Germany, 85 percent of applicants from Serbia and 70 percent from North Macedonia were Roma (Müller, Reinig and Klingholz, 2017). Applications followed a seasonal pattern, rising every autumn and winter and falling in spring and summer, which both sources attribute to precarious living conditions at home that are hardest to bear in winter (Alscher, Obergfell and Roos, 2015; Müller, Reinig and Klingholz, 2017). The post-October 2012 Germany-only decline coincides with the BAMF fast-track procedure documented in Appendix I, indicating that the German government recognized the attraction margin and responded by reducing the discounted benefit gain for this group via shorter procedures rather than via the benefit level itself. Two further instruments targeted the same margin. Return assistance for Western Balkan nationals was suspended or reduced (Alscher, Obergfell and Roos, 2015). Benefits could also be cut under § 1a AsylbLG, but before October 2015 this provision covered only persons who were tolerated or obliged to leave and had entered Germany in order to obtain benefits, so it reached Western Balkan applicants only after their claims had been rejected. No systematic data exist on how widely either measure was applied; the federal government, for instance, holds no figures on the number of persons receiving reduced benefits under § 1a (**?**).

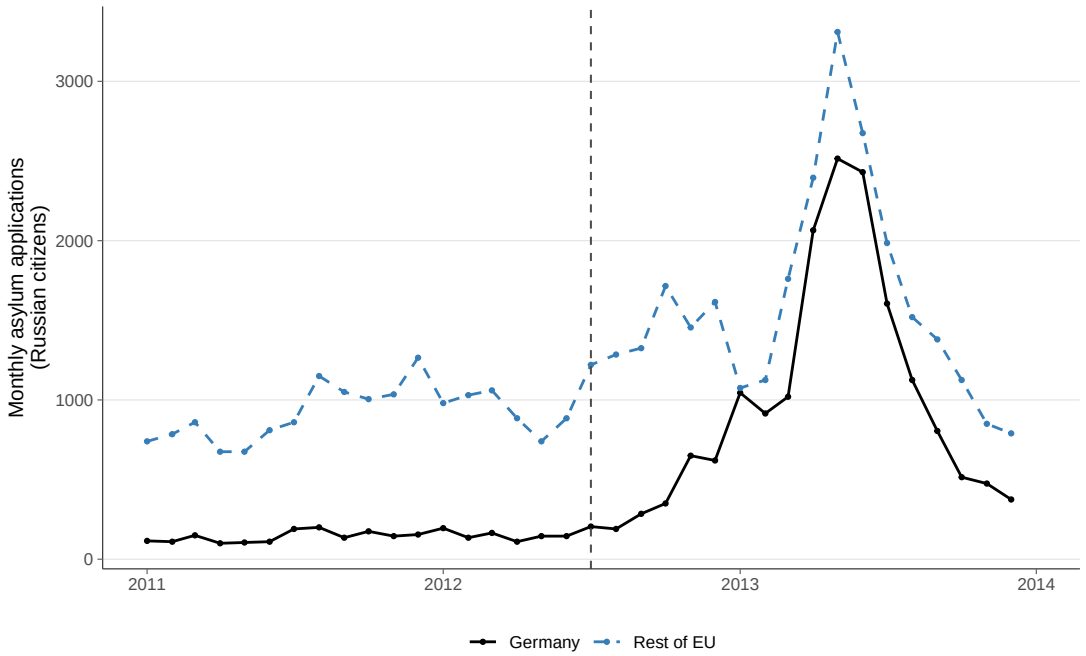
**Russian Chechens.** For Russian citizens (mostly Chechen; Figure A15b, panel (b)), the German rise post-2012 precedes a delayed rise in the rest of the EU+. The temporal sequence is consistent with two compounding mechanisms. First, attraction: a Chechen population already in Poland under temporary protection or tolerated status responded to the German ruling by lodging new applications in Germany. Second, onward migration through the Polish first-application channel: many Chechens who had reached the EU via Poland under the pre-existing Russian visa regime moved on to Germany after the ruling, generating Category 1–Category 1 inward hits from Poland. The Eurodac sender-group decomposition (Section 7.1) directly confirms the second channel: Polish-sender inward hits show the largest proportional post-ruling increase of any sender group ($\hat{\beta} = 1.969$). The Polish press

extensively documented both mechanisms in the 2012–2013 period.

Figure A15: Applications from attraction-prone origins: Germany vs. rest of EU+




(a) Western Balkan origins (RS, XK, BA, AL, ME, MK)




(b) Russian citizens (mostly Chechen)

*Notes:* Monthly first-time applications to Germany (solid) and to the rest of the EU+ (dashed), January 2011–December 2013. *Panel (a):* Western Balkan origins; the parallel post-July 2012 rise points to attraction, not diversion, and the post-October 2012 decline in Germany coincides with the BAMF fast-track (Appendix I). *Panel (b):* Russian citizens; the German rise post-2012 precedes a delayed rise in the rest of the EU+, consistent with attraction combined with onward migration via the Polish first-application channel.

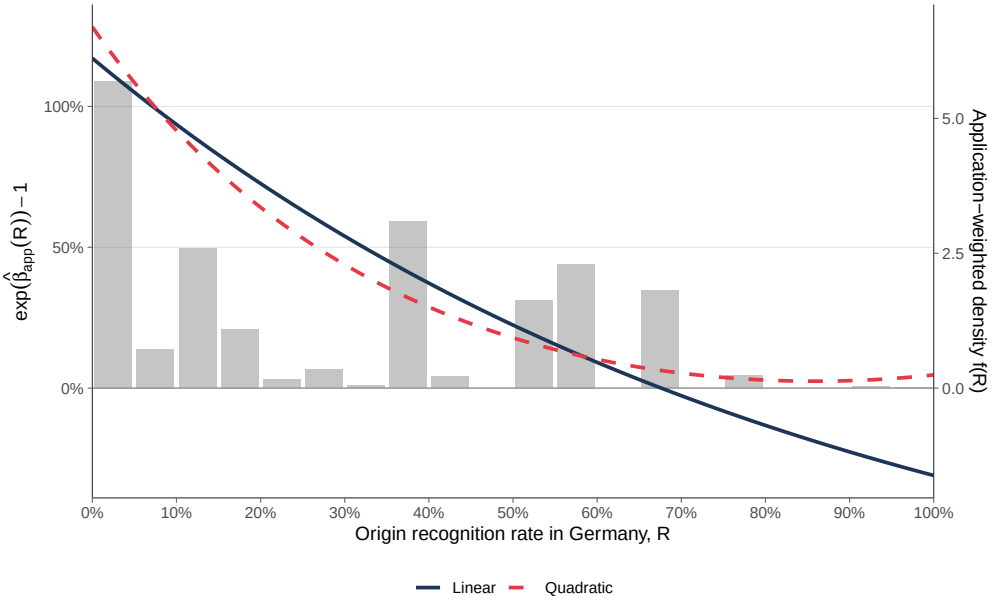
**Translating applications into granted protection.** Let $\hat{\beta}^{\text{app}}(R) = 0.78 + -1.15\,R$ denote the bivariate linear interaction of the post-treatment indicator with the recognition rate, estimated on the

main DiD sample. Let $f(R_o) = M_o / \sum_{o'} M_{o'}$ be the application-weighted empirical distribution of $R_o$ across origins, with weights equal to total Germany applications from origin $o$ over August 2011–June 2012. Assuming that the recognition rate of the marginal applicant from origin $o$ equals the average $R_o$ and that within-origin recognition rates are unchanged, the implied proportional increase in granted refugees is

$$\frac{\Delta G}{G_{\text{base}}} = \frac{\sum_o \left(e^{\hat{\beta}^{\text{app}}(R_o)} - 1\right) R_o \, f(R_o)}{\sum_o R_o \, f(R_o)}. \tag{A11}$$

Most application mass lies at low $R$, where the proportional response is large; the integrand for granted refugees weights this mass by $R$, shifting the relevant range upward. Numerically, equation (A11) yields a predicted increase of 26 percent in granted refugees against an aggregate application increase of 63 percent. The implied refugee-benefit elasticity (% change in refugees over 54 percent benefit rise) is 0.48. Figure A16 plots the integrand kernel and $f(R)$.

