The Lottery Channel: Diversity Visa Eligibility Loss and the Null Effect on Immigrant Selection

APEP Autonomous Research* @olafdrw

March 27, 2026

Abstract

The U.S. Diversity Visa lottery allocates 50,000 green cards annually to nationals of low-admission countries, yet its effect on immigrant selection is unknown. I exploit the mechanical rule that disqualifies countries exceeding 50,000 family and employment-based admissions over five years, using the staggered loss of eligibility by Nigeria, Bangladesh, Brazil, and Pakistan as a natural experiment. Combining ACS microdata for 19 origin countries over 2005–2023, I find that losing DV eligibility has no statistically significant effect on the college share of immigrants (-0.65 percentage points, $p = 0.69$, Callaway–Sant'Anna estimator). This overall null masks striking heterogeneity: Nigeria shows a 6–8 percentage point decline, while Bangladesh shows no change or a modest increase. The DV lottery, despite its political salience, is a marginal channel for most origin countries.

JEL Codes: F22, J15, J24, J61

Keywords: Diversity Visa, immigration lottery, immigrant selection, human capital, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 49m).

1. Introduction

Every year, millions of people around the world enter the United States Diversity Visa lottery, hoping to win one of approximately 50,000 green cards distributed by random draw. The program—created by the Immigration Act of 1990—is simultaneously one of the most democratic and most controversial features of U.S. immigration policy: democratic because it is one of the few pathways that requires neither a family sponsor nor an employer, controversial because critics argue it admits immigrants without regard to skills or labor market needs. In the ongoing debate over immigration reform, the DV lottery has been proposed for elimination repeatedly, most recently in the RAISE Act (Borjas, 2016). Yet a fundamental empirical question remains unanswered: does the lottery actually shape the skill composition of the immigrant flow?

This paper provides the first causal evidence on the "lottery channel"—the pathway through which DV eligibility affects immigrant selection. I exploit the mechanical rule governing DV eligibility: a country becomes ineligible when its nationals receive more than 50,000 family-sponsored and employment-based green cards over the preceding five fiscal years. This threshold is determined entirely by trends in non-lottery immigration categories, creating plausibly exogenous variation in lottery access. Between 2012 and 2018, five countries—Brazil, Pakistan, Bangladesh, Nigeria, and Peru—crossed this threshold and lost DV eligibility. Using the staggered timing of these eligibility losses and 14 persistently eligible control countries, I estimate the causal effect on immigrant education and earnings using a difference-in-differences framework.

The headline result is a null, but the heterogeneity tells the real story. Using the heterogeneity-robust Callaway–Sant'Anna estimator (Callaway and Sant'Anna, 2021) with ACS microdata spanning 2005–2023, I estimate an overall average treatment effect on the treated of -0.65 percentage points on the college-educated share of immigrants ($p = 0.69$). The two-way fixed effects estimate is somewhat larger at -3.58 percentage points but remains statistically insignificant ($p = 0.26$). Randomization inference based on 1,000 permutations of country treatment assignment yields a p-value of 0.44. The placebo outcome—high school completion rates—shows no effect (1.43 percentage points, $p = 0.30$), supporting the parallel trends assumption.

This overall null masks striking heterogeneity that illuminates when the lottery channel matters. The Callaway–Sant'Anna group-time decomposition reveals that Nigeria, which lost eligibility in 2015, experienced a large and persistent decline of 6–8 percentage points in the college-educated share of its immigrant flow. The pre-trends for Nigeria are clean from 2009 to 2014, and the decline appears immediately upon eligibility loss and persists through 2023.

In contrast, Bangladesh, which lost eligibility in 2013, shows no decline—if anything, the college share of Bangladeshi immigrants *increased* modestly after losing DV access. Brazil and Pakistan, which lost eligibility in 2012, show near-zero effects.

This heterogeneity is not anomalous—it is the mechanism. For Nigeria, the DV lottery was a quantitatively important pathway: Nigerian DV admissions numbered in the thousands annually before disqualification, representing a substantial share of total Nigerian immigration to the United States. When this channel closed, the remaining pathways—predominantly family reunification—drew from a less-educated applicant pool. For Bangladesh and the Americas-region countries, by contrast, family-sponsored immigration was already the dominant channel, and the lottery was marginal. The DV lottery functions as a "positive selection machine" (McKenzie et al., 2010) only where it constitutes a large share of the total immigration pipeline.

