Latent Human Capital and the Immigrant Mobility Advantage: A General Equilibrium Analysis

Joseph Kopecky*

20th April 2026

Abstract

Children of low-income immigrants in the United States systematically out-earn children of comparable natives. I develop a dynastic general-equilibrium model to explain this "immigrant mobility advantage" and quantify the macroeconomic costs of the institutional frictions underlying it. The central mechanism is a wedge between latent human capital and realized earnings: immigrants are positively selected on ability but face institutional frictions that decay at rate λ. Calibrated to U.S. data, the model fits key empirical patterns and reveals that immigrant frictions cost the U.S. economy 4.94% of GDP. This loss is split between a static labor-misallocation loss (1.57 pp) and dynamic effects on intergenerational investment in human and physical capital (3.37 pp). I find that friction decay (assimilation) accounts for 87% of the second-generation advantage, with positive selection contributing the remainder. Finally, the model identifies a "sign-flip" threshold at $\lambda \approx 0.25$, beyond which frictions persist too strongly for the mobility advantage to appear. Calibrating λ across ten OECD destinations recovers a distribution of friction values that straddles this threshold, consistent with the heterogeneous mobility gaps documented in the empirical literature.

JEL classification codes: E24, J15, J24, J61, O15

Keywords: Immigration, intergenerational mobility, human capital, assimilation, general equilibrium

*Trinity College Dublin (jkopecky@tcd.ie)

1. Introduction

Why do the children of poor immigrants outperform the children of equally poor natives? This question, motivated by the influential work of Abramitzky, Boustan, Jacome, and Perez (2021), presents a puzzle for standard models of intergenerational mobility. In the United States, second-generation children whose immigrant parents earned at the 25th percentile reach roughly the 50th percentile of the income distribution, climbing approximately 5 percentile points higher than children of native-born parents at the same starting position. Figure 1 displays this "immigrant mobility advantage," a pattern that is robust across historical periods, origin countries, and empirical specifications.

Standard empirical frameworks for understanding intergenerational mobility treat observed parental income as a sufficient statistic for a child's expected outcome (Solon, 1999; Chetty, Hendren, Kline, and Saez, 2014). The theoretical foundation in Becker and Tomes (1986) is in fact more permissive, modeling unobserved endowments that transmit separately from observed earnings — a structure this paper builds on by formalizing the distinction between latent and realized human capital. Under the standard view, two families earning the same income should produce children with similar prospects, regardless of nativity. The immigrant mobility advantage violates this prediction and suggests that observed income is a poor measure of the relevant parental endowment for immigrant families.

This paper develops a mechanism that rationalizes this pattern, built on a distinction between *latent human capital* and *realized market productivity*. While the costly decision to migrate selects for high-ability individuals, institutional frictions (e.g., credential non-recognition, language barriers, and network effects) suppress their market earnings below what their abilities would otherwise command. The result is a systematic misallocation of high-ability labor. This is productive capacity that exists in the economy, but it is not fully utilized. Ultimately, the second-generation outcome is determined by an intergenerational "horse race". Because latent ability naturally reverts toward the population mean across generations, the immigrant mobility advantage manifests only when "assimilation" (the decay of market frictions) occurs rapidly enough to outpace this mean reversion.

This paper makes a dual contribution: it is the first to demonstrate that a simple latent human capital channel can quantitatively reproduce the patterns observed in microdata, and it quantifies the aggregate efficiency cost of the friction. To do so, I construct a dynastic overlapping-generations (OLG) general equilibrium model with parental investment in children's human capital. The key analytical novelty is the introduction of a separate decaying friction, θ_g, that is conceptually distinct from standard intergenerational skill

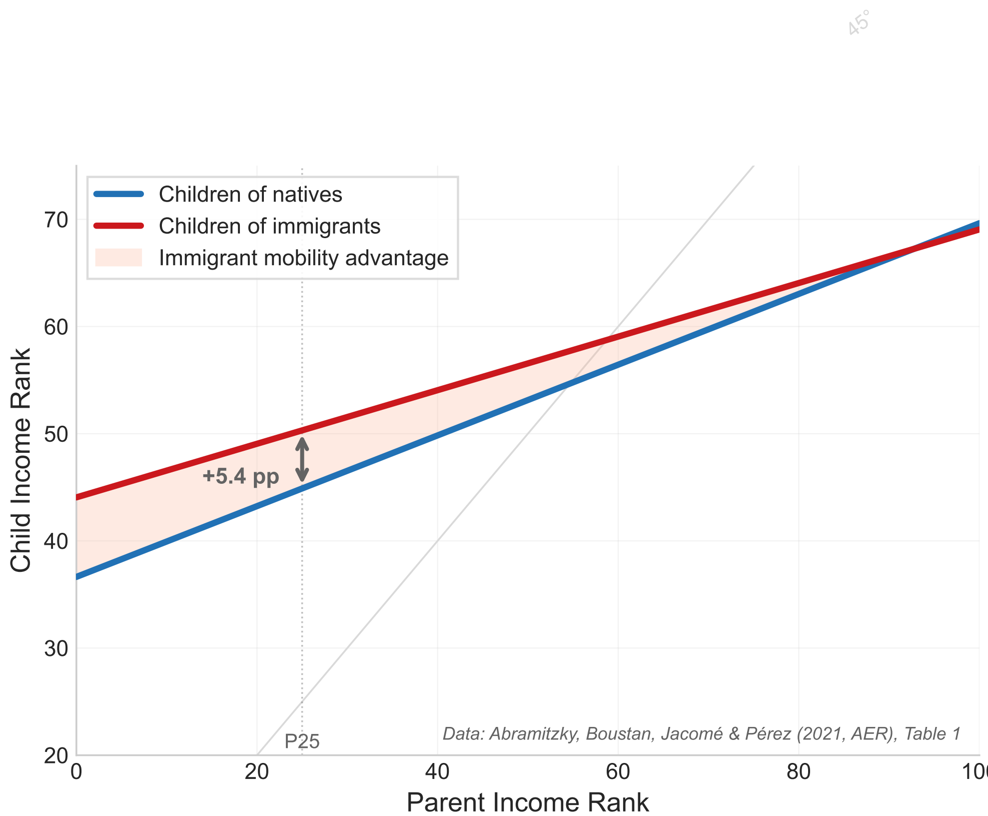

Figure 1: ***The immigrant mobility advantage.*** *Rank-rank regressions estimated by Abramitzky et al. (2021) (Table 1) for native-born and second-generation immigrant children in the U.S.*

persistence, ρ. This structure formalizes the components described above: ρ governs the transmission of latent ability from one generation to the next, while θ_g suppresses its market price. This distinction is what allows the model to produce a second-generation mobility gap for immigrants at the same observed parental income rank as natives, a feature that pure skill persistence cannot generate without an additional friction wedge.

I model the friction as evolving deterministically: $\theta_g = \kappa \lambda^{g-1}$. Here κ represents the initial penalty facing first-generation immigrants, and λ governs the rate of friction decay, or intergenerational assimilation. Calibrated to United States data, the model identifies a friction persistence of $\lambda_{US} = 0.175$. At this value, friction decays rapidly: a first-generation penalty of 30% shrinks to roughly 5% for the second generation. When λ is high, as in many European countries, "outsider" status persists across generations, suppressing both individual mobility and aggregate productivity.

The framework formalizes the under-placement mechanism that Abramitzky et al. (2021) propose informally and embeds it in a quantitative dynastic environment. This builds directly on the lineage of Borjas (1993), but provides the quantitative general-equilibrium discipline needed to decompose the contributions of selection and assimilation—a step that prior structural work does not take.

The contribution of this paper is quantitative rather than purely theoretical. While Borjas (1993) develops the closest theoretical predecessor[1], his framework lacks both a separate decaying friction distinct from skill persistence and the quantitative general-equilibrium discipline needed to decompose the contribution of each channel. To my knowledge, this is the first paper to calibrate a structural dynastic general-equilibrium model to the Abramitzky et al. (2021) mobility moments. Specifically, the model addresses four questions: (i) what fraction of the second-generation mobility advantage reflects under-placement of immigrant fathers versus positive selection on latent ability; (ii) how large is the aggregate output loss from suppressing immigrant productive capacity; (iii) at what threshold of friction persistence does the second-generation gap flip sign, and where does the United States sit relative to that threshold; and (iv) whether a single institutional parameter can rationalize the cross-country variation in second-generation outcomes documented by Boustan, Jensen, Abramitzky, Jácome, Manning, Pérez, Watley, Adermon, Arellano-Bover, Åslund, et al. (2025).

By separating under-placement from intergenerational inheritance as a distinct decaying friction, I am able to decompose the second-generation gap into selection and assimilation components and quantify the aggregate efficiency cost of the resulting misallocation. The model's strength lies in its ability to simultaneously rationalize a cluster of empirical findings from a single, parsimonious mechanism. These include the second-generation mobility advantage at the 25th percentile of parental income, third-generation convergence to the native distribution, first-generation wage penalties, and the broad patterns of origin-group heterogeneity and cross-country variation documented in both modern and historical data.

The calibrated model yields several quantitative insights. The central macroeconomic finding is that immigrant frictions impose a general-equilibrium output loss of 4.94% of U.S. GDP. To isolate the role of the model's dynamic linkages, I perform a static wedge-accounting exercise in the spirit of Hsieh and Klenow (2009); this shows that the direct labor-misallocation effect accounts for 1.57 percentage points of the total loss. The full general-equilibrium magnitude is thus approximately three times larger than the static wedge alone.

[1] That framework is a Markov skill-transmission model where selection at the first generation propagates through a single persistence parameter. It lacks both a separate decaying friction distinct from skill persistence ρ and the general-equilibrium discipline needed to decompose these channels. Appendix A.4 provides a diagnostic on the identifying content of κ within the present framework: a variant with the friction wedge fixed at zero ($\kappa = 0$) cannot simultaneously reproduce a positive second-generation mobility advantage and a below-unity first-generation wage ratio. While this exercise is intended as a diagnostic of the model's internal logic rather than a formal test of the Borjas (1993) specification, it suggests the need for two parameters to explain both empirical moments.

This difference reflects the model's dynamic effects: when frictions suppress immigrant earnings, they simultaneously reduce the resources available for intergenerational human capital investment and induce a subsequent capital-stock adjustment.

A formal decomposition of the second-generation advantage reveals that friction decay (assimilation) accounts for 87% of the improvement in outcomes, while positive selection on latent ability contributes the remaining 13%. This finding carries a sharp policy implication: it suggests that the immigrant mobility advantage is less about "who comes" than about how quickly institutional barriers (credential non-recognition, language gaps) are removed after arrival. While selection provides the opportunity for upward mobility, it is the rate of assimilation that determines whether that potential is realized within a single generation.

Extending the framework cross-country, the model identifies an "emergent" sign-flip threshold at $\lambda^* \approx 0.25$. While the United States ($\lambda_{US} = 0.175$) sits safely within the positive-advantage region, many European destinations calibrated in the cross-country exercise sit beyond this threshold. In these high-persistence regimes, intergenerational frictions are deep enough to erase the selection advantage, causing immigrant children to underperform natives. These results suggest that much of the heterogeneous global evidence on immigrant assimilation, including the "immigrant paradox" in the U.S. and the stagnation observed in parts of Europe, can be rationalized by the interaction of a common selection mechanism with varying speeds of institutional assimilation.

The paper builds on the empirical literature on immigrant intergenerational mobility (Abramitzky et al., 2021; Abramitzky, Boustan, and Eriksson, 2014; Borjas, 1993; Card, 2005; Boustan et al., 2025; Algan, Dustmann, Glitz, and Manning, 2010), on cross-country comparisons of first-generation earnings assimilation (Antecol, Kuhn, and Trejo, 2006), and on the theoretical lineage from Borjas (1993) through Ehrlich and Pei (2021). A small body of structural work touches adjacent territory. Ehrlich and Pei (2021) and Ehrlich and Kim (2015) build dynastic OLG models with endogenous immigration and parental investment in children's human capital, but their focus is the immigration surplus and endogenous growth rather than the second-generation mobility advantage, and neither model features an immigrant-specific friction wedge decaying across generations. Bohn and Lopez-Velasco (2018) embed intergenerational skill transitions in a political-economy model of immigration preferences, targeting median-voter policy outcomes rather than second-generation mobility magnitudes. Within-generation structural analogues of θ include Adda, Dustmann, and Goerlach (2022), who model first-generation wage assimilation jointly with return migration; Lessem and Sanders (2020), whose occupational-upgrading framework formalizes the within-career version of the under-placement idea; and Eckstein and Weiss (2004), who decompose the wage growth of former Soviet Union immigrants in Israel. Llull

(2018) provides the methodological precedent for embedding immigration in a labor-market equilibrium model. On the reduced-form side, Achard (2024) uses linked three-generation French data to estimate intergenerational mobility patterns analogous to those documented by Abramitzky et al. (2021) for the U.S. None of these papers addresses the second-generation mobility advantage structurally or executes a quantitative decomposition of selection versus assimilation.