Figure A16: Translating the application elasticity into granted refugees

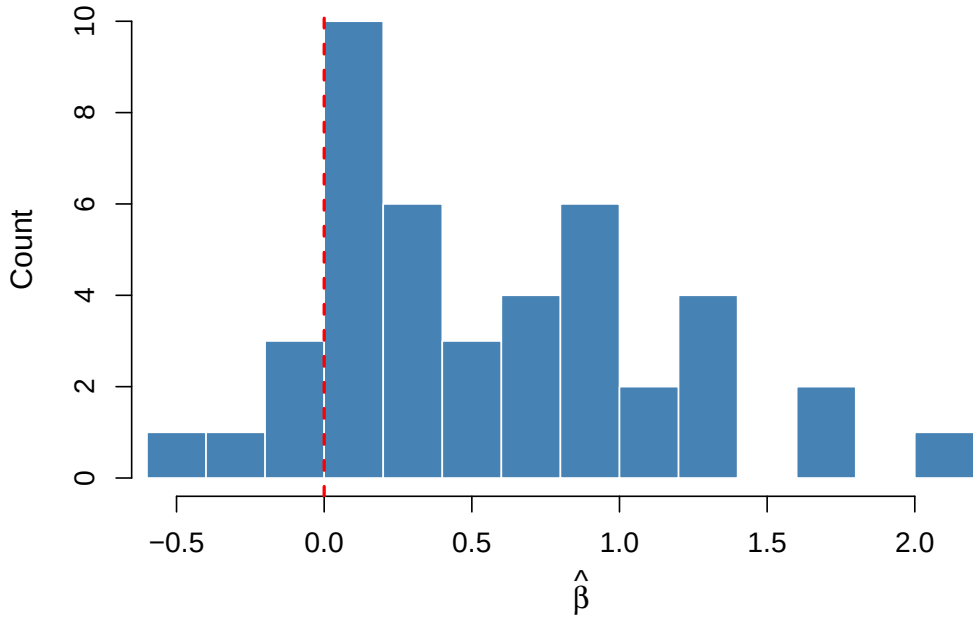



*Notes:* Dark line: proportional post-ruling increase in applications, $\exp(\hat{\beta}^{\text{app}}(R)) - 1$. Shaded histogram (right axis): application-weighted empirical distribution of origin-specific Germany recognition rates $R_o$. Quadratic and bin functional forms (Appendix Table A18) yield similar implied refugee responses (24 percent and 25 percent, respectively).

**Distribution of per-origin effects.** The headline coefficient averages over origins that respond very differently. Figure A17 plots the distribution of the per-origin DiD coefficients, which makes clear that the aggregate effect reflects a subset of strongly responding origins rather than a uniform shift across all of them.

**Visa-free travel within the low-recognition-rate group.** Section 5.1 shows two things separately: that the response is concentrated among origins with a low recognition rate in Germany, and that it is larger for origins whose nationals can reach the Schengen area without a visa. Because the two characteristics overlap heavily, it is worth asking whether they compound. Figure A18 re-estimates the event study within the low-recognition group alone, splitting it by visa status. The visa-free origins

Figure A17: Distribution of estimated effects by origin country

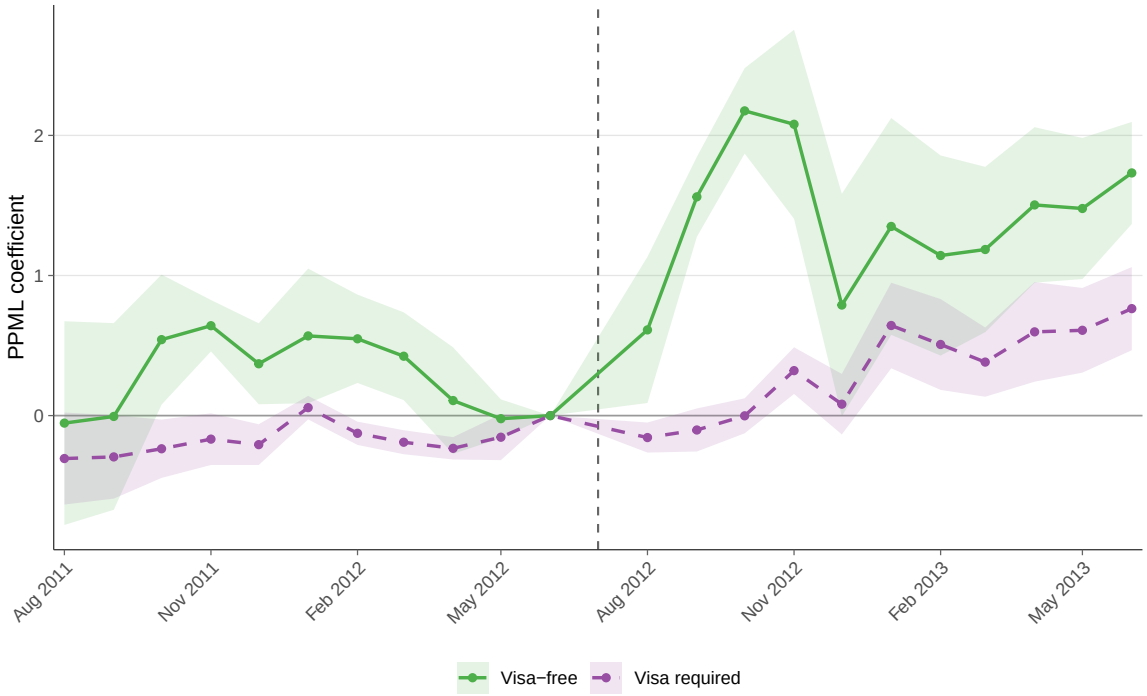



*Notes:* Per-origin PPML coefficients. Sample restricted to origins with $\geq$ 1,000 first-time applications EU-wide in the 11-month pre-period. Haiti ($\hat{\beta} = 19.14$) is excluded from the histogram as an extreme outlier; its point estimate is reported here.

respond several times more strongly than the visa-required ones, and they do so immediately, jumping in the first post-ruling month; the visa-required series is flat through the autumn of 2012 and builds only from the winter onwards. The natural reading is that a low expected protection prospect and a low cost of getting to Germany are complements rather than substitutes: the ruling raises the value of applying, and the origins that can act on it fastest are those whose nationals do not first have to obtain a visa.

Two cautions apply. The visa-free cell holds only 9 origins against 61 visa-required ones, so the visa-free series is estimated on a narrow base and its confidence band is correspondingly wide. And the mirror-image exercise is not available: of the 10 high-recognition origins, only 1 is visa-free, and it contributes a negligible number of applications, so there is no way to check whether the same visa gradient appears among origins with strong protection prospects. What the figure establishes is therefore that the low-recognition response is not driven solely by visa-free origins, since the visa-required series is itself large and precisely estimated; whether the interaction is genuinely second-order, or simply reflects that visa-free access and low recognition rates travel together, cannot be settled with this variation.

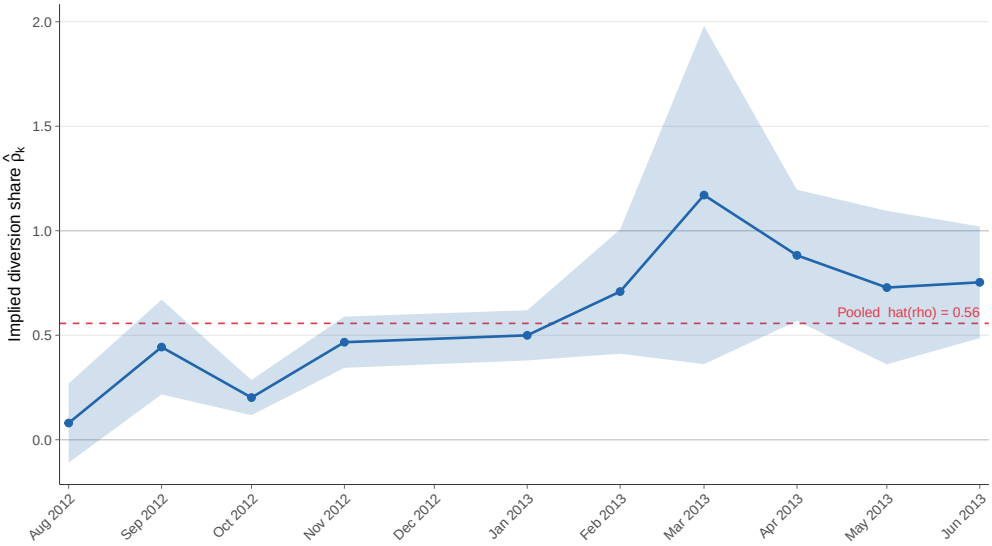
Figure A18: Event study within the low-recognition-rate group, by visa status




*Notes:* PPML event-study estimates from equation (7), restricted to origins with $R_o < 0.33$ and estimated separately by whether the origin's nationals required a visa for Germany in 2012 (DEMIG VISA inward indicator, the same variable as column (3) of Table 3). Visa-free: 9 origins. Visa required: 61 origins. Destination × origin and origin × year-month fixed effects; standard errors clustered at the origin and destination country level; shaded areas are 95 percent confidence intervals. The dashed vertical line marks the ruling (July 2012).