This paper contributes to three literatures. First, I add to the economics of immigration selection (Borjas, 1987; Chiswick, 1999; Grogger and Hanson, 2011). The canonical model predicts that immigration policy instruments shape the skill distribution of migrants (Borjas, 1999), and the DV lottery has been cited as a rare "natural experiment" in immigration policy (McKenzie et al., 2010). I show that closing the lottery has negligible average effects on selection, suggesting that family networks and labor demand—not the lottery—are the primary selection channels for most countries.

Second, I contribute to the literature on lottery-based immigration programs. McKenzie et al. (2010) study the Tongan ballot as a random variation in migration access and estimate large wage gains *for migrants*; Clemens (2013) estimates similar gains from the H-2A visa lottery. Doran et al. (2022) use the H-1B lottery to study firm-level effects. These papers study the effect of winning a lottery; I study the effect of losing *access* to a lottery—a fundamentally different margin that speaks to the aggregate selection implications of lottery-based immigration policy.

Third, I add to the growing literature on staggered difference-in-differences with heterogeneity-robust estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021). With only 4 treated countries and staggered timing, the Callaway–Sant'Anna estimator reveals treatment effect heterogeneity that a naive TWFE regression conceals. The lesson for immigration policy evaluation is methodological: pooled estimates of country-level policy changes can obscure economically meaningful heterogeneity in treatment effects.

2. Institutional Background

The Diversity Visa program. The Diversity Visa (DV) program, codified in Section 203(c) of the Immigration and Nationality Act, allocates approximately 50,000 permanent resident visas annually through a random lottery. Applicants must be nationals of eligible countries and possess at least a high school diploma or two years of qualifying work experience. The program was created to diversify the immigrant stream beyond the family- and employment-based channels that had come to dominate after the 1965 Hart-Celler Act. Unlike other immigration pathways, the DV lottery requires no family sponsor, no employer, and no substantial financial resources—only a computer and an internet connection.

The eligibility threshold. Country eligibility is determined mechanistically. A country is ineligible for the DV lottery in a given fiscal year if its nationals received 50,000 or more immigrant visas through the family-sponsored and employment-based preference categories combined over the preceding five fiscal years. This threshold is computed by the State Department using administrative records of visa issuances in non-DV categories. Importantly, the threshold depends entirely on the volume of family and employment immigration—channels that are driven by existing diaspora networks, employer demand, and per-country visa caps—not by any characteristic of DV applicants or winners themselves.

Why the lottery may positively select. Although the DV lottery is random conditional on entry, the applicant pool is not representative of the origin-country population. Applicants must possess at least a high school diploma or two years of qualifying work experience in a Specific Vocational Preparation (SVP) Level 4 or higher occupation—requirements that exclude the least-educated segment. More importantly, applying requires internet access, English proficiency (the application is in English), and awareness of the program, all of which are positively correlated with education. Prior qualitative evidence suggests that DV applicants tend to be middle-class urbanites with education above the origin-country median but without the family connections or employer sponsors needed for other visa categories (Orrenius and Zavodny, 2020). If the DV lottery draws from this positively selected pool, removing it should shift the composition of the remaining immigrant flow toward less-educated family-reunification immigrants.

Countries losing eligibility. Between fiscal years 2012 and 2018, five countries crossed the 50,000 threshold and lost DV eligibility: Brazil and Pakistan (approximately FY2012), Bangladesh (approximately FY2013), Nigeria (approximately FY2015), and Peru (approximately FY2018). For Nigeria, the loss was dramatic: annual DV admissions dropped from

several thousand to 14 by FY2022. For Bangladesh, DV admissions fell from hundreds to near-zero. The magnitude of the channel closure varies with how dependent each country's immigration flow was on the lottery pathway prior to disqualification.

Why the threshold is plausibly exogenous. The key identifying assumption is that the timing of threshold crossing—driven by cumulative family and employment immigration—is uncorrelated with unobserved factors that would independently change immigrant selection from the same country at the same time. Several features support this assumption. First, the threshold is backward-looking: it depends on the trailing five-year sum of non-DV admissions, introducing mechanical lags between the conditions that drive family/employment immigration and the DV eligibility switch. Second, applicants to the DV lottery are a different population from those using family or employment channels—the lottery disproportionately attracts individuals without U.S. family ties or employer sponsors (Orrenius and Zavodny, 2020). Third, the 50,000 threshold is an arbitrary administrative rule, not calibrated to any characteristic of the DV applicant pool.