The model of latent versus observed productivity relates to the "hidden capital" interpretation in Lubotsky (2007) and the wage assimilation literature initiated by Chiswick (1978). The human capital production function draws on Becker and Tomes (1986) and the quantitative implementation in Lee and Seshadri (2019), with causal empirical support for the intergenerational transmission channel from Oreopoulos, Page, and Stevens (2006), who use compulsory schooling laws to instrument for parental education. The general equilibrium framework with heterogeneous agents follows Aiyagari (1994) and Huggett (1996). The misallocation framing builds on the wedge-accounting tradition initiated by Hsieh and Klenow (2009) and Restuccia and Rogerson (2008), which quantifies aggregate productivity losses from heterogeneous distortions across producers. In a similar spirit, Lamadon, Mogstad, and Setzler (2022) quantify labor-market misallocation from imperfect competition using matched employer-employee tax data, estimating that eliminating labor-market wedges would raise U.S. aggregate output by roughly 3% and welfare by 5% — a methodological parallel to the present exercise applied to a different labor-market friction. A related literature on dynamic misallocation (Moll, 2014) shows that when distortions affect capital accumulation rather than just current allocation, static wedge accounting understates the true welfare cost. This paper applies that insight to immigrant labor-market frictions, where the dynamic margin is intergenerational human capital transmission in addition to capital-stock adjustment. Within the human-capital misallocation literature, the model connects to Hsieh, Hurst, Jones, and Klenow (2019) on the cost of suppressing high-ability talent and to Bell, Chetty, Jaravel, Petkova, and Van Reenen (2019) on the analogous loss of inventive capacity from differential exposure.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 describes the calibration strategy and data sources. Section 4 reports the quantitative results from three experiments. Section 5 exploits variation across origin countries and destination countries to test the mechanism's ability to match heterogeneous assimilation patterns. Section 6 discusses policy implications and directions for future work.

2. Model

2.1. Environment

Time is discrete, $t = 0, 1, 2, \ldots$. The economy is populated by a unit continuum of dynastic households. Each household lives for one period, supplies labor, consumes, saves, invests in their child's human capital, and is replaced by a single descendant. In each period, a fraction m of the population consists of newly arriving first-generation immigrants, with latent human capital drawn from the distribution specified in Section 3, who replace an equal measure of randomly selected exiting households, leaving the total population constant. This replacement formulation isolates the within-immigrant misallocation channel from the additive labor-supply effects emphasized by Borjas (1993) and Card (2005); the implied assumption is that accidental bequests from exiting households accrue to their replacements, with the discount factor β calibrated to match the aggregate capital-output ratio under this structure. Relaxing the replacement assumption (allowing migration to add to aggregate labor supply) would raise the output loss in Section 4.1 rather than lower it, as additional migrants would contribute additional suppressed productive capacity at the calibrated κ.

2.2. Households

A household is characterized by a state vector $s = (a, h, g)$, where $a \in \mathbb{R}_+$ denotes financial assets (bequests), $h \in \mathbb{R}_+$ denotes latent human capital, and $g \in \{0, 1, 2, \ldots, G\}$ denotes generational status ($g = 0$ for natives, $g = 1$ for first-generation immigrants, etc.).

Preferences. The household has CRRA preferences over consumption c and values its child's welfare through an altruism factor β:

$$V(a, h, g) = \max_{c, a', i} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \beta \, \mathbb{E}_{\varepsilon'} \left[V(a', h', g') \right] \right\} \tag{1}$$

where $\sigma > 0$ is the coefficient of relative risk aversion, a' is the bequest to the child, i is investment in the child's human capital, and h' and g' are the child's state variables.

Budget Constraint. The household allocates total resources between consumption, saving, and investment:

$$c + a' + i = (1 + r)\, a + w \cdot h \cdot (1 - \theta_g) \tag{2}$$

where r is the interest rate, w is the wage per efficiency unit of labor, and θ_g is the generation-specific labor market friction. The term $w \cdot h \cdot (1 - \theta_g)$ is the household's realized labor income: latent ability h is discounted by friction θ_g. The non-negativity constraints $a' \geq 0$ and $i \geq 0$ complete the household's feasible set.

2.3. Labor Market Friction

The friction θ_g captures institutional barriers that prevent immigrants from fully monetizing their latent human capital. It evolves deterministically across generations:

$$\theta_g = \begin{cases} 0 & \text{if } g = 0 \quad \text{(native)} \\ \kappa \cdot \lambda^{g-1} & \text{if } g \geq 1 \quad \text{(immigrant descendant)} \end{cases} \tag{3}$$

where $\kappa \in (0, 1)$ is the initial penalty facing first-generation immigrants and $\lambda \in [0, 1)$ governs the persistence of friction across generations.

The parameter λ is the central object of interest. When λ is close to zero, friction decays rapidly: the child of an immigrant ($g = 2$) faces a penalty of only $\lambda\kappa$, which is near zero. When λ is close to one, friction persists indefinitely across generations, producing a permanent "outsider" status. I interpret λ as summarizing the institutional environment (credential recognition systems, language integration programs, labor market openness, discrimination) that determines how quickly immigrant families achieve full economic participation.

2.4. Human Capital Formation

Latent human capital evolves across generations according to a log-linear production function:

$$\ln h' = \rho \ln h + \gamma \ln i + \varepsilon', \qquad \varepsilon' \sim N(0, \sigma_\varepsilon^2) \tag{4}$$

where $\rho \in (0, 1)$ is the intergenerational persistence of ability (capturing genetic endowments and cultural transmission), $\gamma > 0$ is the elasticity of child human capital with respect to parental investment, and ε' is an idiosyncratic ability shock.

Two features of this specification are important. First, $\rho < 1$ ensures mean reversion in ability: absent investment, all dynasties converge to the same long-run mean. Second, $\gamma > 0$ creates a channel through which parental resources affect child outcomes, introducing a role for the budget constraint and hence for the friction θ_g.

2.5. The Key Mechanism: Intuition

Consider two families, one native and one immigrant, with identical observed income at the 25th percentile (y_{P25}). The native family has low income because of low latent ability: $h_{\text{nat}} = y_{P25}/w$. The immigrant family has low income because of friction: $h_{\text{mig}} = y_{P25}/[w(1-\kappa)]$. Crucially, $h_{\text{mig}} > h_{\text{nat}}$.

For the child's expected human capital, taking logs of equation (4) and ignoring investment differences and shocks, the expected human capital gap is:

$$\mathbb{E}[\ln h'_{\text{mig}}] - \mathbb{E}[\ln h'_{\text{nat}}] \approx \rho \left(\ln h_{\text{mig}} - \ln h_{\text{nat}}\right) \quad (5)$$

Meanwhile, the child's friction gap is $\theta_{\text{mig},g=2} - \theta_{\text{nat}} = \lambda\kappa - 0 = \lambda\kappa$.[2] The immigrant child's income advantage, combining the ability advantage with the residual friction penalty, is approximately:

$$\Delta \approx \kappa(\rho - \lambda) \quad (6)$$

This expression reveals the core trade-off. The immigrant child benefits from inherited high ability ($\rho > 0$ amplifies the parent's latent advantage) but suffers from residual friction ($\lambda > 0$ means the penalty has not fully decayed). When $\rho > \lambda$ (that is, when ability is more persistent than friction), the immigrant child ends up ahead. This condition holds in the United States ($\rho = 0.23 > \lambda_{US} = 0.175$), placing the U.S. inside the region where the immigrant advantage obtains. In high-friction environments with $\lambda > \rho$, the condition fails and immigrant children underperform. The linearization above is a useful heuristic for the *direction* of the effect, but it substantially understates the *magnitude* because it ignores nonlinear amplification from differential parental investment (immigrant parents with higher latent h invest more, multiplying the inherited advantage).

Figure 2 illustrates the mechanism. The first generation faces a friction wedge κ that suppresses observed income below latent ability (Panel a); the second generation faces only the residual friction $\lambda\kappa$, which is approximately 5.2% under the baseline calibration (Panel b). The waterfall decomposition is in expected-income space; the resulting net income advantage translates to the 4.6 percentile-point headline gap reported in Table 2 once the income-to-rank transformation is applied. See Experiment C (Section 4.3) for the rank-space

[2] The generational log-wage recovery from the decaying wedge is approximately $\kappa(1-\lambda)$. This object has a reduced-form empirical counterpart: Borjas (1993) estimates a cross-ethnic-group regression of second-generation on first-generation log wages (Table 4, Row 1) and recovers a common intergenerational intercept of roughly 7%, which he attributes verbally to schooling quality, English proficiency, and reduced ties to ethnic enclaves but does not formalize as a structural parameter. The decaying wedge $\theta_g = \kappa\lambda^{g-1}$ in this paper can be read as a structural counterpart of that observation, separating friction decay from the Markov skill-transmission coefficient that his regression slope also estimates.

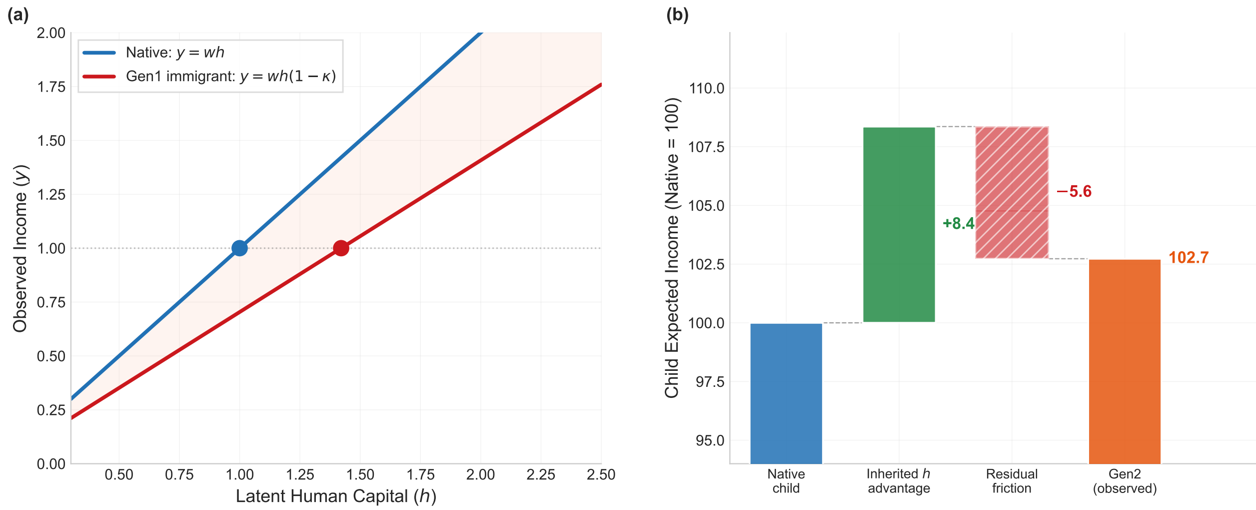

Figure 2: ***The selection-assimilation mechanism.*** *Magnitudes reflect the baseline calibration.* ***Panel (a)****: at any given observed income level (dashed line, normalized to 1), an immigrant has higher latent h than a native because the friction κ suppresses observed earnings below latent productivity. The 25th percentile is the empirical focal point.* ***Panel (b)****: Waterfall decomposition of expected Gen2 income, native child indexed to 100. Inherited latent h (green) raises the immigrant child above the native baseline; residual Gen2 friction λκ (red, hatched) offsets part of the gain, leaving a small net advantage. Income space used as h and the friction κ enter earnings additively and multiplicatively; the rank-space counterpart (larger in magnitude due to nonlinear rank translation) is reported in Table 2 and reconciled with the income-space decomposition in Section 4.3.*

decomposition.

2.6. Production

A representative firm produces a homogeneous final good using physical capital K and effective labor L with a Cobb-Douglas technology:

$$Y = A\,K^{\alpha}\,L^{1-\alpha} \tag{7}$$

where A is total factor productivity and $\alpha \in (0, 1)$ is the capital share. Factor prices are determined competitively:

$$r = \alpha A K^{\alpha-1} L^{1-\alpha} - \delta \tag{8}$$

$$w = (1-\alpha) A K^{\alpha} L^{-\alpha} \tag{9}$$

where δ is the depreciation rate of physical capital.

Effective labor aggregates latent ability net of frictions across all households:

$$L = \int h_s \cdot (1 - \theta_{g_s})\, d\mu(s) \tag{10}$$

where μ is the cross-sectional distribution of household states. The difference between total latent ability $\bar{L} = \int h_s\, d\mu(s)$ and effective labor L is the static labor-misallocation wedge. This is the productive capacity that exists but is not utilized at the observed equilibrium because of institutional frictions. This is the within-period analogue of the Hsieh-Klenow wedge; the full general-equilibrium output loss additionally captures how the wedge suppresses parental investment in children and the capital stock, and is quantified in Section 4.1.