# F Diversion and attraction: additional exhibits

Figure A19: Implied diversion share lower bound over post-ruling period




*Notes:* Per-period $\hat{\rho}_k = \Delta O_k/(\Delta O_k + \Delta N_k)$, where $O$ are onward applicants (previously applied in another EU country) and $N$ are first-time applicants who had not. Shaded band: 95 percent CI (delta method). Dashed line: pooled $\hat{\rho} = 0.56$. Dec 2012 omitted ($\hat{\rho}_5$ explodes).

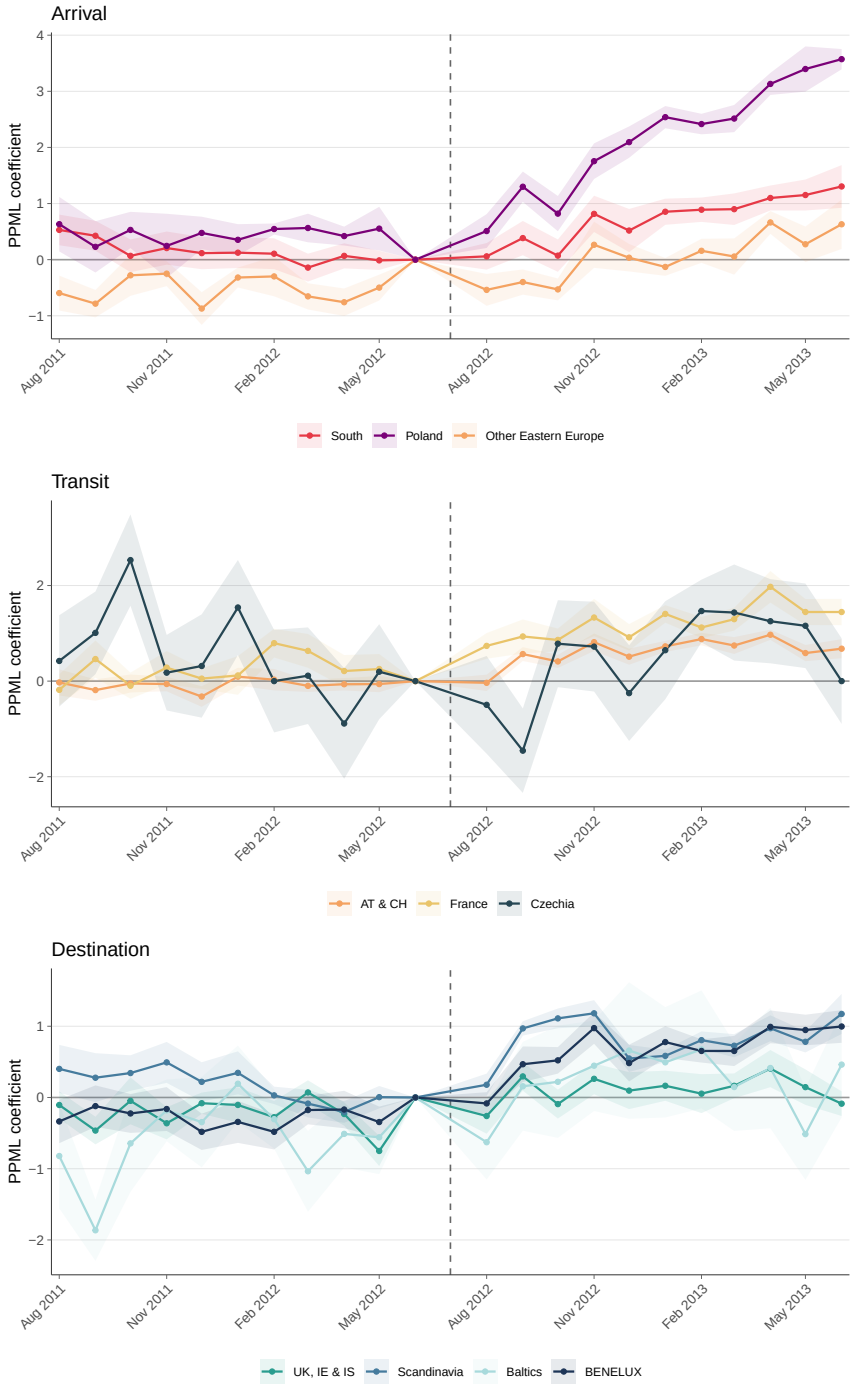
**Eurodac sender-group DiD table.**

Table A9: Eurodac decomposition: single-coef DiD table

| | Never | All inward Cat1 | Arrival | Transit | Destination |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| DiD $\hat{\beta}$ | 0.302*** | 0.938*** | 1.143*** | 0.772*** | 0.807*** |
| | (0.089) | (0.064) | (0.085) | (0.067) | (0.054) |
| *p-values* | | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] |
| Conformal | [0.045] | [0.091] | [0.091] | [0.045] | [0.045] |
| Permutation inference, all destinations | [0.400] | [0.040] | [0.040] | [0.040] | [0.040] |
| Permutation inference, ≥1,000 applications | [0.438] | [0.062] | [0.062] | [0.062] | [0.062] |
| $N$ | 545 | 545 | 545 | 545 | 545 |
| $M_{pre}^{DE}$ | 34,957 | 9,763 | 3,537 | 2,363 | 3,863 |
| $\Delta = M_{pre}(e^{\hat{\beta}} - 1)$ | 12,343 | 15,180 | 7,550 | 2,752 | 4,793 |

*Notes:* PPML single-coefficient DiD (Germany × post) of equation (6) on monthly inward Category 1–Category 1 Eurodac hits, estimated separately for the never-applied group and for each sender group (Arrival, Transit, Destination) defined in Section 7.1. Standard errors (in parentheses) are clustered at the origin and destination country level. $M_{pre}^{DE}$ is the pre-ruling monthly base count in Germany and $\widehat{\Delta} = M_{pre}^{DE}(e^{\hat{\beta}} - 1)$ the implied monthly increase. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. The conformal test runs on the 22-month estimation window rather than on a longer one, because column (1) is built from destination-level total applications and the only series reaching back to 2009 is defined differently; the floor is therefore $1/22 = 0.045$.

Figure A20: Detailed sender-group event study: monthly applications in Germany




*Notes:* PPML event studies on monthly applications in Germany, splitting each of the three main-text sender groups (Figure 9) into finer sub-regions, stacked in three panels. *Arrival*: South = Greece, Italy, Spain, Malta, Cyprus; Poland; Other Eastern Europe = Croatia, Bulgaria, Romania, Hungary, Slovakia, Slovenia. *Transit*: AT & CH = Austria, Switzerland; France; Czechia. *Destination*: UK, IE & IS = United Kingdom, Ireland, Iceland; Scandinavia = Denmark, Finland, Norway, Sweden; Baltics = Estonia, Latvia, Lithuania; BENELUX = Belgium, Netherlands, Luxembourg. Portugal is excluded from this breakdown: it recorded very few monthly hits (0–2 per month for most of the pre-ruling period, rising to single digits only from late 2013), too small a flow to support a meaningful sub-region estimate. Each remaining sub-region's countries sum back exactly to its parent group. Specifications and clustering as in Figure 9. Shaded bands are 95 percent confidence intervals. Austria and Hungary enter here as Eurodac sender countries—their fingerprint records are available regardless of destination-sample coverage—even though both are excluded as *destinations* from the main Eurostat-based estimation sample (Section 3).

# G Robustness

**Total applications and the inclusion of Austria.** The baseline outcome is first-time applications, and Austria is excluded because Eurostat does not disaggregate first-time from repeat applications for Austria over the sample window (Section 3.1). Table A10 re-estimates the main gravity DiD on *total* applications (first-time plus subsequent), which are reported for every destination, Austria included. Column (1) reproduces the baseline first-time estimate; columns (2)–(3) use total applications, first excluding and then including Austria. The coefficient is somewhat smaller on total applications—repeat applications are mechanically less responsive to the ruling—but remains large and highly significant, and including Austria leaves it essentially unchanged.

Table A10: Robustness: total applications and the inclusion of Austria

| | (1) | (2) | (3) |
|---|---|---|---|
| DE × $\mathbb{1}(t > t_0)$ | 0.476*** | 0.386*** | 0.388*** |
| | (0.102) | (0.126) | (0.121) |
| *p-values* | | | |
| Cluster-robust | [<0.001] | [0.002] | [0.001] |
| Conformal | [0.043] | [0.022] | [0.022] |
| Permutation inference, all destinations | [0.360] | [0.538] | [0.519] |
| Permutation inference, ≥1,000 applications | [0.250] | [0.412] | [0.389] |
| Observations | 25,843 | 27,642 | 29,139 |
| Outcome | First-time apps | Total apps | Total apps |
| Includes Austria | No | No | Yes |
| geo-citizen fixed effects | ✓ | ✓ | ✓ |
| citizen-ym fixed effects | ✓ | ✓ | ✓ |

*Notes:* PPML gravity DiD (*DE* × *Post*) with destination × origin and origin × year-month fixed effects; two-way clustering on destination and origin. Window August 2011–June 2013 (July 2012 excluded); non-EU origins. Column (1): first-time applications, baseline destinations (excl. Austria). Columns (2)–(3): total applications (first-time + subsequent), excluding and then including Austria. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. The conformal test uses the same three samples extended back to August 2009, so $T = 46$ and the floor is 0.022.