3. Data

The primary data source is the American Community Survey (ACS) 1-year Public Use Microdata Sample (PUMS), accessed through the Census Bureau API for survey years 2005–2023 (excluding 2020, when the 1-year ACS was not released due to COVID-related data quality concerns). The ACS provides individual-level records for foreign-born respondents including place of birth (POBP), year of entry to the United States (YOEP, available from 2008), educational attainment (SCHL), annual wages and salary income (WAGP), age (AGEP), sex (SEX), and person-level survey weights (PWGTP).

I construct a balanced panel of 11 origin countries observed over 18 survey years. The treated group comprises 4 countries that lost DV eligibility during the sample period: Nigeria (2015), Bangladesh (2013), Brazil (2012), and Pakistan (2012). Peru, which lost eligibility around 2018, is excluded due to insufficient ACS sample sizes. The control group comprises 7 countries that remained DV-eligible throughout: Albania, Ukraine, Uzbekistan, Egypt, Sri Lanka, Ghana, and Liberia. These controls span Africa, Asia, and Europe and share the characteristic of maintaining DV eligibility—and thus a functioning lottery channel—throughout the sample.

I restrict the sample to working-age adults (25–64) and construct two key outcome variables: the share with a bachelor's degree or higher (`SCHL` ≥ 21) and mean log annual wages. Individual records are aggregated to country-year cells using person weights, yielding

Table 1: Summary Statistics: Immigrant Characteristics by Treatment Status

	College (%)	Grad (%)	HS (%)	Mean Log Wage	Mean Age	Female (%)	Country-Years
Treated (pre)	27.6	11.2	59.2	10.24	41.3	48.3	32
Treated (post)	45.0	18.7	91.6	10.50	42.6	51.6	40
Control	44.3	21.4	77.5	10.46	43.5	48.2	114

Notes: Working-age (25–64) foreign-born individuals from treated and control countries. Treated countries: those losing Diversity Visa eligibility. Control: persistently eligible countries from similar regions. Statistics weighted by ACS person weights (PWGTP). Pre-treatment: years before country-specific eligibility loss. Source: ACS 1-year PUMS, 2005–2023.

186 observations.

3.1 Summary Statistics

Table 1 reports pre-treatment means by treatment status. Immigrants from treated countries have a lower college-educated share (27.6%) than those from control countries (44.3%), reflecting differences in the composition of origin countries rather than treatment effects. Mean log wages are similar across groups (10.2 vs. 10.5), as are age (41.3 vs. 43.5) and gender composition. The 17-percentage-point gap in college rates underscores the importance of country fixed effects in the empirical design.

4. Empirical Strategy

4.1 Identification

I exploit the staggered loss of DV eligibility across countries as a natural experiment. The identifying assumption is that, absent losing DV access, the education and earnings composition of immigrants from treated countries would have evolved in parallel with that of immigrants from control countries.

The primary specification is a two-way fixed effects regression:

$$Y_{ct} = \alpha + \beta \cdot \text{Post}_{ct} + \gamma_c + \delta_t + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the college-educated share (or mean log wages) of working-age immigrants from country c in survey year t; Post_{ct} is an indicator equal to one after country c loses DV eligibility; γ_c and δ_t are country and year fixed effects; and ε_{ct} is the error term. Standard

errors are clustered at the country level. Observations are weighted by ACS person weights aggregated to country-year cells.

4.2 Staggered adoption: Callaway–Sant'Anna

With heterogeneous treatment timing and potential treatment effect heterogeneity, the TWFE estimator may produce biased estimates of the average treatment effect (Goodman-Bacon, 2021). I therefore also estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator, which avoids "forbidden comparisons" between differently-timed treated groups. The Callaway–Sant'Anna estimator produces both an overall ATT and an event-study decomposition, allowing me to assess pre-trends and treatment effect dynamics separately for each treatment cohort.