2.7. Equilibrium

Definition 1 (Recursive Competitive Equilibrium)**.** A recursive competitive equilibrium consists of a value function $V(a, h, g)$, policy functions $\{c(s), a'(s), i(s)\}$, factor prices $\{r, w\}$, and a stationary distribution μ^* over household states such that:

(i) **Household Optimality:** Given prices (r, w), the policy functions solve the Bellman equation (1) subject to the budget constraint (2), the human capital law of motion (4), and the friction dynamics (3).

(ii) **Firm Optimality:** Factor prices satisfy the marginal productivity conditions (8)–(9).

(iii) **Capital Market Clearing:** $K = \int a_s\, d\mu^*(s)$.

(iv) **Labor Market Clearing:** $L = \int h_s(1 - \theta_{g_s})\, d\mu^*(s)$.

(v) **Stationarity:** The distribution μ^* is invariant under the transition dynamics induced by the policy functions, the human capital law of motion, and the demographic process (new immigrant inflow at rate m).

2.8. Observed Earnings and the Rank Mapping

The model delivers latent human capital h, factor prices (r, w), and generation-specific friction wedges θ_g for every household. The calibration targets in Section 3 and the cross-country exercises in Section 5, however, are rank-rank slopes and conditional percentile gaps rather than levels. This subsection states the mapping from the equilibrium objects to these rank-based observables; it also introduces one structural parameter, σ_w, that governs within-cohort dispersion in first-generation observed earnings without entering any Section 2 equation directly.

Observed earnings. Household i's observed labor income is

$$y_i = w\, h_i\, (1 - \theta_{g_i}) \cdot \xi_i, \qquad \xi_i = \begin{cases} \exp(\eta_i), \quad \eta_i \sim \mathcal{N}(0, \sigma_w^2) \text{ i.i.d.} & \text{if } g_i = 1, \\ 1 & \text{otherwise.} \end{cases} \tag{11}$$

The σ_w component adds an idiosyncratic, first-generation-specific multiplicative shock to observed earnings. It captures within-cohort heterogeneity in the realized wedge (e.g., individual-level language proficiency, credential match with specific employers, network-specific discrimination, and related barriers) that differs across Gen1 immigrants at the same latent h but does not transmit to children. Natives and later-generation immigrants face no such dispersion in the observation equation; their earnings depend only on latent productivity and the generation-specific friction. Children of Gen1 parents inherit latent h through the human-capital law of motion (Equation (4)), not the realized earnings shock, so ξ_i leaves no intergenerational trace. This is what distinguishes σ_w from the latent-ability persistence parameter ρ: ρ governs how h transmits across generations; σ_w governs how observed earnings scatter around latent earnings within the first-generation cohort in a given period. Gen2 and Gen3+ dispersion is not modeled separately because their residual frictions ($\lambda\kappa$ and $\lambda^2\kappa$, roughly 5% and 1% at the U.S. calibration) are small and plausibly more homogeneous; any additional within-cohort scatter at those generations would be largely absorbed into σ_ϵ, the intergenerational shock on latent h in Equation (4).

Rank mapping. For a simulated cross-section of N households, I compute parent and child percentile ranks against the pooled joint distributions:

$$r_i^{\text{par}} = \frac{1}{N} \cdot \#\{j : y_j^{\text{par}} \leq y_i^{\text{par}}\} \times 100, \tag{12}$$
$$r_i^{\text{ch}} = \frac{1}{N} \cdot \#\{j : y_j^{\text{ch}} \leq y_i^{\text{ch}}\} \times 100. \tag{13}$$

The reference distributions pool all households, natives and all migrant generations, in the parent cross-section and in the child cross-section respectively. Children's incomes are computed from inherited h and the same friction structure as their parents' cohort one generation forward, without any additional observational noise:

$$y_i^{\text{ch}} = w\, h_i^{\text{ch}}\, (1 - \theta_{g_i^{\text{ch}}}), \tag{14}$$

where h_i^{ch} is the child's latent human capital drawn from Equation (4) and g_i^{ch} is the child's generation index (one above the parent's, capped at the absorbing native state). This pooled-population rank construction matches the empirical convention in Abramitzky et al. (2021) and Boustan et al. (2025), who rank children against the full observed income distribution of the relevant generation rather than within subgroups.

3. Calibration

3.1. Data Sources

Five moments discipline the calibration, drawn from three data sources.

Opportunity Insights. The baseline intergenerational mobility statistics for native-born Americans come from Chetty, Hendren, Jones, and Porter (2020). I use the national-level rank-rank transition matrices (Table 1, pooled race) to compute the native rank-rank slope of approximately 0.33.

Abramitzky, Boustan, Jacome, and Perez (2021). The immigrant mobility advantage is measured using Table 1 of Abramitzky et al. (2021), which reports the regression:

$$\text{Rank}_{\text{child}} = 36.64 + 7.42 \cdot \mathbf{1}[\text{2nd gen}] + 0.33 \cdot \text{Rank}_{\text{parent}} - 0.08 \cdot (\mathbf{1}[\text{2nd gen}] \times \text{Rank}_{\text{parent}})$$

At the 25th percentile of parental income, the native child's expected rank is $36.64 + 0.33 \times 25 = 44.9$ and the immigrant child's expected rank is $(36.64 + 7.42) + (0.33 - 0.08) \times 25 = 50.3$. The gap of 5.4 percentile points is the primary mobility-advantage target. The second-generation rank-rank slope of $0.33 - 0.08 = 0.25$ serves as a secondary mobility target disciplining the overall steepness of the immigrant parent-child relationship.

IPUMS ACS. The median immigrant-to-native wage ratio is computed from IPUMS American Community Survey data (2015–2019) on prime-age workers (25–55). The conditional median ratio is approximately 0.92, corresponding to an 8% wage gap. Unlike the prior calibration, κ is not fixed externally but estimated jointly with the other structural parameters; the wage ratio provides the primary identifying moment for κ. Appendix B reports the supporting Mincerian regressions, which yield a conditional foreign-born coefficient of -4.3% after controlling for education and an unconditional coefficient of -9.3%. The calibrated $\kappa = 0.296$ sits near the upper end of the entry-cohort penalty range (15–30%) documented by Borjas (1993) and Lubotsky (2007), well above the cross-sectional Mincerian estimates.

3.2. Calibration Strategy

Five parameters are calibrated jointly by simulated method of moments (SMM): β (altruism/discount factor), λ (friction persistence), κ (first-generation penalty), $\bar{h}_{\text{mig}}$ (mean latent human capital of new immigrants), and σ_w (cross-sectional wage process dispersion). Ability persistence ρ is fixed externally at the Lee and Seshadri (2019) value of 0.23 (their Table 7, ρ_a, calibrated to match the Chetty et al. intergenerational earnings slope of 0.34). The remaining parameters (γ, δ, α, σ, σ_ε) are set to standard values from the literature.

The objective function is a weighted sum of squared relative errors on the five targeted moments. I optimize via adaptive Nelder-Mead with a tight convergence tolerance; the final reported moments are averaged over 20 independent simulation seeds with 400,000 agents per seed to minimize Monte Carlo noise. The model is solved by value function iteration with Howard's policy improvement on a discretized state space.[3]

Section 3.4 formalizes the identification argument by mapping each calibrated parameter to its dominant moment channel. Table 1 reports the full parameter vector.

3.3. Model Fit

Although the system is just-identified, the five moments span distinct parts of the data. A simple wedge mechanism fitting them jointly is itself a result that validates the theoretical mechanism that has been described informally in the literature. Table 2 reports the fit of the calibrated model to the five targeted moments. The second-generation rank-rank slope is matched within 1%, and the capital-output ratio and immigrant-to-native wage ratio are matched within 5%. The native rank-rank slope misses by approximately 7%, a tension induced by the external pin of ρ and discussed in Section 3.4. The second-generation gap at P25 is the largest remaining miss at 14%, reflecting the high sampling variance of this moment, which is computed within a narrow 10-percentile-point parental bin of first-generation immigrants.

The calibrated value $\lambda_{US} = 0.175$ implies that the second generation retains $0.175 \times 0.296 = 0.052$, or a 5.2% friction, a roughly six-fold reduction from the first generation's 29.6% penalty. By the third generation, the residual friction is $0.175^2 \times 0.296 = 0.009$, or 0.9%, consistent with the empirical finding that third-generation immigrants are economically indistinguishable from natives (Abramitzky et al., 2021).

[3] State space: 50 grid points for assets and 75 for latent human capital. Choice space at each state: 80 savings alternatives and 50 investment alternatives. Grid convergence was verified across both state and choice dimensions; reported moments are stable to further refinements.

Table 1: *Model Parameters*

Parameter	Description	Value	Target / Source
Panel A: Internally calibrated (SMM)			
β	Discount factor (altruism)	0.967	$K/Y = 3.0$
λ_{US}	Friction persistence (US)	0.175	Gen2 gap at P25 = 5.4 pp
κ	Gen1 penalty	0.296	Imm/nat median wage ratio = 0.92
$\bar{h}_{\text{mig}}$	Mean immigrant latent h	0.759	Joint (κ, wage ratio, Gen2 gap)
σ_w	Gen1 obs.-earnings disp.	0.332	Gen2 rank-rank slope = 0.25
Panel B: Set from data / literature			
ρ	Ability persistence	0.23	Lee and Seshadri (2019), Table 7 (ρ_a)
γ	HC inv. elast.	0.30	Cunha, Heckman, and Schennach (2010) Lee and Seshadri (2019)
α	Capital share	0.33	National accounts
δ	Depreciation rate	0.05	Standard (per generation)
σ	Risk aversion (CRRA)	2.0	Standard
σ_ε	Ability shock volatility	0.20	Central value (Lee and Seshadri, 2019)

Notes: Panel A parameters are calibrated jointly by simulated method of moments; Panel B parameters are set from data or standard values in the quantitative macro literature. The "Target / Source" column lists the primary identifying moment for each calibrated parameter (full identification argument in Section 3.4) or the external source for each fixed parameter.

3.4. Identification

With ρ externally pinned at the Lee-Seshadri (2019) value, the calibration is formally just-identified: five free parameters (β, λ, $\bar{h}_{\text{mig}}$, σ_w, κ) mapped to five SMM moments. Each parameter has a dominant identification channel. The capital-output ratio pins β, since it is essentially unaffected by any other parameter. The immigrant-to-native wage ratio then separates $\bar{h}_{\text{mig}}$ (which raises it) from κ (which lowers it); the second-generation gap at P25 resolves this pair through their joint effect on mobility and identifies λ as the residual lever on intergenerational friction decay; the second-generation rank-rank slope identifies σ_w through the observation equation (11) (higher σ_w scatters Gen1 parents' observed incomes at the same latent h, flattening the mapping from Gen1 parent rank to Gen2 child outcome); and the native rank-rank slope disciplines the overall transmission strength. Appendix A.5 reports a local Jacobian of the five moments with respect to the five free parameters, evaluated at the v6.2 calibration: the dominant-moment column pattern supports the argument in the preceding paragraph, the system has full effective rank 5/5, and a single weak direction combines λ, $\bar{h}_{\text{mig}}$, σ_w, and κ in a way that is approximately orthogonal to the model's headline outputs. Appendix A.1 then probes robustness of the

Table 2: *Moment Fit (20-seed averages with standard errors)*

Moment	Model	SE	Target	% Error
Capital-output ratio (K/Y)	2.83	0.00	3.00	−5.5
Native rank-rank slope	0.307	0.002	0.33	−6.8
Gen2 gap at P25 (pp)	+4.64	1.02	+5.4	−14.0
Imm/nat median wage ratio	0.949	0.004	0.92	+3.2
Gen2 rank-rank slope	0.251	0.004	0.25	+0.5

Notes: Model moments are averages across 20 simulation seeds with 400,000 agents per seed; the SE column reports the cross-seed standard error. % Error is $100 \times (\text{Model} - \text{Target})/\text{Target}$. Targets are as described in Section 3.1.

calibrated parameters to ±5% target perturbations: each calibrated parameter responds primarily to its hypothesized identifying moment, and the headline numbers are insensitive to small changes in any single target. Appendix A.3 reports an unconstrained variant in which ρ is included in the SMM optimization; the resulting SMM-selected value is $\rho = 0.255$ (within 11% of the external anchor), which restores the native rank-rank slope fit to within 1% and reduces the Gen2-gap miss to about 10%, while the static wedge shifts modestly from 1.57% to 1.34%. As an additional diagnostic on the identifying content of the wedge itself, Appendix A.4 reports a variant in which κ is fixed at zero and the remaining five parameters are freely re-optimized; this variant yields an immigrant-to-native wage ratio of 1.054 against the target of 0.92, flipping above unity, confirming that positive selection and skill persistence alone cannot simultaneously produce a positive Gen2 advantage and a below-unity first-generation wage ratio.