**More demanding fixed effects and macro controls.** A further check varies the controls of the main specification along two dimensions at once. Table A11 reports the coefficient under four progressively more demanding fixed-effect structures, each estimated without and with a destination × quarter unemployment-rate control. The estimate is stable across all eight columns: neither richer seasonality controls nor local labor-market conditions account for the rise in applications in Germany.

**Aggregation and the Germany-vs-control collapse.** Table A12 re-estimates the main DiD as the data are progressively collapsed: from the full origin × destination × month panel (column 1), to a destination × month panel summed over origins (column 2), to a two-unit panel of Germany versus the combined rest of the EU+ over months (column 3). The coefficient is stable across all three levels (0.44–0.49). Because all fixed effects in the main specification are origin-nested, $\hat{\beta}$ is a weighted average of origin-specific contrasts, which is why the collapse barely moves it.

Table A11: Robustness: fixed-effect structures and the unemployment control

| | Year-month + pair FE | | Baseline | | + Dest. × month FE | | + Dest. × origin × month FE | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| DE × $\mathbb{1}(t > t_0)$ | 0.490*** | 0.399*** | 0.476*** | 0.430*** | 0.479*** | 0.474*** | 0.437*** | 0.435*** |
| | (0.071) | (0.098) | (0.102) | (0.110) | (0.106) | (0.092) | (0.101) | (0.086) |
| *p-values* | | | | | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] | [<0.001] |
| Conformal | [0.022] | [0.022] | [0.043] | [0.043] | [0.043] | [0.043] | [0.043] | [0.043] |
| Permutation inference, all destinations | [0.280] | [0.478] | [0.360] | [0.435] | [0.400] | [0.261] | [0.360] | [0.261] |
| Permutation inference, ≥1,000 applications | [0.312] | [0.462] | [0.250] | [0.308] | [0.250] | [0.231] | [0.250] | [0.308] |
| Destination × Origin FE | Yes | Yes | Yes | Yes | Yes | Yes | | |
| Year-month FE | Yes | Yes | | | | | | |
| Origin × Year-month FE | | | Yes | Yes | Yes | Yes | Yes | Yes |
| Destination × month-of-year FE | | | | | Yes | Yes | | |
| Destination × Origin × month-of-year FE | | | | | | | Yes | Yes |
| Unemployment-rate control | ✗ | ✓ | ✗ | ✓ | ✗ | ✓ | ✗ | ✓ |
| Observations | 25,577 | 23,235 | 25,577 | 21,588 | 25,537 | 21,548 | 16,013 | 13,466 |

*Notes:* PPML; "month" denotes month-of-year. Each pair of columns is one fixed-effect structure estimated without (odd columns) and with (even columns) a destination × quarter unemployment-rate control (Eurostat `une_rt_q`). Cols (1)–(2): year-month and destination × origin FE. Cols (3)–(4): baseline (destination × origin and origin × year-month FE). Cols (5)–(6): add destination × month-of-year FE. Cols (7)–(8): destination × origin × month-of-year FE, which subsumes the destination × origin FE and is identified off year-on-year variation within month-of-year, so its estimation sample is smaller. The unemployment control is available only for the subsample of destinations with Eurostat unemployment data, so even-numbered columns are estimated on fewer observations. Two-way clustering on destination and origin, except cols (1)–(2), where the stripped fixed-effect structure makes the two-way correction degenerate and clustering is on destination only. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. The conformal counterfactual uses each column's own fixed-effect structure, on the window extended back to August 2009 ($T = 46$, floor 0.022). It omits the unemployment control, since the counterfactual is a prediction device and the validity of the test rests on the symmetry of the procedure rather than on the counterfactual being correctly specified; the even columns keep their restriction to unemployment-covered destinations.

**EASY registration counts for Germany.** A residual concern is that the application response reflects a change in how, or how quickly, Eurostat records applications rather than a change in arrivals. Germany's BAMF keeps a separate administrative count, the EASY system (*Erstverteilung von Asylbegehrenden*), which registers asylum seekers when they first present at a reception center—before, and independently of, the formal application that Eurostat counts. We collapse the bilateral panel to the destination × month level (summing first-time applications over origins) and re-run the Germany-versus-controls event study with destination and year-month fixed effects, once with Germany's Eurostat applications and once replacing them with the EASY count. Figure A21 shows that the two event studies almost coincide: both are flat before the ruling and rise to a comparable plateau after it. The single-coefficient collapsed DiD is 0.492 on Eurostat applications and 0.544 on EASY registrations (a 72 percent increase), both close to the main estimate. The response therefore reflects more asylum seekers reaching Germany, not a recording convention.

**Alternative estimators: applications per 10,000 inhabitants.** Table A13 re-estimates the main treatment effect using applications per 10,000 inhabitants as the outcome across six alternative specifications. The exhibits that follow show what each column is fitting. Figure A22 plots the raw series behind the two linear difference-in-differences columns: Germany against the comparison group, averaged in the two ways those columns weight it. Figure A23 gives the corresponding event study,

Table A12: Main estimate at different levels of aggregation

| | (1) Main (orig.×dest.×month) | (2) Destination × month | (3) Germany vs. control × month |
|---|---|---|---|
| DE × $\mathbb{1}(t > t_0)$ | 0.476*** | 0.492*** | 0.486*** |
| | (0.102) | (0.075) | (0.071) |
| *p-values* | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] |
| Conformal | [0.043] | [0.022] | [0.022] |
| Permutation inference, all destinations | [0.360] | [0.280] | [–] |
| Permutation inference, ≥1,000 applications | [0.250] | [0.312] | [–] |
| Implied effect $e^{\hat{\beta}} - 1$ | 61.0% | 63.5% | 62.7% |
| Origin × destination FE | Yes | | |
| Origin × year-month FE | Yes | | |
| Destination FE | | Yes | |
| Unit FE (Germany vs. rest) | | | Yes |
| Year-month FE | | Yes | Yes |
| Observations | 25,577 | 545 | 44 |

*Notes:* PPML DiD (Germany × post) on first-time applications, main DiD window (August 2011–June 2013, July 2012 excluded). Column (1) is the main specification; columns (2)–(3) sum applications to coarser units before estimating. Columns (1)–(2) cluster two-way (destination and origin, or destination and year-month); column (3) has two cross-sectional units, so a heteroskedasticity-robust standard error is reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. Permutation inference is not available in column (3): collapsing the control group into a single "Rest of EU+" unit leaves one unit to reassign the treatment to, so there is no placebo distribution to rank the estimate against. The conformal test survives the collapse, because it permutes along time rather than across units. Each column's counterfactual uses its own fixed-effect structure on the window extended back to August 2009 ($T = 46$).

unweighted and population-weighted. Figures A24 and A25 give the synthetic-control and synthetic difference-in-differences fits, for the short and the long pre-period respectively, and Tables A14 and A15 report the weights behind all four of those fits.

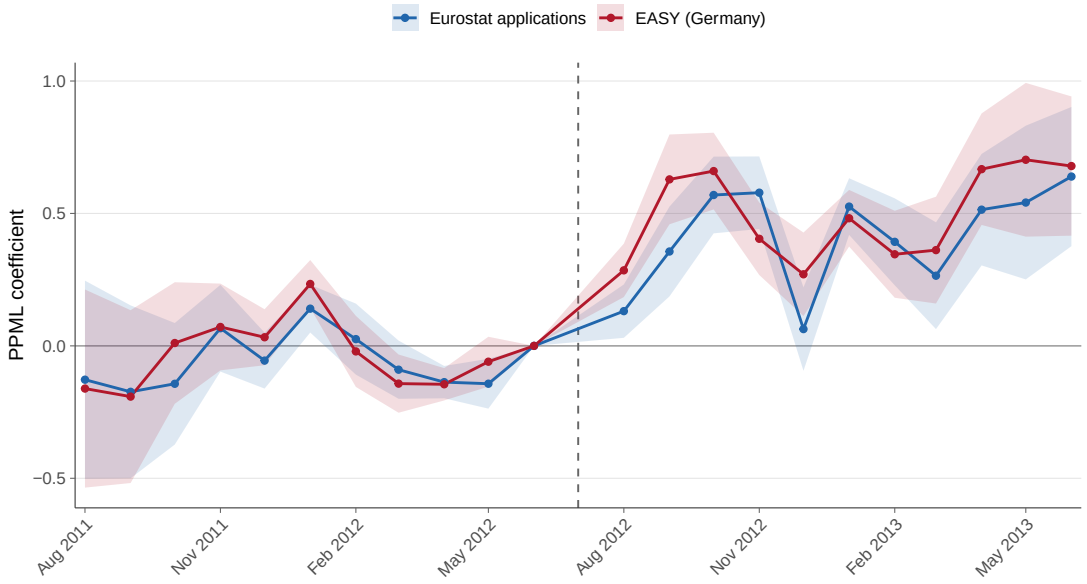
**Alternative pre- and post-treatment windows.** Table A16 re-estimates the ruling effect across alternative pre- and post-treatment windows of 6, 11 (baseline), and 18 months.