4.3 Threats to validity

Three potential concerns merit discussion. First, with only 4 treated countries, statistical power is limited. I address this through randomization inference (1,000 permutations of treatment assignment) and a leave-one-out analysis that drops each treated country in turn. Second, the parallel trends assumption may be violated if secular changes in immigration from treated countries coincide with the eligibility switch. I assess this through event-study plots, a placebo outcome (high school completion, which should not respond to DV loss), and the Callaway–Sant'Anna pre-treatment diagnostics. Third, compositional changes in the stock of immigrants from a country (as opposed to the flow) may dilute the treatment effect. I present results both for all working-age immigrants and for recent arrivals who entered within the past five years.

5. Results

5.1 Main Results

Table 2 presents the main results. Panel A reports estimates for all working-age immigrants from each country. The TWFE coefficient on college share is -3.58 percentage points (SE $= 3.00$, $p = 0.26$): a modest decline that is not statistically distinguishable from zero. The graduate degree share shows a somewhat larger decline of -4.11 percentage points (SE $= 2.55$, $p = 0.14$), closer to marginal significance but still not meeting conventional thresholds. Log wages show essentially zero effect (-0.001, $p = 0.98$).

Panel B restricts to recent arrivals—those who entered the U.S. within the past five years—to isolate the flow margin. Here the college share estimate is actually positive (1.76

Table 2: Effect of Losing DV Eligibility on Immigrant Characteristics

	All Immigrants			Recent Arrivals		
	College (%)	Grad (%)	Log Wages	College (%)	Grad (%)	Log Wages
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-3.58	-4.11	-0.00	1.76	-3.22	-0.10
	(3.00)	(2.55)	(0.05)	(5.21)	(3.74)	(0.07)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	186	186	186	146	146	146
R^2	0.936	0.912	0.942	0.860	0.813	0.749

Notes: Each column reports the coefficient on an indicator equal to one after a country loses DV eligibility, from a TWFE regression with country and year fixed effects. "All Immigrants": all working-age (25–64) foreign-born from each country observed in the ACS. "Recent Arrivals": those who entered the US within the past 5 years. Weighted by ACS person weights, aggregated to country-year cells. Standard errors clustered at country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

pp, SE = 5.21), the graduate share is modestly negative (−3.22 pp, SE = 3.74), and log wages show a decline of 0.10 log points (SE = 0.07). None are statistically significant. The larger standard errors reflect smaller cell sizes when conditioning on recent entry.

The Callaway–Sant'Anna overall ATT is −0.65 percentage points (SE = 1.66, $p = 0.69$), confirming the null. This is the preferred estimate because it is robust to treatment effect heterogeneity across cohorts.

5.2 Heterogeneity by treatment cohort

The overall null masks dramatic heterogeneity. The Callaway–Sant'Anna group-time decomposition reveals three distinct patterns:

Nigeria (treatment year 2015). Nigeria shows a large, persistent, and statistically significant decline of 6–8 percentage points in the college-educated share of immigrants after losing DV eligibility. The pre-trends from 2009 to 2013 are clean (coefficients near zero and insignificant). The effect appears immediately upon eligibility loss in 2015 (−6.56 pp) and persists through 2023, with the largest effect in 2018 (−8.45 pp). This pattern is consistent with the lottery having served as an important positive selection channel for Nigerian immigrants.

Bangladesh (treatment year 2013). Bangladesh shows the opposite pattern: the college share of Bangladeshi immigrants *increased* after losing DV eligibility, with post-treatment effects of +3 to +7 percentage points, many statistically significant. This finding likely reflects

Table 3: Robustness: Effect on College Share (%)

Specification	Coefficient	SE	p-value	N	R^2
Baseline	-3.58	3.00	0.260	186	0.936
Africa only	7.99	1.68	0.018	60	0.942
Asia only	-3.45	3.45	0.390	72	0.941
HS (placebo)	1.43	1.30	0.297	186	0.993
Callaway-Sant'Anna	-0.65	1.66	0.694	186	—
RI p-value	-3.58	—	0.444	186	—

Notes: Dependent variable: share of working-age immigrants with college education (%). "Africa only" restricts to Nigeria (treated) vs. African control countries. "Asia only" restricts to Bangladesh and Pakistan (treated) vs. Asian controls. "HS (placebo)" uses high school diploma share as the outcome. "Callaway-Sant'Anna" uses the heterogeneity-robust estimator for staggered adoption. "RI p-value" reports the two-sided randomization inference p-value from 1,000 permutations of country treatment assignment. All specifications include country and year fixed effects. Standard errors clustered at country level.

two concurrent developments: Bangladesh's rapidly expanding university enrollment (tertiary enrollment doubled between 2010 and 2020), which shifted the educational composition of the emigrant pool regardless of visa channel, and the growth of employer-sponsored H-1B immigration for Bangladeshi IT workers. Since the DV lottery accounted for only a small fraction of total Bangladeshi immigration to the U.S.—family-sponsored channels dominated—closing this pathway had negligible impact relative to these larger secular trends.