3.5. Robustness

Appendix A reports two robustness exercises. First, an Andrews-Gentzkow-Shapiro (2017) target sensitivity exercise perturbs each of the five calibration targets by ±5% and re-runs the full SMM optimization, tracing how the calibrated parameters and headline numbers respond to small changes in the empirical inputs. Second, a single-parameter sensitivity check perturbs σ_w (the parameter without a direct empirical counterpart) by ±20% and recomputes the moments and the static wedge. Both exercises show that the headline static wedge and the sign-flip threshold are quantitatively invariant within the relevant range of parameter and target uncertainty; the qualitative conclusions (positive Gen2 advantage, dominant assimilation channel, large concentrated first-generation loss) are unaffected.

Table 3: *Untargeted Moments (Overidentification Check)*

Moment	Model	Empirical	Source
Gen3 gap at P25 (pp)	$+0.41$	≈ 0	Abramitzky et al. (2021)
Gen2 gap at P65–75 (pp)	$+0.65$	$+1.42$	ABJP linear spec, $p = 75$

3.6. Untargeted Moment Validation

As an overidentification check, I compare the model's predictions on two moments that were *not* included in the calibration targets against their empirical counterparts (Table 3). The third-generation gap at P25 tests whether the model produces the empirical convergence-by-third-generation pattern documented by Abramitzky et al. (2021). The second-generation gap at P65–75 tests the model's cross-sectional shape: because the P25 moment is targeted, the P65–75 gap would be trivially close to the ABJP linear extrapolation ($7.42 - 0.08 \cdot p$) only if the model's implied profile were itself linear. The model instead produces a steeper profile than the linear ABJP specification (see Section 4.4), making the P65–75 gap a genuine out-of-sample moment that carries information about the model's cross-sectional fit.

The third-generation gap is essentially at parity ($+0.4$ pp), corroborating that the mechanism produces convergence to natives by the third generation as documented empirically by Abramitzky et al. (2021). The second-generation gap at P65–75 falls below the empirical extrapolation by approximately 0.8 pp, the signature of the slope discrepancy between the model's profile and the ABJP linear specification discussed in detail in Section 4.4. I flag this gap as a known limitation of the first-order model rather than a surprise: the model deliberately does not match the empirical cross-sectional shape at higher parental ranks because doing so would require the heterogeneous-friction extension also discussed in Section 4.4.

4. Results

I use the calibrated model to conduct three experiments. Experiment A anatomizes the U.S. output loss along two dimensions, by generation within the static wedge, and by economic channel within the full general-equilibrium loss, and explores how the loss responds to counterfactual changes in the fundamental parameters (κ, λ, $\bar{h}_{\text{mig}}$); this establishes where the policy leverage lies and how large the stakes are. Experiment B then maps how outcomes vary across the empirically observed range of friction persistence λ, locating the threshold at which the immigrant mobility advantage flips sign and setting up the

Table 4: *Static Labor-Misallocation Wedge by Generation (U.S. baseline)*

Generation	Share	Mean h	θ	Eff. labor	First-best	Share of loss
Native	85.7%	1.000	0.000	1.000	1.000	0%
Gen1	5.0%	1.379	0.296	0.970	1.379	87.1%
Gen2	4.7%	1.059	0.052	1.004	1.059	11.1%
Gen3+	4.5%	1.005	0.009	0.997	1.005	1.8%

Notes: All human-capital quantities are normalized to the native population mean ($\bar{h}_{\text{Native}} \equiv 1.000$); raw calibration values can be recovered by multiplying by the native mean of 0.573 in the model's internal h-index. Effective labor is $\bar{h}_g(1-\theta_g)$ and first-best is $\bar{h}_g$; share of loss is $n_g\bar{h}_g\theta_g$ divided by the total labor misallocation $\sum_g n_g\bar{h}_g\theta_g$ and is invariant to the normalization. Shares may not sum to 100% due to rounding.

cross-country analysis in Section 5. Experiment C decomposes the immigrant mobility advantage itself into selection and assimilation components.

4.1. Experiment A: Output Loss Anatomy and Counterfactuals

Experiment A anatomizes the U.S. output loss along two complementary dimensions. First, I decompose the static labor-misallocation wedge, which is the direct suppression of migrant productivity at the observed equilibrium distribution, into by-generation contributions. This static measure is the Hsieh-Klenow equivalent for the present setting: $L^{\text{FF}} - L$, where first-best effective labor $L^{\text{FF}} = \int h_s \, d\mu(s)$ is latent labor absent frictions. Summing each generation's contribution $n_g \cdot \bar{h}_g \cdot \theta_g$ (population share times mean latent human capital times friction wedge) shows where the static loss sits. Second, I solve the full general-equilibrium counterfactual in which $\kappa = 0$ is imposed and the model is re-solved to steady state, letting the human capital distribution and capital stock adjust endogenously. The difference between the static wedge and the full-GE loss reveals the magnitude of dynamic accumulation channels, suppressed intergenerational investment in children's human capital and suppressed household savings, which cross-sectional wedge accounting cannot capture by construction.

Within the static labor-misallocation wedge reported in Table 4, 87% is attributable to the first generation. First-generation immigrants face a 29.6% friction wedge that suppresses their market productivity well below their latent ability. This concentrated loss illustrates the core economic insight: the problem is not immigrant quantity but the suppression of immigrant productivity at arrival, which decays but does not disappear by the second generation (5.2% residual friction) and becomes negligible by the third generation and beyond (0.9%). The 87/11/2 breakdown across the first, second, and third generations means that any policy lever acting directly on κ (the first-generation friction) has roughly eight times the potential aggregate impact of a policy lever acting only on λ (the rate of intergenerational decay), though *both* frictions create distortions with quantitatively meaningful macroeconomic consequences.

The static wedge itself amounts to 1.57% of GDP, or roughly $440 billion at U.S. 2024 output levels. This figure represents only the direct labor misallocation visible to cross-sectional wedge accounting, however. When the model is re-solved to steady state under $\kappa = 0$, letting the human capital distribution and capital stock adjust endogenously, the full general-equilibrium output loss is 4.94%, approximately $1.4 trillion annually at 2024 levels and roughly three times the static wedge.[4] This magnitude sits within the range of moderate financial-friction estimates in the misallocation literature (Moll, 2014), above typical welfare-cost estimates for major tax distortions (1–3% of GDP), and substantially below the manufacturing TFP gaps documented by Hsieh and Klenow (2009) for developing economies, as would be expected for an advanced economy with comparatively mild distortions.

The 4.94% total decomposes into three economically distinct channels reported in Table 5. The static labor misallocation (1.57 pp, 32% of the total) corresponds to what a Hsieh and Klenow (2009)-style wedge analysis would deliver: the direct productivity loss from the observed distribution of frictions. The intergenerational human capital response (0.90 pp, 18%) reflects the suppression of parental investment in children's human capital: first-generation immigrants with suppressed earnings invest less in their children, lowering the latent human capital of future generations relative to the frictionless benchmark. The equilibrium capital response (2.47 pp, 50%) captures the downstream adjustment of the capital stock: when κ suppresses immigrant earnings, aggregate savings are lower and the capital-output ratio sits below what a frictionless economy would sustain.

The first two channels are labor-market primitives: they are the direct and dynamic effects of κ on immigrant labor productivity within and across generations. The third channel is economically downstream: the capital stock rises because wages and income rise, and households save more. Static wedge accounting, by construction, cannot see any of the last two channels, and the ratio of total loss to static wedge (roughly 3:1) is itself a quantitative statement about the magnitude of dynamic amplification in this setting. The dynamic misallocation literature (Restuccia and Rogerson, 2008; Moll, 2014) shows a

[4]The GE output loss is reported throughout as $(Y^{\text{ff}} - Y)/Y$, the gain from eliminating frictions expressed as a percentage of baseline GDP. This is the measure whose channel decomposition sums naturally to $1.57 + 0.90 + 2.47 = 4.94$ percentage points (Table 5). The alternative Hsieh-Klenow convention $1 - Y/Y^{\text{ff}}$ yields the numerically smaller 4.71% and is the measure reported in the "GE loss" columns of Tables 6 and 7, whose notes define the denominator explicitly.

Table 5: *Decomposition of the Full Steady-State Output Loss*

Channel	Loss (pp of GDP)	Share	Interpretation
Static labor misallocation	1.57	32%	Hsieh and Klenow (2009)-style wedge at baseline
Intergenerational HC transmission	0.90	18%	Parental investment response
Equilibrium capital response	2.47	50%	Savings/wealth adj. (downstream)
Total GE output loss	**4.94**	**100%**	Full GE counterfactual

Notes: Decomposition uses a labor-first causal ordering (wedge → investment → capital), reflecting the model's equilibrium structure: the capital stock adjusts in response to changes in effective labor rather than being shocked independently. The static component is computed at the baseline equilibrium with capital and distribution held fixed; subsequent channels are computed by sequentially allowing the human capital distribution and then the capital stock to adjust to their $\kappa = 0$ steady-state values.

similar phenomenon for firm-level distortions with dynamic capital accumulation; here the analogous mechanism operates through intergenerational human capital transmission in addition to capital-stock adjustment.

I now compare the U.S. baseline against five counterfactuals: (i) frictionless ($\kappa = 0$), (ii) full single-generation assimilation ($\lambda = 0$), (iii) higher immigrant selection ($\bar{h}_{\text{mig}}$ raised 20%), (iv) lower immigrant selection ($\bar{h}_{\text{mig}}$ lowered 20%), and (v) doubled barriers (κ doubled). Each is solved in full general equilibrium with endogenous prices and distribution. Table 6 reports both the static labor-misallocation wedge evaluated at each scenario's own equilibrium and the full GE output loss measured against the proper frictionless counterpart for that scenario.

Two patterns merit comment. First, the gains from eliminating barriers are concentrated on κ, not λ. Removing all initial frictions ($\kappa = 0$) raises GDP by 4.95% above the U.S. baseline and drives the GE output loss to zero. Setting $\lambda = 0$ (full single-generation assimilation) only reduces the GE output loss from 4.71% to 4.13% — a 0.58-percentage-point reduction, less than one-eighth of what removing κ achieves. The reason follows mechanically from the by-generation decomposition above: the largest source of static labor misallocation is the first-generation cohort itself, whose 29.6% penalty is untouched by changes to λ, which governs only intergenerational decay. The same asymmetry holds along the dynamic channels: suppressed Gen1 earnings are the dominant source of both reduced parental investment in children's human capital and reduced household savings. Policy design that targets first-generation labor market integration therefore yields far larger aggregate returns than policy that waits for convergence across generations.

Second, selection and friction reduction are strictly complementary policy levers. Raising

Table 6: *Experiment A: Counterfactual Scenarios*

Scenario	Y	K/Y	w	r	Static (%)	GE loss (%)
U.S. baseline	**0.9560**	**2.83**	**1.1181**	**0.0667**	**1.57**	**4.71**
No barriers ($\kappa = 0$)	1.0033	2.90	1.1314	0.0639	0.00	—
Full assimilation ($\lambda = 0$)	0.9619	2.83	1.1189	0.0665	1.37	4.13
High selection ($+20\%\ \bar{h}_{\text{mig}}$)	0.9819	2.86	1.1240	0.0654	1.81	5.45
Low selection ($-20\%\ \bar{h}_{\text{mig}}$)	0.9292	2.79	1.1100	0.0684	1.30	4.04
Doubled barriers (2κ)	0.9075	2.78	1.1064	0.0692	3.20	9.55

Notes: Each row is a full general-equilibrium solve with endogenous prices and steady-state distribution. Only the indicated parameter differs from the U.S. baseline; all others are held at calibrated values. "Static (%)" is the static wedge, $(Y^{\text{ff}} - Y)/Y$, at each row's own equilibrium, holding capital and distribution fixed — the Hsieh and Klenow (2009)-style measure. "GE loss (%)" is $1 - Y_{\text{row}}/Y^{\text{ff}}_{\text{proper}}$, where $Y^{\text{ff}}_{\text{proper}}$ is the $\kappa = 0$ counterfactual at the same selection level. For the baseline, $\lambda = 0$, and doubled-barriers rows, the frictionless benchmark is $Y = 1.0033$. For the $\pm 20\%\ \bar{h}_{\text{mig}}$ rows, the benchmarks are $Y = 1.0385$ and $Y = 0.9683$ respectively, from separate GE solves at $\kappa = 0$ and the corresponding selection level.