**Excluding small origin countries.** Table A17 re-estimates the main specification after progressively dropping origin countries with few EU+-wide applications in the 11-month pre-reform window (Aug. 2011–Jun. 2012, summed across all EU+ destinations), to check that the result is not driven by a handful of origins with thin, noisy counts.

**Leave-one-out jackknife.** The baseline coefficient is robust to dropping any single control destination or origin one at a time (Figure A26).

**Recognition-rate functional form.** The heterogeneity result interacts the treatment with the recognition rate linearly. Table A18 shows that the gradient is not an artifact of that choice: adding a quadratic term, or replacing the linear term with flexible 25-percentage-point bins, leaves the same pattern in place, with the response declining steeply as the recognition rate rises.

Figure A21: Collapsed destination × month event study: Eurostat applications versus EASY registrations




*Notes:* PPML event study at the destination × month level (first-time applications summed over origins), with destination and year-month fixed effects and Germany × event-time indicators; standard errors clustered on destination, with 95 percent confidence bands shaded. The reference month is June 2012; the dashed line marks the ruling (July 2012). The blue series uses Germany's Eurostat first-time applications; the red series replaces them with the BAMF EASY registration count, which is a national monthly total and is therefore not broken down by origin.

**Alternative definitions of the recognition rate.** The recognition-rate gradient could in principle reflect Germany's own pre-ruling recognition decisions rather than a general property of the origin, or be sensitive to how far back the pre-ruling window used to measure it reaches. Table A19 re-estimates the interaction under five measures of the origin recognition rate: Germany's pre-ruling rate over the full H1 2010–H1 2012 window (the baseline), the same Germany-specific rate over the narrower H1 2011–H1 2012 (6 quarters) and H2 2011–H1 2012 (4 quarters) windows, the whole-EU pre-ruling rate, and a predicted rate that scales the post-ruling non-German rate by the pre-ruling Germany-to-non-Germany ratio, so that it uses no German recognition decision from the post period. The negative gradient is present and of similar magnitude under all five, so it reflects which origins respond rather than Germany's own recognition behavior or the particular pre-period window chosen.

Table A13: Alternative estimators: applications per 10,000 inhabitants

| | (1) Linear DiD Unweighted | (2) Linear DiD Weighted† | (3) SC Short pre (Aug 2011) | (4) SDiD Short pre (Aug 2011) | (5) SC Long pre (Aug 2009) | (6) SDiD Long pre (Aug 2009) |
|---|---|---|---|---|---|---|
| DE $\times \mathbb{1}(t > t_0)$ | 0.428*** | 0.345*** | 0.414** | 0.383*** | 0.430*** | 0.454*** |
| | (0.099) | (0.128) | (0.174) | (0.092) | (0.156) | (0.091) |
| *p-values* | | | | | | |
| Cluster-robust | 0.000 | 0.007 | | | | |
| Conformal | | | 0.045 | 0.136 | 0.021 | 0.064 |
| Permutation inference | 0.160 | 0.160 | | | | |
| RMSPE ratio | | | 0.120 | 0.040 | 0.080 | 0.040 |
| Effect, % of DE pre-reform mean | 86.3% | 69.5% | 83.5% | 77.1% | 86.8% | 91.4% |
| Observations | 550 | 550 | 550 | 550 | 1,175 | 1,175 |

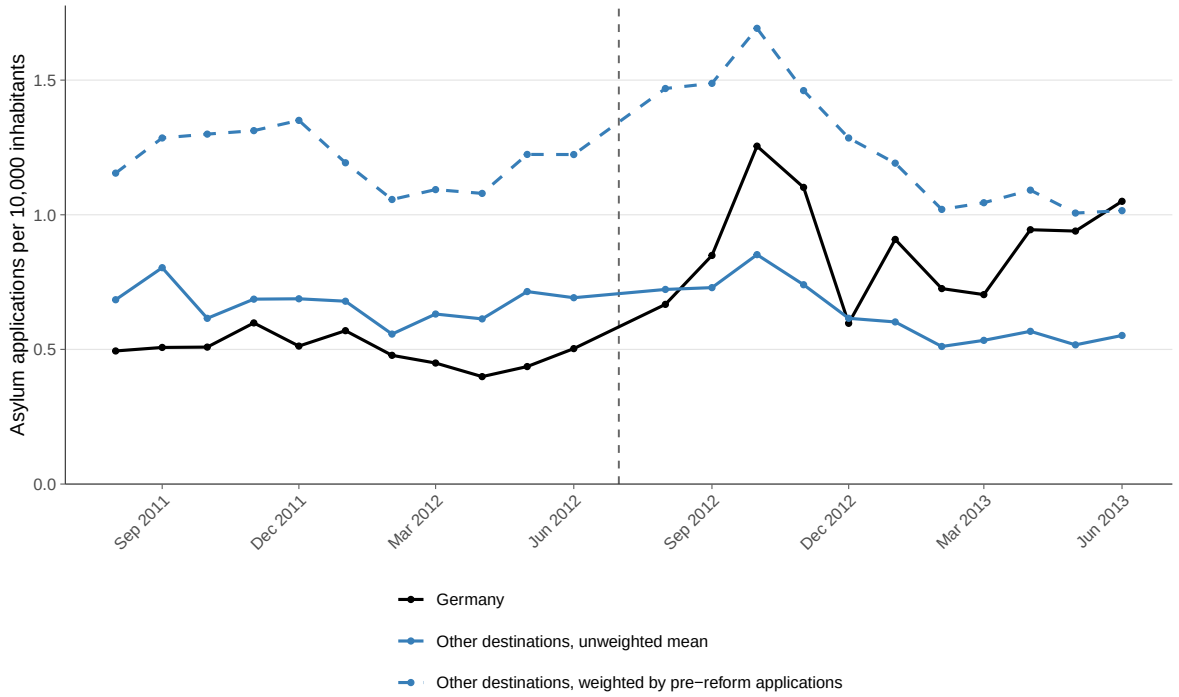
*Notes:* Outcome: monthly applications per 10,000 inhabitants; all columns estimated at the destination × month level (collapsed over origin). Hungary is dropped from every column: it has real data in only 6 of 22 short-pre-window months, the rest zero-filled by panel balancing, which would otherwise corrupt both its own placebo fit and its role as a donor for every other destination (24 donor destinations remain). †Column (2) weights by each destination's total applications in the 11-month pre-reform window (Aug. 2011–Jun. 2012). For columns (3)–(6), the point estimate is fit on the full 24-donor pool, as always; **the standard error (and the significance stars) is computed from the placebo distribution of Arkhangelsky et al. (2021) after removing the five donor destinations with the highest pre-treatment RMSPE** (consistently drawn from Belgium, Finland, Luxembourg, Norway, Sweden, and Switzerland, the donors whose own synthetic-control fit is markedly worse than Germany's and which otherwise dominate the placebo variance without being informative). Three further rows report complementary $p$-values that do *not* feed into the stars: "Permutation inference" re-assigns the ruling to each of the 24 donor destinations in turn, re-estimates the same linear DiD, and ranks Germany's $|t|$ against the resulting placebo distribution, $p = (1 + \#\{|t_j| \geq |t_{\text{DE}}|\})/25$; the permuted statistic is studentized rather than the raw $\hat{\beta}$ because, with destinations of very different size, the null distribution of $\hat{\beta}$ depends on which unit is pseudo-treated, as MacKinnon and Webb (2020) show. It is reported only for columns (1)–(2): the synthetic-control columns have no per-donor standard error to studentize by, and already carry the two rank-based $p$-values below. "RMSPE ratio" is the rank-based test of Abadie, Diamond and Hainmueller (2010) (Germany's post/pre-treatment RMSPE ratio ranked against the same ratio for all 24 donors, $p = \text{rank}/25$); "conformal" is the moving-block permutation test of Chernozhukov, Wüthrich and Zhu (2021) on Germany's own residual series over time, which does not involve the donor pool at all. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (stars from the RMSPE-screened SE only).