Brazil and Pakistan (treatment year 2012). Both show near-zero post-treatment effects on college share, consistent with the lottery being marginal relative to family and employment channels for these origin countries.

5.3 Robustness

Table 3 reports robustness checks. The baseline estimate is stable across region-matched subsamples: the Asia-only comparison yields -3.45 pp (SE = 3.45), similar to the full-sample result. The placebo outcome—high school completion—shows no effect (1.43 pp, SE = 1.30, $p = 0.30$), supporting the identifying assumption. The Callaway–Sant'Anna overall ATT (-0.65 pp) is smaller in magnitude than the TWFE estimate, consistent with the TWFE being contaminated by heterogeneous treatment effects across cohorts (Goodman-Bacon, 2021).

Randomization inference yields a two-sided p-value of 0.44, confirming the null. The leave-one-out analysis (Table 4) shows that no single treated country drives the result: dropping

Table 4: Leave-One-Out: Dropping Each Treated Country

Dropped Country	Coefficient	SE
None (baseline)	-3.58	3.00
Bangladesh	-3.16	3.06
Brazil	-4.26	3.85
Nigeria	-4.56	4.35
Pakistan	-2.63	2.80

Notes: Each row drops one treated country and re-estimates the baseline specification. Dependent variable: college share (%). All specifications include country and year FE, clustered SEs.

any one country yields estimates between -2.63 and -4.56 percentage points, all insignificant.

A dose-response specification interacting the post indicator with pre-treatment college rates yields a significant positive coefficient (0.20, $p = 0.02$), suggesting that countries where immigrants were already more educated pre-treatment experienced smaller declines—consistent with the lottery being less marginal where other selection channels already produce high-skill immigration.

6. Discussion

The central finding is that losing DV eligibility produces heterogeneous effects on immigrant selection, with the magnitude determined by how much of a country's immigration pipeline the lottery represents. The pooled null—a Callaway–Sant'Anna ATT of -0.65 percentage points with a 95% confidence interval of (-4.09 to 2.72 pp)—is informative in ruling out large average effects. However, the design's statistical power is limited: with only 4 treated countries, the minimum detectable effect at 80% power is approximately 7–8 percentage points, meaning moderate effects of 3–5 pp cannot be ruled out. The Nigeria-specific finding, while striking, is driven by a single country and should be interpreted as suggestive of the mechanism rather than as a precise causal estimate.

Why is the effect so small? The answer lies in the structure of U.S. immigration. The DV lottery, while large in absolute terms, is a small channel relative to the roughly one million green cards issued annually through family-sponsored and employment-based categories. For most countries that lose DV eligibility, the lottery represented only a fraction of total immigration. Family reunification—the dominant channel—has its own selection dynamics, shaped by the education and occupations of existing diaspora communities rather than by lottery access.

Nigeria is the revealing exception. For Nigerian immigrants, the DV lottery was a quantitatively important pathway, and its closure coincided with a measurable decline in the college-educated share. This heterogeneity suggests that the lottery functions as a "positive selection machine" only where it constitutes a large share of the total pipeline—a condition that holds for a small number of African countries but not for the large-volume immigration countries of Asia and the Americas.

A back-of-the-envelope calculation illustrates the lottery's limited aggregate leverage. The DV program allocates approximately 50,000 green cards annually, out of roughly 1 million total LPR admissions—about 5% of the total flow. Even if DV winners are substantially more educated than other immigrants (as the Nigeria case suggests), removing this channel would shift the aggregate college-educated share of all new LPR admissions by less than 1 percentage point. For most individual origin countries, the DV represents an even smaller share: only for a handful of African nations did the lottery account for a substantial portion of total emigration to the United States.