$\bar{h}_{\text{mig}}$ by 20% alone delivers a 2.71% GDP gain (to $Y = 0.9819$); removing frictions alone delivers 4.95% (to $Y = 1.0033$); pursuing both (a higher-selection admission regime paired with $\kappa = 0$) yields 8.63% (to $Y = 1.0385$), roughly one percentage point above the sum of the individual effects. Equivalently, higher-skilled immigration leaves more untapped potential when frictions remain: the GE output gap widens from 4.71% to 5.45% because all three loss channels (direct wedge, suppressed intergenerational investment, and suppressed capital accumulation) scale with latent human capital. The symmetric case confirms the mechanism: lowering $\bar{h}_{\text{mig}}$ by 20% narrows the GE output gap to 4.04%, because less latent capacity is available to be lost to friction. This supermodularity provides a sharp efficiency argument for pairing any skill-based admission reform with concurrent friction-reducing measures, and against evaluating either lever in isolation.

Examining the investment policy functions at first-generation migrant states reveals that first-generation parents invest more in their children when friction is higher. This compensating response arises because forward-looking parents anticipate friction decay: when the child's friction will be $\lambda\kappa$ rather than κ, the return to investment is elevated relative to the low-friction benchmark. In the high-friction counterfactual, parents invest more aggressively because they understand that the reduction from κ to $\lambda\kappa$ will unlock a larger share of their child's potential human capital. This "investment motive" partially offsets the direct income loss from higher friction but cannot fully compensate for the aggregate productivity loss.

The 4.94% headline should be read as a lower bound on the full welfare cost of immigrant frictions. The model treats aggregate labor supply as fixed along three margins: migration is modeled as replacement rather than additive inflow (new arrivals replace exiting households, keeping population constant); native labor supply is inelastic; and the migration rate m is exogenous. Each assumption biases the headline downward. Additive migration would raise both the static wedge channel (more suppressed migrant labor) and the human capital channel (more Gen2+ descendants benefiting from intergenerational transmission). Endogenous native labor supply under a standard Frisch elasticity would partially redistribute the capital channel toward labor-side channels as natives respond to the roughly 1.2% wage increase at $\kappa = 0$. Endogenous migration quantity (a more substantial extension) would further raise the total loss by admitting additional workers into the more productive frictionless economy. A richer labor-supply environment therefore points uniformly to a larger aggregate cost of frictions than the 4.94% figure captures.

4.2. Experiment B: The Friction Persistence Sweep

Experiment A established that the 4.94% U.S. GE output loss decomposes into static labor misallocation, suppressed intergenerational human capital investment, and downstream capital accumulation, and that κ is the dominant policy lever along all three channels. Experiment B turns to the secondary lever: how outcomes vary as λ (the rate at which friction decays across generations) changes, holding κ fixed at its calibrated value. This exercise has two purposes. First, it quantifies how much additional output loss is generated by slow intergenerational decay on top of the first-generation floor that Experiment A identified. Second, it locates the threshold at which the immigrant mobility advantage reverses sign, setting up the cross-country mapping exercise in Section 5, where empirical cross-country variation is mapped into variation in λ.

I solve the model for ten values of λ spanning the range from near-frictionless intergenerational decay ($\lambda = 0.05$) to extreme persistence ($\lambda = 0.85$), holding all other parameters at their U.S. calibrated values. Table 7 reports the results.

Two patterns emerge from these results (see also Figure 3). The λ dimension alone accounts for a meaningful share of the total GE output loss. With κ held at its calibrated value, the static wedge ranges from 1.41% at near-frictionless intergenerational decay to 2.85% at extreme persistence, and the full GE output loss rises from 4.23% to 8.60% across the same range. Moving in the opposite direction of the U.S. baseline, into the range characteristic of Continental European destinations ($\lambda \in [0.25, 0.54]$, see Table 11), the GE output loss rises from 4.71% (U.S.) to between 5.1% and 6.6% depending on the destination. Persistent intergenerational friction is a quantitatively significant macroeconomic cost, especially once the dynamic accumulation channels are included.

Table 7: *Experiment B: Varying Friction Persistence (λ)*

λ	Y/Y_{US}	Gen2 Gap (pp)	Static wedge (%)	GE loss (%)	Interpretation
0.05	1.005	+11.45	1.41	4.23	Near-frictionless
0.10	1.003	+8.37	1.47	4.47	Fast assimilation
0.15	1.002	+6.13	1.53	4.56	Fast assimilation
0.175	**1.000**	**+4.64**	**1.57**	**4.71**	**US baseline**
0.25	0.997	+0.00	1.67	5.04	Sign-flip threshold
0.30	0.995	−2.31	1.75	5.21	Past threshold
0.40	0.989	−6.05	1.91	5.77	Moderate European
0.50	0.983	−12.29	2.09	6.32	High friction
0.70	0.969	−20.24	2.50	7.66	Very high friction
0.85	0.959	−26.99	2.85	8.60	Extreme persistence

Notes: Each row is a full GE solve at the indicated λ with all other parameters held at their U.S. calibrated values (in particular $\kappa = 0.296$). Y/Y_{US} is steady-state output relative to the baseline. Gen2 gap is averaged across simulation seeds. Static wedge is the $(Y^{\text{ff}} - Y)/Y$ measure at each row's own equilibrium; GE loss is $1 - Y_\lambda/1.0033$ against the common $\kappa = 0$ frictionless benchmark (which is independent of λ since $\theta_g = \kappa\lambda^{g-1} = 0$ when $\kappa = 0$). The sign-flip threshold $\lambda^* \approx 0.25$ is found directly at the $\lambda = 0.25$ row, where the model produces a Gen2 gap of $+0.00$ pp.

A useful property of the decomposition emerges from comparing the two loss columns: the GE-to-static-wedge ratio is approximately 3.0 across the entire λ range, varying only between 2.97 and 3.07. This near-constant amplification indicates that the dynamic accumulation channels (suppressed intergenerational investment plus capital-market response) scale approximately linearly with the static wedge across the parameter range studied, and is not an artifact of the U.S. baseline calibration. A rough rule of thumb emerges: the full welfare cost of immigrant frictions in this class of models is approximately three times the cross-sectional wedge that a Hsieh and Klenow (2009)-style analysis would deliver.

A second pattern is that the immigrant mobility advantage reverses sign at a threshold close to the U.S. value. At $\lambda = 0.175$ the second generation enjoys a $+4.64$ pp advantage; at $\lambda = 0.30$ the advantage has collapsed to -2.31 pp. Solving the model directly at $\lambda = 0.25$ places the sign-flip threshold at exactly that value ($+0.00$ pp gap, Table 7), approximately 0.075 above the calibrated U.S. value. A modest deterioration in friction persistence would suffice to erase the immigrant upward mobility pattern that the U.S. currently exhibits.

4.3. Experiment C: Decomposing Selection and Assimilation

The immigrant mobility advantage could arise from two channels: *selection* (immigrants at a given income percentile have higher latent ability than natives at the same percentile)

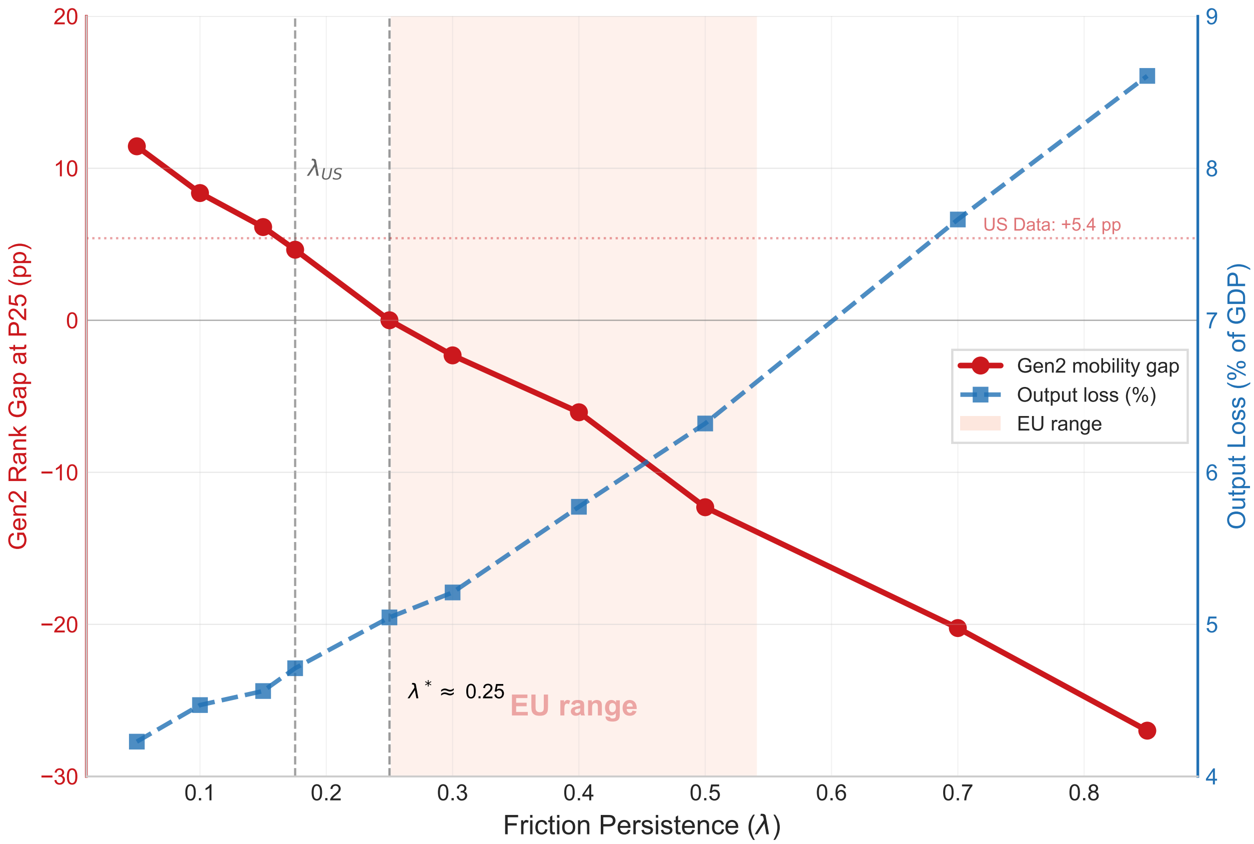

Figure 3: ***Friction persistence and the immigrant mobility advantage.*** *Gen2 mobility gap (left axis) and GDP relative to U.S. baseline (right axis) as functions of friction persistence λ. The vertical dashed line marks the sign-flip threshold $\lambda^* \approx 0.25$; the shaded region spans the calibrated λ values for Continental European destination countries (Table 11). The U.S. baseline ($\lambda_{US} = 0.175$) sits inside the positive-gap region but within the range of plausible institutional variation observed across OECD destinations.*

and *assimilation* (the friction θ decays from κ to $\lambda\kappa$ between the first and second generation). Experiment C disentangles these channels using a 2×2 Shapley decomposition, computed two ways: (i) a *fixed-price* version that holds equilibrium prices, policy functions, and the child income distribution at U.S. baseline values while toggling the channels, and (ii) a *full general-equilibrium* version that re-solves the model to steady state at each corner and measures the simulated Gen2 gap at P25 in the resulting equilibrium. The fixed-price version is a local channel decomposition valid at the baseline configuration; the GE version is a global counterfactual that allows prices, the human-capital distribution, and the capital stock to adjust endogenously in each corner.

The four scenarios are:

- **Both channels:** Migrant parent has h_{mig} (selected) and second-generation friction is $\lambda\kappa$ (assimilated).

- **Selection only:** Migrant parent has h_{mig} but friction persists ($\lambda = 1$, no assimilation).
- **Assimilation only:** Migrant parent has h_{nat} (no selection) but second-generation friction decays to $\lambda\kappa$.
- **Neither:** Migrant parent has h_{nat} and friction persists at κ.

Table 8 reports the second-generation rank gap at P25 under each scenario for both versions, along with the Shapley decomposition of the total improvement from "neither" to "both." Both versions agree that assimilation is the dominant channel, but they differ quantitatively: the fixed-price decomposition yields a 20/80 split (selection/assimilation), while the GE version shifts the split further toward assimilation at 13/87. The direction of the shift is informative: equilibrium responses *amplify* rather than attenuate the assimilation channel. Mechanically, the "neither" counterfactual is less extreme in GE (-33.25 pp) than in fixed-price (-44.1 pp) because GE allows the immigrant child's latent human capital to still accumulate partially through ρ-persistence even under permanent friction, cushioning the worst-case outcome. The total improvement from "neither" to "both" is therefore smaller in GE (37.9 pp vs. 49.1 pp fixed-price), and since the assimilation channel's absolute contribution is preserved, its share of the smaller total is larger.

Across both methodologies, selection alone (positive selection with permanent friction) yields strongly negative Gen2 gaps (-39 FP, -32.5 GE), while assimilation alone (common latent ability but friction decays) produces near-native outcomes (-9.7 FP, -4.2 GE). Friction decay is therefore the binding constraint: it is what allows the latent-ability advantage to manifest in observed outcomes. This conclusion is robust across both local (fixed-price) and global (full-GE) decompositions, and the GE version strengthens the quantitative claim rather than weakening it.