Table A14: Synthetic control donor weights

| Country | Short pre | Long pre |
|---|---|---|
| Netherlands | 0.415 | 0.060 |
| Ireland | 0.379 | 0.006 |
| Spain | 0.000 | 0.343 |
| Lithuania | 0.000 | 0.176 |
| Slovenia | 0.075 | 0.121 |
| Denmark | 0.033 | 0.089 |
| Switzerland | 0.080 | 0.032 |
| Belgium | 0.000 | 0.043 |
| Sweden | 0.004 | 0.042 |
| Portugal | 0.004 | 0.025 |
| Luxembourg | 0.001 | 0.020 |
| Greece | 0.000 | 0.019 |
| United Kingdom | 0.000 | 0.011 |
| Bulgaria | 0.000 | 0.009 |
| Slovakia | 0.006 | 0.000 |
| Italy | 0.000 | 0.003 |
| Finland | 0.002 | 0.000 |
| *Pre-treatment RMSPE* | 0.0176 | 0.0566 |

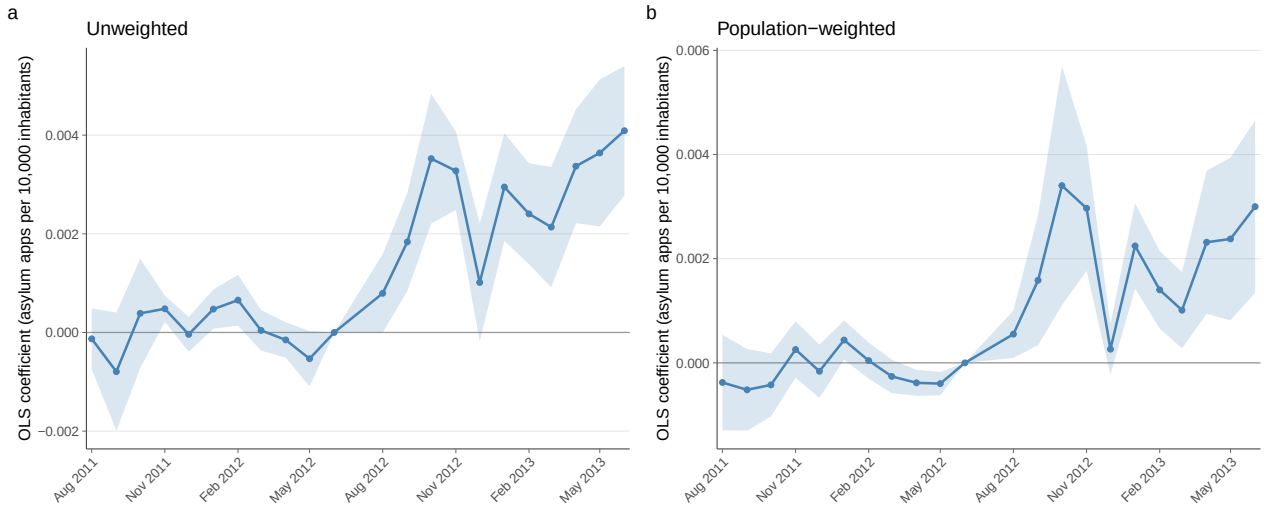
*Notes:* Donor weights $\omega$ for the synthetic-control fits in Figures A24 and A25, panel (a) of each: "Short pre" fits the weights on August 2011–June 2012, "Long pre" on August 2009–June 2012. The donor pool is the same 24 destinations in both columns; destinations receiving no weight in either fit are omitted from the table. The final row reports the root mean squared prediction error of the fitted synthetic unit over its own pre-treatment months, which is larger for the long pre-period simply because the fit is asked to match many more of them.

Figure A22: Applications per 10,000 inhabitants: Germany and the comparison group



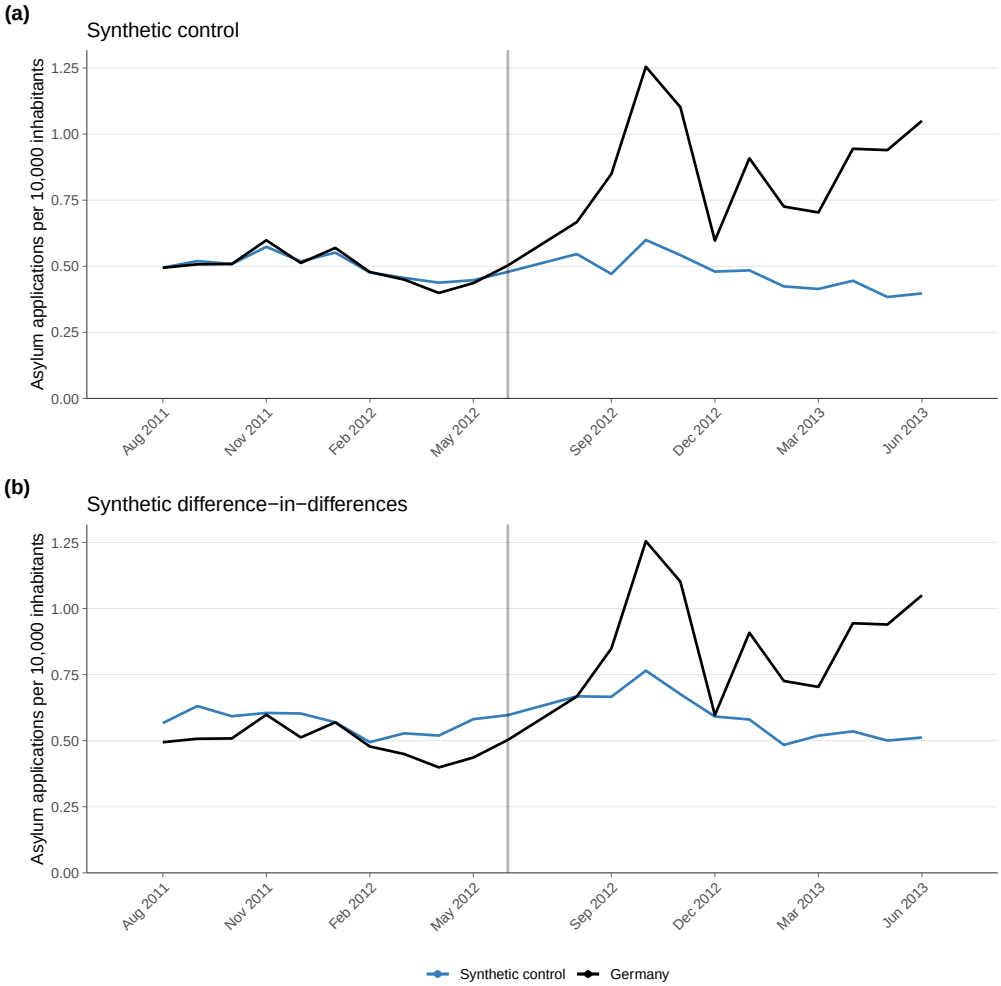

*Notes:* Monthly first-time asylum applications per 10,000 inhabitants, summed over origins within each destination. The two comparison series are two averages of the same 24 donor destinations: the solid line gives every destination equal weight, which is the implicit comparison in column (1) of Table A13, and the dashed line weights each destination by its total applications over the 11 pre-reform months (August 2011–June 2012), which is the comparison in column (2). The weighted average sits well above the unweighted one throughout because the largest destinations by application volume also run the highest rates; what matters for the estimate is not the level gap but the divergence after the ruling. The dashed vertical line marks the ruling (July 2012).

Figure A23: Event study: applications per 10,000 inhabitants (OLS)