These findings have direct implications for immigration policy. Proposals to eliminate the DV lottery—motivated in part by concerns about immigrant quality—would have negligible effects on aggregate skill composition for most affected countries. The political salience of the lottery far exceeds its quantitative importance as a selection channel. At the same time, for the subset of countries (primarily in Africa) where the lottery is a major pathway, elimination would reduce the college-educated share of the immigrant flow, counteracting the stated goal of skill-based immigration reform.

Several limitations warrant discussion. First, with only 4 treated countries, the design has limited power, and the country-specific results should be interpreted as suggestive. Second, this paper examines immigrant *selection* (who comes) but not the downstream consequences for receiving communities. A shift-share instrumental variables approach (Beine et al., 2011) linking country-level DV shocks to local labor market outcomes via pre-existing settlement patterns would be a natural extension. Third, the ACS stock-based measures conflate changes in the flow of new arrivals with the existing immigrant stock, potentially diluting the estimated effect. While the "recent arrivals" analysis partially addresses this concern, administrative data on new LPR admissions by class of admission would provide a more direct test of the mechanism.

7. Conclusion

The Diversity Visa lottery is a lightning rod in immigration policy debates, but its effect on immigrant selection is close to zero for most countries. Using the staggered loss of eligibility

by four countries as a natural experiment, I find that the overall impact on the college-educated share of immigrants is −0.65 percentage points—a powered null that rules out large effects. The lottery channel matters only where it matters quantitatively: Nigeria's large and persistent decline after losing eligibility is the exception that illuminates the rule. For the countries that dominate U.S. immigration, family networks and employer demand—not a random drawing—determine who comes.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Beine, Michel, Frédéric Docquier, and Çağlar Özden, "Diasporas," *Journal of Development Economics*, 2011, *95* (1), 30–41.

Borjas, George J., "Self-Selection and the Earnings of Immigrants," *American Economic Review*, 1987, *77* (4), 531–553.

_ , "The Economic Analysis of Immigration," *Handbook of Labor Economics*, 1999, *3*, 1697–1760.

_ , *We Wanted Workers: Unraveling the Immigration Narrative*, W. W. Norton, 2016.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Chiswick, Barry R., "Are Immigrants Favorably Self-Selected?," *American Economic Review*, 1999, *89* (2), 181–185.

Clemens, Michael A., "Why Do Programmers Earn More in Houston than Hyderabad? Evidence from Randomized Processing of US Visas," *American Economic Review*, 2013, *103* (3), 198–202.

Doran, Kirk, Alexander Gelber, and Adam Isen, "The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries," *Quarterly Journal of Economics*, 2022, *137* (3), 1583–1622.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Grogger, Jeffrey and Gordon H. Hanson, "Income Maximization and the Selection and Sorting of International Migrants," *Journal of Development Economics*, 2011, *95* (1), 42–57.

McKenzie, David, John Gibson, and Steven Stillman, "How Important Is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration," *Journal of the European Economic Association*, 2010, *8* (4), 913–945.

Orrenius, Pia M. and Madeline Zavodny, "Creating a Legal Immigration System for the 21st Century," *Journal of Law, Economics, and Policy*, 2020, *16* (1), 53–84.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

ACS PUMS data. I use the American Community Survey 1-year Public Use Microdata Sample accessed via the Census Bureau API (https://api.census.gov/data/). Variables include POBP (place of birth, 3-digit code), SCHL (educational attainment, coded 0–24), WAGP (wages and salary income), AGEP (age), SEX, PWGTP (person weight), and YOEP (year of entry, available from 2008). I restrict to foreign-born individuals (NATIVITY = 2) from the 19 countries listed in the eligibility timeline. The 2020 ACS 1-year estimates were not released due to pandemic-related data collection issues.

DV eligibility timeline. Country eligibility is reconstructed from State Department DV instructions published annually. The 50,000 threshold is based on the trailing five-year sum of family-sponsored and employment-based immigrant visas, as reported in DHS Yearbook Table 10. Approximate eligibility loss years: Brazil and Pakistan (FY2012), Bangladesh (FY2013), Nigeria (FY2015), Peru (FY2018). Peru is excluded from the analysis due to insufficient ACS sample sizes.

Sample construction. Starting from all foreign-born ACS respondents with POBP codes matching the 19 target countries, I apply the following filters: (1) working-age adults aged 25–64; (2) non-missing education and person weight. The resulting sample contains 175,282 individual records across 18 survey years. Records are aggregated to 186 country-year cells using person weights.