A first-order analytical approximation to the immigrant advantage in *log-income* units takes the form $\Delta \log y \approx \kappa(\rho - \lambda)$, which for the U.S. calibration gives $0.296 \times (0.23 - 0.175) = 0.016$: a roughly 1.6% expected-income gap at the P25 parent-income level. Converted into percentile-rank units through the model's simulated child-income density at P25, this log-income gap corresponds to a rank gap on the order of one percentile point — still substantially smaller than the full-model gap of 4.64 pp. The residual roughly three percentage points come from the nonlinear amplification of differential parental investment: immigrant parents with higher latent human capital invest more in their children, multiplying the inherited advantage beyond what the direct ρ-persistence would imply. This reconciles the income-space waterfall in Figure 2 with the rank-based moment targeted in Table 2. The approximation nonetheless highlights the correct qualitative condition: the sign of the immigrant mobility advantage depends on whether ρ exceeds λ.

Table 8: *Experiment C: Selection vs. Assimilation Decomposition (Fixed-Price vs. Full GE)*

	Fixed-price		*Full GE re-solve*	
	Assim ($\theta = \lambda\kappa$)	No Assim ($\theta = \kappa$)	Assim ($\theta = \lambda\kappa$)	No Assim ($\theta = \kappa$)
Selection ($h = h_{\text{mig}}$)	+5.0	−39.0	+4.65	−32.52
No Selection ($h = h_{\text{nat}}$)	−9.7	−44.1	−4.21	−33.25
Shapley decomposition of the "neither"→"both" improvement:				
Total improvement (pp)	49.1		37.9	
Selection contribution	+9.9 pp (**20%**)		+4.79 pp (**12.6%**)	
Assimilation contribution	+39.2 pp (**80%**)		+33.10 pp (**87.4%**)	

Notes: All gaps are second-generation rank gaps at P25. *Fixed-price* columns hold equilibrium prices, policy functions, and the child income distribution at U.S. baseline values while toggling the selection (h_{mig} vs. h_{nat}) and assimilation ($\theta = \lambda\kappa$ vs. $\theta = \kappa$) channels; the gap is computed using an analytical representative-agent shortcut at P25. *Full GE re-solve* columns re-solve the model to steady state at each corner's parameter values and measure the simulated Gen2 gap at P25 in each new equilibrium (20 simulation seeds per solve, 400,000 agents). Shapley contributions are the average of the two orderings in which each channel is added to the "neither" baseline. Replication script: `model/run_ge_decomposition_all.py`.

In the United States, this condition holds with a modest margin ($0.23 > 0.175$). The sign-flip threshold from the full model ($\lambda^* \approx 0.25$) is above $\lambda_{US} = 0.175$, but still within the range of institutional variation observed across OECD countries. In a high-friction environment with $\lambda > \rho$, the condition fails and immigrant children underperform natives.

4.4. Testable Implications

Beyond rationalizing the findings to which it was calibrated, the model produces two testable predictions for patterns the empirical literature has not yet examined directly. It also generates a U.S. second-generation profile across the full parent income distribution that matches ABJP's linear specification in direction, at the targeted P25 level (fit by calibration to within 0.8 pp), and in qualitative shape, with one quantitative caveat on slope magnitude that I discuss below.

(i) The calibrated first-generation friction sits near entry-cohort empirical estimates, consistent with the model's one-period abstraction. The calibrated $\kappa = 0.296$ exceeds the cross-sectional immigrant wage penalty estimated from IPUMS Mincerian regressions (4–9%, see Appendix B) and sits near the upper end of entry-cohort estimates derived from longitudinal data (15–30%, e.g. Borjas, 1993, Lubotsky, 2007). The model's one-period-per-generation structure imposes a single friction value on the entire first-generation life, so κ does not have a direct cross-sectional or longitudinal counterpart; it is identified here by the second-generation mobility gap, through its role in suppressing Gen1 earnings and thus parental investment in children's human capital. The comparison to empirical ranges is therefore a plausibility check on magnitude rather than an identification argument. That κ lands closer to entry-cohort estimates than to cross-sectional Mincerian gaps is consistent with the model's implicit timing: parental investment decisions in the model are made when the Gen1 parent is economically active, and are most strongly shaped by earnings near the start of their career — the period when empirical entry-cohort penalties are largest. A richer model with explicit years-since-arrival dynamics would decompose κ into an entry penalty plus a within-lifetime decay, mapping more directly onto longitudinal wage assimilation profiles (Lubotsky, 2007); the uniform one-period κ here is a parsimonious reduced form for those within-cohort dynamics, and I return to this extension in the conclusion.

(ii) Cross-country variation in λ should generate comparative differences in the slope of the second-generation profile. Destinations with more persistent intergenerational friction (higher λ) should exhibit *flatter* second-generation profiles across the parent income distribution, because persistent friction prevents the second-generation advantage from materializing at any parent rank. Lower-λ destinations should show steeper downward-sloping profiles — positive at the bottom, attenuating and eventually flipping sign at the top. This cross-country *slope* prediction can be tested against existing harmonized rank-rank regressions (e.g. Boustan et al., 2025) without requiring new functional-form assumptions. The caveat is that the model's U.S. slope is itself somewhat steeper than ABJP's linear estimate (-0.12 vs. -0.08/percentile), so the prediction is sharper about *comparative* slope differences across destinations than about absolute slope levels.

Figure 4 compares the model-predicted profile (red points) to ABJP's linear specification (dashed grey) for the United States. The two agree on direction, at the targeted P25 level (the model hits $+4.64$ pp against a $+5.4$ pp target), and on qualitative shape: a monotonically declining advantage from $+7.0$ pp at P5–15 to -3.6 pp at P85–95, with a sign flip near P70. The slope is steeper in the model (-0.12/percentile) than in ABJP's linear fit (-0.08), because at high parent income immigrants are less likely to be borrowing-constrained and the latent-h gap relative to natives shrinks, amplifying the uniform friction wedge and accelerating the sign flip. Two readings of the discrepancy are consistent with the model. First, ABJP's linear specification imposes constant slope by construction; their binned ventile data (their Figure 2 Panel D) show some tail curvature that may mask a genuine flattening or sign change at the top. Second, the uniform-κ assumption may

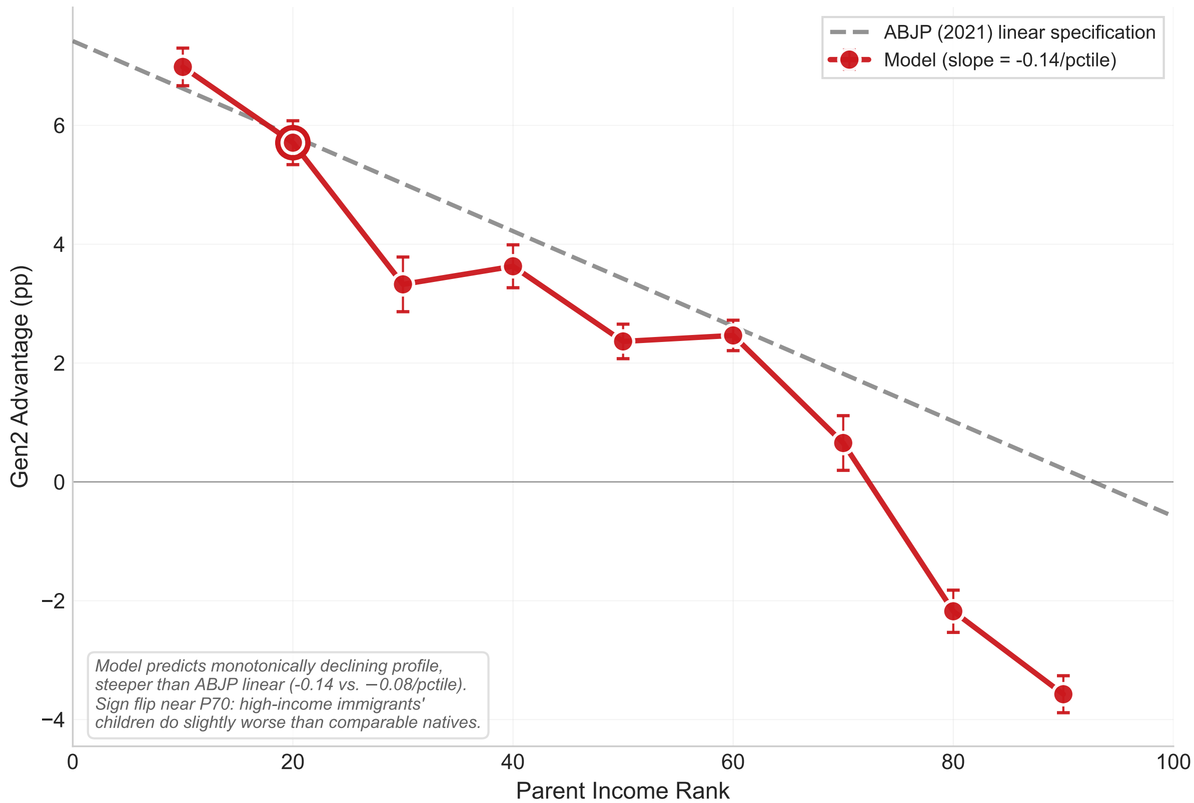

Figure 4: ***The second-generation advantage profile: model vs. empirical pattern.*** *Red points and connecting line: model-predicted second-generation advantage in 10-percentile-point bins of parent income, computed from the baseline calibration with multi-seed averaging (20 simulation seeds × 400,000 agents); 95% confidence intervals based on cross-seed variance are shown. Dashed grey line: the linear-in-rank specification from Abramitzky et al. (2021) Table 1,* $7.42 - 0.08 \cdot p$*. The model predicts a monotonically declining profile that is qualitatively consistent with the ABJP linear specification: both show the largest second-generation advantage at low parent income, declining toward zero and eventually turning negative at high parent income. The model's slope (*-0.12*/percentile) is steeper than the empirical linear specification (*-0.08*/percentile), with a sign flip near P70. The level at P25 is pinned by calibration construction.*

overstate friction for high-h immigrants whose skills are more internationally portable; allowing $\kappa(h) = \kappa \cdot (h_{\text{ref}}/h)^{\kappa_1}$ would flatten the upper-tail slope toward the data. I leave this heterogeneous-friction extension outside the baseline because the five calibrated moments cannot separately identify a friction gradient; a conditional-variance moment at a given parent rank would provide the discipline that κ_1 requires.

The model complements the existing empirical literature on three fronts: rationalizing the facts on which it is calibrated, generating new out-of-sample targets (longitudinal entry penalties and cross-country slope comparisons), and isolating uniform κ as the structural assumption whose relaxation could close the remaining slope gap at high parent income. I discuss additional theoretical extensions in the conclusion.

5. Origin Heterogeneity

The baseline calibration treats all immigrants as a single group with common entry penalty κ and friction persistence λ. In practice, both parameters likely vary across origin countries, reflecting differences in language distance, credential portability, and institutional reception. This section exploits this variation through a pair of mapping exercises: it calibrates group-specific friction parameters — origin-specific (κ_g, λ_g) for within-U.S. variation and destination-specific λ_j for cross-country variation — to the relevant second-generation moments, and examines whether the resulting distribution of friction parameters is plausible in magnitude and institutionally coherent.

5.1. Data and Method

For each origin group g, I set κ_g from Mincerian wage regressions on IPUMS American Community Survey data (2015–2019):

$$\ln(\text{hourly wage}) = a + b_g \cdot \mathbf{1}[\text{born in country } g] + f(\text{age}) + \psi(\text{education}) + u$$

estimated by weighted least squares on prime-age workers (25–55) with IPUMS person weights. The comparison group is US-born natives. The conditional specification (including education dummies) isolates the residual wage gap after controlling for human capital composition, ensuring that κ_g captures institutional friction rather than skill differences. Full regression results, including unconditional specifications and a discussion of the credential non-recognition penalty, are reported in Appendix B.

The target second-generation gap for each origin comes from the Opportunity Insights data (Chetty et al., 2020), which reports predicted kid family income rank (kfr) at the 25th percentile of parental income by parent country of birth (Table 6a). I compute the gap as $\text{kfr}_{P25}(\text{country } g) - \text{kfr}_{P25}(\text{US-born})$. I hold $(\beta, \rho, \gamma, \sigma, \sigma_\varepsilon, \delta, \alpha)$ at their baseline values and search over $\lambda_g \in [0.01, 0.90]$ via bisection to match each origin's observed second-generation gap.

5.2. Results

Table 9 reports the results for three modern origin groups where the model successfully matches the target. The fourth (Vietnam) is a boundary case discussed below. The entry penalty κ_g ranges from 0.14 (Mexico) to 0.26 (Cuba), and the calibrated λ_g varies accordingly.