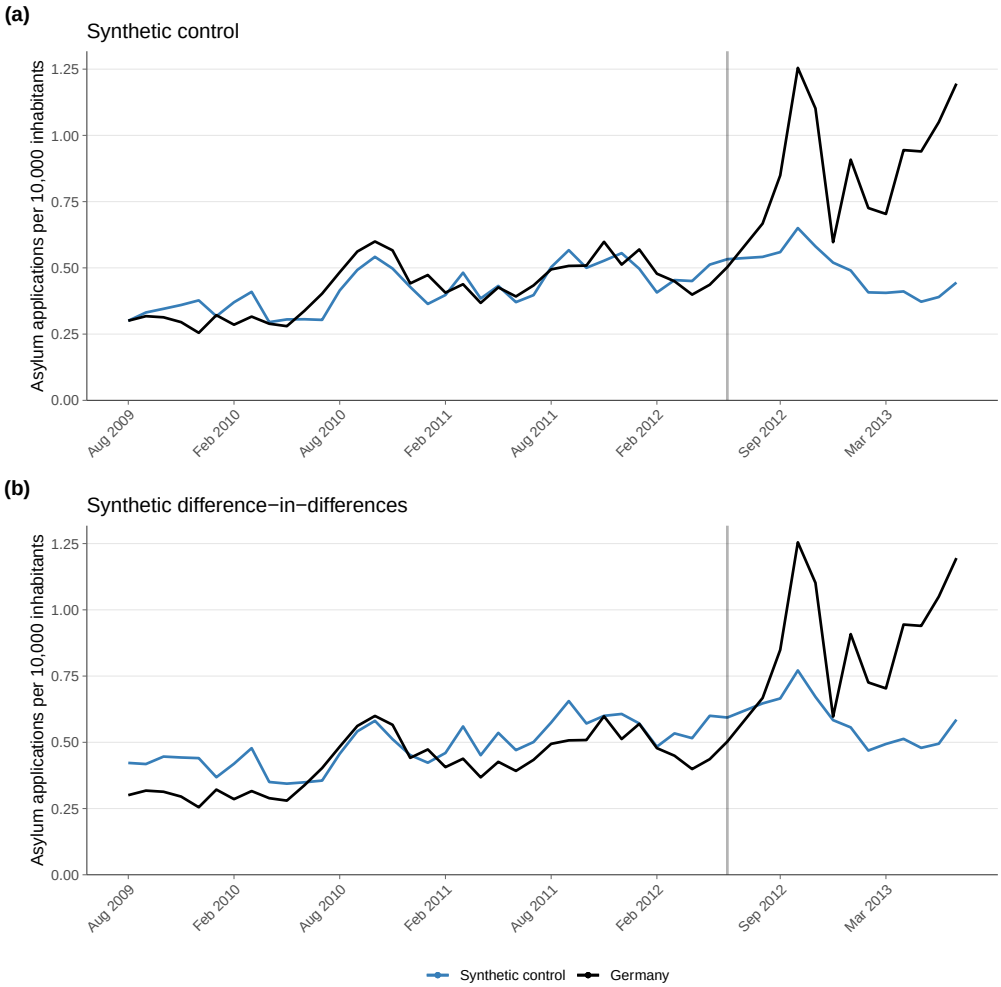

*Notes:* OLS event study on applications per 10,000 inhabitants. Panel (a): unweighted. Panel (b): population-weighted (complementing panel a). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure A24: Synthetic control and synthetic DiD: applications per 10,000 inhabitants




*Notes:* Fitted paths for columns (3) and (4) of Table A13. Panel (a): synthetic control. Panel (b): synthetic difference-in-differences. Both are fit on the destination × month panel of applications per 10,000 inhabitants over a pre-period beginning in August 2011, with Hungary excluded and 24 donor destinations remaining. The vertical line marks the ruling (July 2012); donor weights are in Tables A14 and A15.

Figure A25: Synthetic control and synthetic DiD: long pre-period (from August 2009)




*Notes:* As Figure A24, except that the pre-period begins in August 2009 rather than August 2011, giving almost three times as many pre-treatment months on which to fit the donor weights. Panel (a): synthetic control. Panel (b): synthetic difference-in-differences. These are the fits behind columns (5) and (6) of Table A13. A longer matching window is a more demanding test of the design, since the synthetic unit must now track Germany through the 2010–2011 run-up as well as through the immediate pre-reform months; both panels show it doing so, and both show the same post-ruling divergence.

Table A15: Synthetic DiD donor and time weights

| *Panel A. Donor weights $\omega$* | | |
|---|---|---|
| Country | Short pre | Long pre |
| Denmark | 0.040 | 0.050 |
| Lithuania | 0.043 | 0.048 |
| Netherlands | 0.047 | 0.045 |
| Slovenia | 0.043 | 0.047 |
| Switzerland | 0.047 | 0.031 |
| Sweden | 0.037 | 0.047 |
| Spain | 0.043 | 0.047 |
| Belgium | 0.046 | 0.046 |
| Slovakia | 0.043 | 0.046 |
| Ireland | 0.045 | 0.046 |
| Bulgaria | 0.044 | 0.046 |
| Portugal | 0.043 | 0.045 |
| Czech Republic | 0.044 | 0.044 |
| France | 0.044 | 0.044 |
| United Kingdom | 0.044 | 0.044 |
| Romania | 0.043 | 0.044 |
| Estonia | 0.042 | 0.044 |
| Finland | 0.043 | 0.043 |
| Poland | 0.043 | 0.040 |
| Latvia | 0.042 | 0.042 |
| Norway | 0.042 | 0.022 |
| Greece | 0.041 | 0.035 |
| Italy | 0.041 | 0.034 |
| Luxembourg | 0.010 | 0.021 |
| *Pre-treatment RMSPE* | 0.0887 | 0.0936 |
| *Panel B. Time weights $\lambda$* | | |
| Period | Short pre | Long pre |
| Aug 2009 | — | 0.023 |
| Jan 2010 | — | 0.118 |
| Mar 2010 | — | 0.472 |
| Oct 2011 | 0.624 | 0.009 |
| Jun 2012 | 0.376 | 0.378 |

*Notes:* Weights for the synthetic difference-in-differences fits in Figures A24 and A25, panel (b) of each, on the same two pre-periods as Table A14. Panel A: donor weights. Unlike synthetic control, synthetic difference-in-differences spreads weight almost evenly across the whole donor pool rather than concentrating it on a few destinations, which is a design feature of the estimator and not a sign of a poorly determined fit. Panel B: time weights, listing every pre-reform month that receives a weight above 0.001 in either fit; a month outside the short pre-period has no entry in that column. The two fits agree on June 2012, the month immediately before the ruling, as the most informative single pre-period month.

Table A16: Alternative pre/post windows

| | Pre-period (post = baseline) | | | | Post-period (pre = baseline) | | |
|---|---|---|---|---|---|---|---|
| | 18 mo | Baseline | 6 mo | Jun 2012 | 6 mo | Baseline | 18 mo |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| DE × $\mathbb{1}(t > t_0)$ | 0.533*** | 0.476*** | 0.412*** | 0.386*** | 0.372*** | 0.476*** | 0.655*** |
| | (0.116) | (0.102) | (0.101) | (0.117) | (0.098) | (0.102) | (0.121) |
| *p-values* | | | | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [0.001] | [<0.001] | [<0.001] | [<0.001] |
| Conformal | [0.069] | [0.045] | [0.118] | [–] | [0.059] | [0.045] | [0.036] |
| Permutation inference, all destinations | [0.200] | [0.360] | [0.360] | [0.600] | [0.480] | [0.360] | [0.280] |
| Permutation inference, ≥1,000 applications | [0.188] | [0.250] | [0.250] | [0.500] | [0.250] | [0.250] | [0.188] |
| Observations | 34,822 | 25,577 | 18,869 | 12,653 | 18,621 | 25,577 | 34,277 |

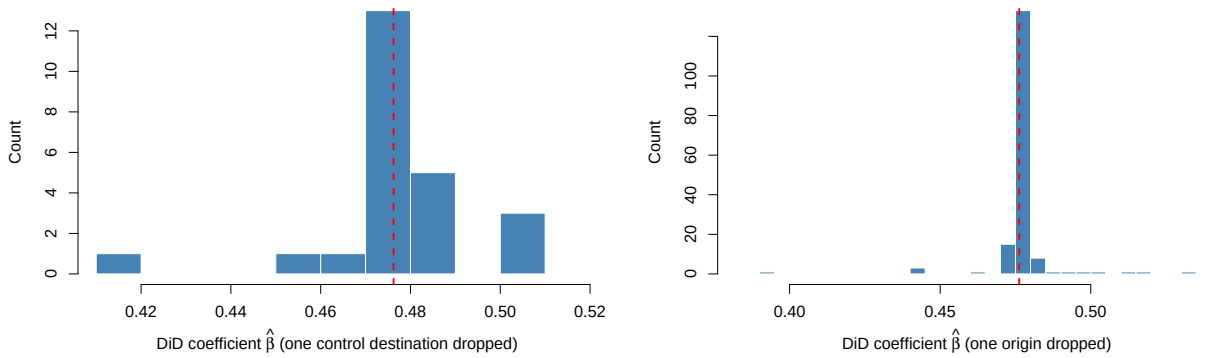
*Notes:* PPML. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. The conformal test here runs on each column's own window rather than on a common long pre-ruling period, since the window is what the table varies; the floor therefore moves across columns, from 1/29 to 1/17, and column (4) admits no test at all because a single pre-ruling month leaves nothing to rotate.