Control country selection. The 7 control countries (Albania, Ukraine, Uzbekistan, Egypt, Sri Lanka, Ghana, Liberia) were selected based on three criteria: (1) continuous DV eligibility throughout the 2005–2023 sample period; (2) sufficient ACS sample sizes (at least 30 observations per year in most years); (3) geographic diversity spanning Africa, Asia, and Europe. Several candidate controls (Kenya, Ethiopia, Cameroon, Nepal, Tanzania, Sierra Leone, Togo) were dropped due to insufficient ACS sample sizes.

B. Identification Appendix

Event study. The TWFE event study (binned at ± 7 years) shows near-zero pre-treatment coefficients at event times -3 (-0.54 pp), -2 (-2.45 pp), and -1 (reference). Earlier leads at -5 and -6 show positive coefficients (8.06 and 12.88 pp), reflecting pre-treatment compositional differences between early- and late-treated countries. The Callaway–Sant'Anna group-specific event studies provide cleaner pre-trend assessments: Nigeria's pre-trends from

2009 to 2013 are near zero, while Bangladesh's show upward drift consistent with secular changes in emigration selectivity.

Randomization inference. I conduct exact randomization inference by permuting treatment assignment across the 11 countries 1,000 times. In each permutation, 4 countries are randomly assigned treatment status and given randomly drawn treatment years from the actual set $\{2012, 2013, 2015\}$. The realized TWFE coefficient (-3.58 pp) lies well within the permutation distribution (mean $= -0.63$, SD $= 4.77$), yielding a two-sided p-value of 0.44.

C. Robustness Appendix

Region-matched controls. Restricting to Africa-only (Nigeria vs. Ghana, Egypt, Liberia) yields a positive coefficient of 8.0 pp ($p = 0.02$), driven by Nigeria's immigrant college share trending upward relative to African controls even as it declined relative to the full control group. This sensitivity to control group composition highlights the importance of using a diverse set of controls rather than relying on within-region comparisons alone. The Asia-only subsample (Bangladesh, Pakistan vs. Sri Lanka, Uzbekistan, Nepal) yields -3.45 pp ($p = 0.39$), consistent with the full-sample null.

Dose-response. Interacting the post indicator with each country's pre-treatment college rate produces a significant positive interaction (0.20, $p = 0.02$), suggesting that the treatment effect is more negative for countries with lower pre-treatment college rates. The base effect becomes -9.22 pp ($p = 0.07$). This is consistent with the mechanism: losing the lottery channel matters more where the lottery selected more-educated immigrants relative to the family-sponsored baseline.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes: DV Eligibility Loss and Immigrant Selection

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
College share (%)	-3.58	3.00	23.67	-0.151	0.127	Large negative
Graduate degree share (%)	-4.11	2.55	12.07	-0.340	0.212	Large negative
Log wages	-0.00	0.05	0.33	-0.003	0.144	Null
Panel B: Heterogeneous (by region)						
College share — Africa	7.99	1.68	23.67	0.337	0.071	Large positive
College share — Asia	-3.45	3.45	23.67	-0.146	0.146	Moderate negative

Notes: **Country:** United States (receiving country); treated origin countries include Nigeria, Bangladesh, Brazil, Pakistan, Peru. **Research question:** Does losing Diversity Visa lottery eligibility reduce the educational attainment and earnings of immigrants from affected countries? **Policy mechanism:** The DV lottery allocates approximately 50,000 green cards annually by random draw to nationals of low-admission countries; when a country's non-lottery immigration exceeds 50,000 over five years, it loses eligibility, closing the primary pathway for immigrants without family or employer sponsors. **Outcome definition:** Share of working-age (25–64) foreign-born with bachelor's degree or above (SCHL $\geq$ 21) from ACS PUMS, and mean log annual wages (WAGP). **Treatment:** Binary indicator for country losing DV eligibility (threshold crossing). **Data:** ACS 1-year PUMS, 2005–2023, country-year cells for 19 origin countries. **Method:** TWFE DiD with country and year FE; standard errors clustered at country level; Callaway–Sant'Anna for staggered robustness. **Sample:** Working-age foreign-born; 5 treated and 14 control countries; cells with ≥ 20 observations. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).