Table 9 reveals substantial cross-origin variation in both parameters. Cuba faces the largest conditional wage penalty (26%) but the fastest friction decay ($\lambda = 0.18$), consistent

Table 9: *Origin Heterogeneity: Calibrated Friction Persistence by Country*

Origin	κ_g	λ_g	Gen2 Gap (pp)	Target (pp)	Static wedge (%)	GE loss (%)	Sel. (%)	Assim. (%)
US Aggregate	*0.296*	*0.175*	+4.64	+5.4	*1.57*	*4.71*	*13*	*87*
Mexico	0.14	0.304	+1.94	+2.0	0.82	2.46	29	71
Dominican Rep.	0.18	0.256	+2.06	+2.7	1.01	3.04	22	78
Cuba	0.26	0.179	+6.00	+6.2	1.38	4.13	16	84

Notes: κ_g from IPUMS ACS 2015–2019 conditional wage regressions; Gen2 gap targets from Opportunity Insights Table 6a. For each origin, λ_g is found by bisection to match the target, holding other parameters at U.S. baseline. GE loss is against the common frictionless benchmark ($Y = 1.0033$) and measures the cost of each origin's friction parameters applied to a U.S.-calibrated economy. Shapley shares are from the full-GE 2×2 re-solve (Table 8).

with its +6.2 pp Gen2 advantage; Mexico has the smallest penalty ($\kappa = 0.14$) paired with slower decay ($\lambda = 0.30$), matching its much smaller empirical gap. The Dominican Republic sits in between. For these three origin groups both κ_g and the observed Gen2 gap rank in the same order (Cuba $>$ DR $>$ Mexico), while λ_g ranks inversely; the three-group sample is too small to separate the independent contribution of friction decay from that of the first-generation penalty.

The selection–assimilation decomposition (full-GE Shapley shares, Table 9) reveals a clean pattern linked to the entry penalty κ_g. Cuba, with the highest $\kappa_g = 0.26$, mirrors the U.S. aggregate at 16/84; the Dominican Republic ($\kappa_g = 0.18$) sits at 22/78; and Mexico ($\kappa_g = 0.14$) at 29/71 shows the largest selection share across all origin groups. The pattern is interpretable: when the first-generation penalty is smaller, less of the total improvement to "both channels" comes from friction decay (there is less friction to decay), so the selection channel's relative contribution rises. This variation is masked by the fixed-price Shapley, which yields a narrower 18–22% range across these three origins; the GE re-solve reveals the underlying heterogeneity. Static wedges scale mechanically with the friction wedge $\kappa_g \cdot$ (share first generation), ranging from 0.82% (Mexico) to 1.38% (Cuba); these fall below the aggregate U.S. static wedge of 1.57% because each origin counterfactual holds only that group's parameters distinct from baseline, while the aggregate U.S. calibration blends all origins. GE output losses scale proportionally, from 2.46% (Mexico) to 4.13% (Cuba), each maintaining the roughly 3× amplification over the static wedge documented in Section 4.2. As with the cross-country exercise, these GE magnitudes are best interpreted as the efficiency cost of applying each origin's friction parameters to the U.S.-calibrated economy rather than as country-of-origin welfare measurements.

These origin-specific λ_g translate each group's reduced-form Gen2 gap into a model-implied measurement of intergenerational friction persistence. This is a structural object with empirical counterparts such as ethnic enclave persistence, credential portability, and origin-specific third-generation convergence speeds. Future empirical work could measure these counterparts independently to validate or falsify the mechanism, motivating extensions that could sharpen the model's quantitative predictions.

5.3. Historical Corroboration: The Age of Mass Migration

A further application of the mechanism is to the Age of Mass Migration (1880–1920). Abramitzky et al. (2014) document the intergenerational outcomes of sons of immigrants from that era using linked historical census records, and report large second-generation advantages for certain origin groups. I hold the baseline household-side parameters at their modern U.S. values and substitute κ_g with occupation-score-based wage penalties from Abramitzky et al. (2014); for each historical origin I then bisect λ_g to match that group's documented Gen2 gap. This is an illustrative exercise rather than a full recalibration: period-specific capital-output ratios, native rank-rank slopes, and wage dispersion are not available at the historical sample's resolution, so I retain the modern baseline for those moments.

Table 10 reports results for two origin groups. Italian immigrants faced a first-generation occupation-score penalty of 38%, among the largest of the era, paired with a modest Gen2 advantage of +2.5 pp. The model delivers this combination at $\lambda_g = 0.186$, close to the modern U.S. aggregate $\lambda_{US} = 0.175$. Russian and Eastern European immigrants of the same cohort faced a smaller first-generation penalty of 19%, yet their sons realized a Gen2 advantage of +11 pp — the "fast convergence" pattern that motivates much of the Abramitzky et al. (2014) analysis. Matching this combination requires $\lambda_g = 0.039$, roughly a fifth of the modern U.S. value and indicating near-frictionless intergenerational decay for that group.

Two features of the historical results stand out. First, the full-GE selection–assimilation decomposition is broadly in line with the modern U.S. 13/87 baseline: assimilation accounts for 87% of the Italian advantage and 84% of the Russian/E. European advantage. The mechanism's qualitative message — that friction decay is the dominant channel — is stable across a century of institutional change, shifts in origin composition, and differences in measurement technology (occupation scores in the historical data versus earnings percentiles today). Second, the static-wedge scale is on the same order of magnitude as the modern U.S. baseline of 1.57%, at 2.04% for Italy (driven by higher κ) and 0.89% for Russia/Eastern Europe (driven by lower κ paired with near-zero λ); the corresponding GE output losses (6.11% and 2.60%) preserve the roughly 3× amplification ratio. The relative

Table 10: *Historical Origins: Age of Mass Migration (1880–1920)*

Origin	κ_g	λ_g	Gen2 Gap (pp)	Target (pp)	Static wedge (%)	GE loss (%)	Sel. (%)	Assim. (%)
Italy	0.38	0.186	+2.32	+2.5	2.04	6.11	13	87
Russia / E. Europe	0.19	0.039	+11.09	+11.0	0.89	2.60	16	84

Notes: κ_g and Gen2 targets from Abramitzky et al. (2014)'s occupation-score wage regressions; other parameters held at modern U.S. baseline. λ_g is found by bisection to match the target. GE loss is against the common frictionless benchmark ($Y = 1.0033$) and measures the cost of these historical parameters applied to a modern U.S.-calibrated economy. Shapley shares are from the full-GE 2×2 re-solve (Table 8).

magnitudes reflect the general-equilibrium tension documented in Experiment A: within an origin group, the output cost of the friction wedge scales primarily with κ and the first-generation population share, with λ playing a secondary role.

The historical exercise is best interpreted as a corroboration rather than a test. Because κ_g is set externally and only λ_g is bisected to the Gen2 gap, the within-cohort fit is mechanical. The substantive content is twofold: the λ_g values required to match historical data are positive but modest, consistent with the U.S. being a low-friction regime across eras; and the full-GE selection–assimilation decomposition is stable at roughly 15/85 across radically different immigrant cohorts. A full recalibration to the 1880–1920 period would require historical capital-output ratios, period-specific native rank-rank slopes, and entry-cohort wage data that are beyond the reach of currently available linked administrative records; I view this as a natural extension.

5.4. Boundary Cases

Vietnam illustrates a useful limitation of the model. Vietnamese immigrants face only a 6% conditional wage penalty (the smallest in the sample), yet their children enjoy the largest second-generation advantage (+17.1 pp). The model cannot match this combination: even at the minimum feasible λ, $\kappa = 0.06$ generates a second-generation gap of only +9.4 pp, well below the +17.1 pp target. The mechanical reason is that such a small friction wedge cannot generate enough latent-ability advantage at P25 to produce a 17-point second-generation gap, regardless of how fast the friction decays.

The Vietnamese pattern likely reflects positive selection on unobservable characteristics (cultural emphasis on educational investment, refugee-cohort motivation, or favorable selection through refugee screening) that the model would attribute to $\bar{h}_{\text{mig}}$ rather than κ. This group's second-generation advantage operates primarily through the selection channel in ways that a single friction parameter cannot capture, and provides a natural motivation for extending the model to allow origin-specific human capital distributions or a Roy-model selection block.

5.5. Destination Heterogeneity: Mapping Institutional Variation

The preceding exercises exploit variation across immigrant *origin* groups within the United States. Turning to cross-country variation, I calibrate a country-specific λ_j for each destination by bisection to match its observed second-generation gap. This extends to the intergenerational margin a comparative tradition previously focused on first-generation wage and employment assimilation (Antecol et al., 2006). Countries differ in their institutional environments (credential recognition, labor market flexibility, anti-discrimination enforcement), and the model summarizes these through a single parameter λ. The exercise asks what value of λ a U.S.-calibrated model requires to rationalize each destination's observed Gen2 gap.

The sign-flip threshold $\lambda^* \approx 0.25$ plays a central interpretive role here. It is *emergent* from the U.S. calibration: given the five parameters disciplined by U.S. moments in Section 3, the model's Gen2 gap crosses zero at $\lambda \approx 0.25$, but this crossing is not a calibration. That the threshold lands within the range of plausible institutional variation across OECD destinations is itself a non-trivial feature of the calibration. The model also delivers a monotone Gen2-gap-in-λ mapping with a unique threshold — not a definitional property, since GE adjustments to wages, interest rates, and the pooled-rank distribution could in principle reverse the direct friction effect, and a poorly-structured mechanism could yield multiple threshold crossings. Given this monotonicity, bisecting each country's observed gap to a corresponding λ_j is well-defined. The informative content of the exercise lies in the *magnitudes* of the required λ_j (which could have been implausible or dispersed) and their *distributional structure* (clustering by institutional type).

I exploit the harmonized cross-country estimates of Boustan et al. (2025), who estimate rank-rank regressions analogous to Abramitzky et al. (2021) for second-generation immigrants in 15 destination countries using linked administrative or survey data. For each country j, I compute the second-generation gap at the 25th percentile of parental income as $\hat{\beta}_{m,j} + \hat{\beta}_{mp,j} \times 25$, where $\hat{\beta}_{m,j}$ and $\hat{\beta}_{mp,j}$ are the estimated coefficients on the immigrant father dummy and its interaction with parental rank, respectively.[5]

[5] Regression coefficients are taken from the country-specific appendix tables in Boustan et al. (2025): Table C.1.3 (Denmark), Table C.10.24 (Norway), Table C.3.8 (Austria), Table C.6.3 (Germany), Table C.14.3 (UK), and Table C.4.19 (Canada). For France, Sweden, and the Netherlands, I use the figure legend coefficients from Figures C.5.28, C.11.37, and C.9.15. Germany and the UK rely on survey data with small immigrant samples ($N = 206$ and 263, respectively), so their estimates are imprecise.

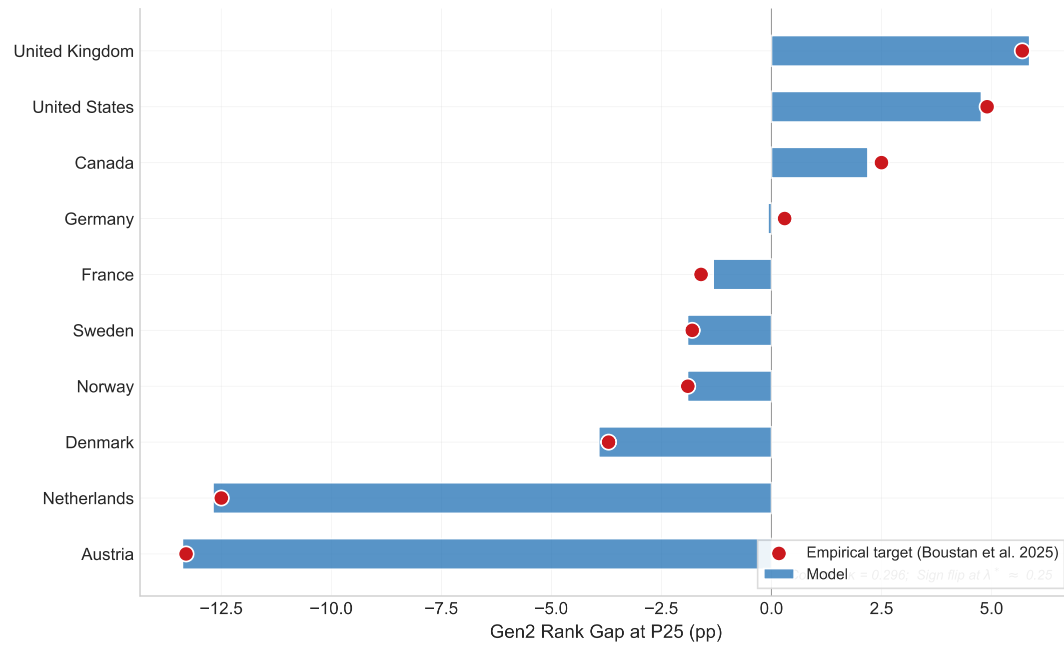

Figure 5: ***Cross-country distribution of second-generation gaps.*** *Bars are model-generated second-generation gaps at each country's calibrated λ_j (sorted by λ_j, smallest to largest); red dots are the empirical targets from Boustan et al. (2025).*

To isolate the role of institutional friction persistence, I hold $\kappa = 0.296$ for all countries at the U.S. estimated value. This common-κ assumption is supported by Han and Hermansen (2024), who show that the conditional immigrant wage gap in Norway is 19–21%, in the same range as the U.S. value. With κ held fixed, cross-country variation in the second-generation gap maps directly into variation in λ.