Table A17: Robustness: excluding small origin countries

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | All origins | ≥100 apps | ≥500 apps | ≥1,000 apps |
| DE × $\mathbb{1}(t > t_0)$ | 0.476*** | 0.476*** | 0.476*** | 0.471*** |
| | (0.102) | (0.102) | (0.103) | (0.105) |
| *p-values* | | | | |
| Cluster-robust | [<0.001] | [<0.001] | [<0.001] | [<0.001] |
| Conformal | [0.043] | [0.043] | [0.043] | [0.043] |
| Permutation inference, all destinations | [0.360] | [0.360] | [0.400] | [0.400] |
| Permutation inference, ≥1,000 applications | [0.250] | [0.250] | [0.250] | [0.250] |
| Origin countries | 169 | 85 | 63 | 44 |
| Observations | 25,577 | 24,198 | 20,862 | 16,149 |

*Notes:* PPML. Columns (2)–(4) drop origin countries whose total first-time applications across all EU+ destinations in the 11-month pre-reform window fall below the stated threshold. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Inference.* The bracketed rows give $p$-values for the reported coefficient under four procedures, set out in Appendix C: the cluster-robust value implied by the standard error above it; the moving-block conformal test of Chernozhukov, Wüthrich and Zhu (2021), which rotates the treated unit's log gap between observed and counterfactual outcomes through every starting point; and permutation inference (MacKinnon and Webb, 2020), which reassigns the treatment to each control destination in turn and ranks Germany's $|t|$ among the draws, reported for the full placebo pool and for destinations with at least 1,000 pre-ruling applications. The conformal test uses the same columns extended back to August 2009, so $T = 46$ and the floor is 0.022; the window is not what this table varies.

Figure A26: Leave-one-out jackknife of the main DiD coefficient




(a) Drop one control destination (b) Drop one origin

*Notes:* The main Poisson DiD re-estimated dropping one control destination (panel a, 24 estimates) or one origin (panel b, 168 estimates) at a time. Histograms of the resulting DiD coefficients; the dashed line is the full-sample estimate.

Table A18: Recognition-rate functional form

| | (1) Linear | (2) + Quadratic | (3) 25-pp bins |
|---|---|---|---|
| DE × $\mathbb{1}(t > t_0)$ | 0.775*** | 0.825*** | 0.717*** |
| | (0.134) | (0.143) | (0.124) |
| DE × $\mathbb{1}(t > t_0)$ × Recognition rate | -1.15*** | -1.86* | |
| | (0.271) | (0.990) | |
| DE × $\mathbb{1}(t > t_0)$ × Recognition rate$^2$ | | 1.08 | |
| | | (1.39) | |
| DE × $\mathbb{1}(t > t_0)$ × Recog 25–49% | | | -0.545*** |
| | | | (0.108) |
| DE × $\mathbb{1}(t > t_0)$ × Recog 50–74% | | | -0.604*** |
| | | | (0.144) |
| DE × $\mathbb{1}(t > t_0)$ × Recog 75–100% | | | -0.685*** |
| | | | (0.118) |
| Destination × Origin FE | Yes | Yes | Yes |
| Origin × Year-month FE | Yes | Yes | Yes |
| Observations | 23,261 | 23,261 | 23,261 |

*Notes:* PPML. Col 3 reference bin: 0–24 percent. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A19: Recognition-rate gradient under alternative definitions

| | (1) Germany pre-period (H1 2010–H1 2012) | (2) Germany H1 2011–H1 2012 | (3) Germany H2 2011–H1 2012 | (4) EU-wide pre-period (H1 2010–H1 2012) | (5) Predicted post-period |
|---|---|---|---|---|---|
| DE × $\mathbb{1}(t > t_0)$ | 0.775*** | 0.776*** | 0.764*** | 0.804*** | 0.778*** |
| | (0.134) | (0.128) | (0.125) | (0.149) | (0.136) |
| × Recognition rate | −1.147*** | −1.108*** | −1.032*** | −1.177*** | −0.956*** |
| | (0.271) | (0.220) | (0.199) | (0.305) | (0.217) |
| $N$ | 23,261 | 22,706 | 22,378 | 24,754 | 23,261 |

*Notes:* PPML with the main-DiD fixed effects and two-way clustering. Each column interacts the treatment with a different origin recognition-rate measure: (1) Germany, H1 2010–H1 2012 (baseline); (2) Germany, H1 2011–H1 2012 (6 quarters); (3) Germany, H2 2011–H1 2012 (4 quarters); (4) all EU+ destinations, H1 2010–H1 2012; (5) predicted, $(\text{recog}^{DE}_{\text{pre}}/\text{recog}^{\neg DE}_{\text{pre}}) \times \text{recog}^{\neg DE}_{\text{post}}$, with the post window 2012 Q3–2013 Q2. Samples differ slightly because the predicted measure is missing for origins without non-German decisions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

# H Application-benefit elasticities using full EU+ variation

As a preliminary benchmark, Appendix Table A20 reports a fully continuous specification on the full destination × origin × month panel from 2000 onwards, regressing monthly applications on the log real benefit (harmonized asylum seeker benefit for a single adult) with destination × origin, origin × year-month, and destination × month-of-year fixed effects, plus a destination log(GDP per capita, PPP) control to net out the general economic-attractiveness confound that more generous benefits tend to coincide with richer destinations. The full-panel elasticity is 1.82; excluding Germany it is 1.13. These cross-country elasticities are likely biased by confounding policies and shocks beyond GDP alone—which is exactly why the rest of the paper focuses on the German court ruling, where the source of variation is plausibly exogenous.

Table A20 reports the continuous-variation specifications for asylum applications, all destinations (1) and excluding Germany (2). The regressor is the log real (own-HICP, 2012 prices) EUR harmonized asylum seeker benefit for a single adult (full-in-kind-only countries lifted by the cross-country mean food top-up, see Appendix A.2), plus the log real destination GDP per capita (PPP, World Bank WDI) control. Coefficients are elasticities. "Excl. DE" drops Germany. The analogous specification for first-instance decisions yields 1.40 for positive decisions (cited in Section 8); it is estimated on the coarser quarterly origin × destination decisions panel (Eurostat `migr_asydcfstq`) and omitted from the table for brevity.

Panel B re-estimates the same specifications dropping the 2014–2017 refugee-crisis years, a period of large origin shocks and correlated destination policy responses; the benefit elasticity is not driven by the crisis period.

Table A20: Benefit elasticity of asylum applications

| | (1) All | (2) Excl. DE |
|---|---|---|
| *Panel A: Full sample (2000–2025)* | | |
| Log real asylum benefit | 1.82*** | 1.13** |
| | (0.439) | (0.455) |
| Log GDP per capita (dest., PPP) | 1.23 | 0.719 |
| | (1.04) | (0.976) |
| Observations | 474,341 | 438,095 |
| *Panel B: Excluding 2014–2017 refugee-crisis years* | | |
| Log real asylum benefit | 1.41*** | 0.845*** |
| | (0.405) | (0.298) |
| Log GDP per capita (dest., PPP) | 1.14** | 1.18*** |
| | (0.522) | (0.417) |
| Observations | 361,732 | 332,349 |

*Notes:* PPML; coefficients are elasticities with respect to the log real asylum seeker benefit. FE: destination × origin, origin × year-month, destination × month-of-year. Both columns control for log real destination GDP per capita (PPP, World Bank WDI). "Excl. DE" drops Germany. Two-way clustering on destination and origin. Panel B drops the 2014–2017 refugee-crisis years. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

# I Policy Response: The Western Balkans

Western Balkan citizens had visa-free Schengen access from December 2009 (North Macedonia, Montenegro, Serbia) and December 2010 (Albania, Bosnia–Herzegovina), following EU visa liberalisation (Council of the European Union, 2009; European Stability Initiative, 2013). Visa-free access is not unconditional: the European Commission periodically monitors whether partner countries meet the

associated requirements—including migration-related criteria such as the rate of unfounded asylum applications—and publishes reports that can inform decisions to suspend visa exemptions (European Commission, 2025). When applications from these countries surged following the July 2012 ruling, the German government responded through two distinct measures. First, from October 24, 2012—the date the Federal Interior Ministry publicly announced accelerated processing for Serbian and Macedonian applicants—the BAMF administratively reallocated its case workers to focus exclusively on Western Balkan claims, reducing average processing times from approximately three months to nine days (European Stability Initiative, 2013). Second, in November 2014, Germany formally designated Serbia, Bosnia–Herzegovina, and North Macedonia as "safe countries of origin" (*sichere Herkunftsstaaten*) under the Asylum Procedure Act (Bundesrepublik Deutschland, 2014). These measures are distinct from the *Westbalkan-Regelung* (§ 26 Abs. 2 BeschV, in force January 2016), which is a labor-migration instrument operating outside asylum law (Bundesrepublik Deutschland, 2015).

The Western Balkan experience indicates that a large country raising its benefits may increase the total number of EU asylum applications—not merely divert applicants from other destinations. Those who responded to the German benefit increase from this region were predominantly economic migrants for whom staying in their origin country was a viable alternative, not applicants who would have filed claims elsewhere regardless. The Germany-only application series in Figure A15a declines after October 2012, consistent with a deterrence effect of accelerated processing rather than the benefit level itself.