Table 11 reports the results. For each country, I solve the full general equilibrium model and search over λ to match the observed second-generation gap, exactly as in the origin-group exercise. Because cross-country comparability requires a single, consistently measured Gen2 moment across destinations, I use the Boustan et al. (2025) estimate for every country including the United States. For the U.S., Boustan et al. (2025) reports a harmonized gap of +4.9 pp, which differs from the +5.4 pp estimate from Abramitzky et al. (2021) used as the baseline U.S. calibration target in Table 2. The resulting U.S. row in Table 11 ($\lambda = 0.172$, Gen2 gap +4.77, static wedge 1.56%) therefore differs very slightly from the main calibration ($\lambda_{US} = 0.175$, Gen2 gap +4.64, static wedge 1.57%); this reflects the switch in moment definition for cross-country consistency, not a separate calibration of the U.S. baseline.

Calibrated λ varies by a factor of nearly 3.5 across destinations (see also Figure 5), from 0.157 (UK) to 0.539 (Austria), with the countries sorting naturally into three clusters: Anglophone destinations with low friction ($\lambda \approx 0.16$–0.23; UK, US, Canada), Continental

Table 11: *Destination Heterogeneity: Calibrated Friction Persistence across Countries*

Destination	λ	Gen2 Gap (pp)	Target[a] (pp)	Static wedge (%)	GE loss (%)	Sel. (%)	Assim. (%)
Positive second-generation advantage ($\lambda < \lambda^*$):							
United Kingdom†	0.157	+5.87	+5.7	1.54	4.60	13	87
United States	0.172	+4.77	+4.9	1.56	4.69	13	87
Canada	0.227	+2.19	+2.5	1.64	4.94	14	86
Negative or near-zero Gen2 gap ($\lambda > \lambda^*$):							
Germany†	0.252	−0.08	+0.3	1.68	5.05	13	87
France	0.274	−1.32	−1.6	1.71	5.12	13	87
Norway	0.289	−1.91	−1.9	1.73	5.17	14	86
Sweden	0.289	−1.91	−1.8	1.73	5.17	14	86
Denmark	0.348	−3.93	−3.7	1.82	5.48	15	85
Netherlands	0.509	−12.69	−12.5	2.11	6.38	16	84
Austria	0.539	−13.38	−13.3	2.16	6.58	17	83

Notes: All countries calibrated with $\kappa = 0.296$ and remaining parameters at U.S. baseline; λ is found by bisection to match the target. GE loss is against a common frictionless benchmark ($Y = 1.0033$); it measures the GDP cost of country j's λ applied to a U.S.-calibrated economy, not country j's own GDP loss. [a]Target from Boustan et al. (2025), Specification 1 (sons), at P25; the U.S. row uses the Boustan et al. (2025) estimate of +4.9 pp rather than the +5.4 pp used in Table 2. †Germany and UK estimates use small samples ($N = 206$ and 263); treat as suggestive. Shapley shares are from the full-GE 2×2 re-solve (see Table 8).

and Scandinavian European destinations with moderate friction ($\lambda \approx 0.25$–0.35; Germany, France, Norway, Sweden, Denmark), and high-friction destinations ($\lambda \approx 0.51$–0.54; Netherlands, Austria).

The calibrated λ_j values cluster by broad institutional type. The three Anglophone destinations (UK, US, Canada) sit below the U.S.-derived sign-flip threshold at $\lambda_j \in [0.16, 0.23]$; six Continental European and Scandinavian destinations sit immediately above it at $\lambda_j \in [0.25, 0.35]$; and the two high-friction destinations (Netherlands, Austria) sit far above at $\lambda_j \in [0.51, 0.54]$. Germany is the decisive borderline case — its calibrated $\lambda_j = 0.252$ sits essentially *at* the U.S.-derived threshold, consistent with its near-zero empirical gap of +0.3 pp. The informative content of this exercise is the joint structure: a threshold λ^* that emerges from U.S. moments alone, required λ_j values that are plausible in magnitude, and a cross-country distribution that lines up with broad institutional type (common-law Anglophone at the low end, codified-credential European destinations in the middle, the high-friction tail at the top). This is consistent with the model's structural interpretation of λ as reflecting cross-country differences in credential portability, labor market openness, and anti-discrimination enforcement.

As a first external check on this structural interpretation, I correlate each destination's calibrated λ_j with its MIPEX 2020 overall integration score — a pre-existing cross-country index of immigrant integration policy constructed independently of the mobility moments used here.[6] The Pearson correlation is -0.48 (Spearman $\rho = -0.40$): destinations with stronger overall integration policies tend to require lower calibrated λ_j, in the direction the model would predict. At $n = 10$ the relationship is not statistically significant ($p = 0.25$ for Spearman), but its economic magnitude is non-trivial: the OLS slope implies that moving from the lowest-MIPEX country in the sample (Austria, 46) to the highest (Sweden, 86) would shift the calibrated λ from roughly 0.39 to 0.20, comparable to the cross-country variation we recover from observed Gen2 gaps. Notably, MIPEX's *Labour Market Mobility* subscore alone — which measures first-generation access conditions (credential recognition, language training, anti-discrimination in hiring) — is essentially uncorrelated with λ_j ($r = -0.06$), while the broader overall index shows the negative relationship. This pattern is consistent with the model's interpretation of λ as governing *intergenerational* friction persistence: what empirically predicts the second-generation gap is not the policy regime facing first-generation arrivals in isolation, but the broader integration environment (schools, anti-discrimination, access to nationality, social mobility institutions) in which second-generation children grow up. A formal test would require a larger cross-section of destinations or a longer panel of MIPEX vintages.

Translating this institutional variation into output magnitudes, the implied GE output loss ranges from 4.60% (UK) to 6.58% (Austria), with Continental European destinations clustering at 5.0–5.5% and the high-friction tail above 6%. Because only λ varies across destinations here (all other parameters are held at U.S. baseline), these numbers are best read as the GDP cost, in a U.S.-calibrated economy, of adopting country j's institutional friction persistence. This isolates the contribution of institutional persistence from other cross-country differences.

Underlying this threshold result, the full-GE selection–assimilation decomposition is remarkably stable across all ten destinations: assimilation contributes 83–87% of the Gen2 outcome, with selection contributing the remaining 13–17%. The high-friction tail (Netherlands 84%, Austria 83%) shows slightly lower assimilation shares than the low-friction destinations (UK, U.S., Germany, France all at 86–87%), but the narrow range

[6] MIPEX (Migrant Integration Policy Index) is a 0–100 composite measure published by the Migration Policy Group that aggregates de jure policy measures across eight domains: labour market mobility, family reunification, education, political participation, permanent residence, access to nationality, anti-discrimination, and health. 2020 overall scores for the ten destinations range from 46 (Austria) to 86 (Sweden), with the United States at 73. See https://www.mipex.eu.

indicates that the dominance of the assimilation channel is robust across OECD destinations.

6. Conclusion

This paper develops a quantitative dynastic model to explain the immigrant mobility advantage documented by Abramitzky et al. (2021) and to assess its macroeconomic implications. The central mechanism is simple: immigrants are positively selected on latent ability but face institutional frictions that suppress their market earnings. When these frictions decay rapidly across generations (as in the United States), immigrant children leverage their inherited ability advantage and outperform comparable natives. When frictions persist (as in many European countries), the same positively selected immigrants produce children who underperform.

Four quantitative findings emerge. The misallocation induced by immigrant labor market frictions imposes a general-equilibrium output loss of 4.94% of U.S. GDP, roughly $1.4 trillion at 2024 output levels. This loss decomposes into three channels: a static labor-misallocation wedge of 1.57 percentage points (the Hsieh and Klenow (2009)-style measure, concentrated overwhelmingly in the first generation — 87% of the static loss comes from the 5% of the population that constitutes the first-generation cohort), a 0.90 pp suppression of intergenerational human capital investment, and a 2.47 pp downstream adjustment of the aggregate capital stock. Friction decay (assimilation) accounts for 80–87% of the second-generation mobility advantage, with positive selection contributing the remaining 13–20%; the range spans a local fixed-price decomposition (80%) and the full general-equilibrium re-solve (87%). The advantage would flip sign at a friction-persistence threshold of $\lambda^* \approx 0.25$, approximately 0.08 above the calibrated U.S. value of $\lambda_{US} = 0.175$. Applied across destinations, calibrating country-specific λ_j to match each country's observed Gen2 gap produces a distribution of required friction-persistence values that clusters coherently by institutional type: Anglophone destinations (UK, US, Canada) at $\lambda_j < 0.25$, Continental and Scandinavian European destinations in the moderate middle ($\lambda_j \in [0.25, 0.35]$), and a high-friction tail (Netherlands, Austria) at $\lambda_j \approx 0.51$–0.54. The sign-flip threshold $\lambda^* \approx 0.25$ is emergent from the U.S. calibration rather than targeted, and happens to fall within the empirical range of cross-country institutional variation. Once λ_j is bisected to the observed gap, its sign alignment with λ^* is mechanical under common-κ and is not itself an independent out-of-sample prediction; what the exercise does establish is that plausible institutional variation across OECD destinations maps to a narrow and internally coherent region of the model's parameter space.

These findings carry two policy implications. Because assimilation dominates the second-generation outcome, policies that accelerate the decay of institutional frictions — credential recognition, language training, anti-discrimination enforcement, housing integration — generate larger returns than admission policies focused solely on selecting high-ability immigrants. Eliminating the first-generation penalty κ entirely would close the GE output gap in full, whereas fully accelerating intergenerational decay ($\lambda = 0$) would close only 0.58 percentage points of it (Table 6); though small in relative terms, this corresponds to roughly \$160 billion in annual foregone production. The κ lever dominates, but both are first-order. A second, complementary implication: the efficiency case for skill-based admission policy depends on pairing it with friction-reducing measures. Selection and friction reduction are strictly complementary policy levers — raising $\bar{h}_{\text{mig}}$ by 20% alone delivers a 2.7% GDP gain, removing frictions alone delivers 4.95%, while pursuing both yields 8.6%, roughly one percentage point above the sum of the individual effects. Positive selection and friction reduction are complements, not substitutes.

Beyond these findings, the structural framework produces model-implied estimates of intergenerational friction persistence by origin (Table 9) and destination (Table 11). These λ estimates have empirical counterparts — ethnic enclave persistence, credential recognition quality, language-distance gradients, origin-specific third-generation convergence speeds — that are plausibly measurable but currently under-used in the mobility literature. External measurement would provide out-of-sample discipline on the mechanism: corroboration would strengthen the first-order story, while divergence would be diagnostic about which structural features the current framework is missing.

Several extensions to the model itself are within reach. Relaxing the uniform-κ assumption in favor of an ability-gradient $\kappa(h) = \kappa \cdot (h_{\text{ref}}/h)^{\kappa_1}$ would flatten the steep upper-tail second-generation profile toward the ABJP linear specification, at the cost of requiring an additional conditional-variance moment for identification. Endogenizing λ — by allowing parents to invest in reducing their children's friction through neighborhood, language, and network choices — would speak to recent evidence that immigrant families' location choices contribute to second-generation outcomes (Abramitzky et al., 2021). A Roy-model selection block would endogenize the migration decision and allow analysis of how skill-based versus family-based admission interacts with assimilation frictions, speaking directly to the literature initiated by Borjas (1993). Allowing origin-specific latent human capital distributions $\bar{h}_{\text{mig},g}$ would improve the model's ability to capture groups like Vietnam (Section 5.4) where selection dominates friction. Additional features — endogenous fertility, neighborhood effects, marriage markets, fiscal redistribution — could each amplify or attenuate the mechanisms identified here.

These model moves would be strengthened by complementary empirical inputs. Period-specific historical capital-output ratios, native rank-rank slopes, and entry-cohort wage data would enable a proper recalibration to the historical periods studied by Abramitzky et al. (2014), rather than the illustrative fixed-parameter exercise reported in Section 5.3. Longitudinal first-generation wage data linked to second-generation outcomes would discipline a within-lifetime assimilation extension, mapping the years-since-arrival convergence pattern documented by Lubotsky (2007) — whose linked longitudinal design addresses the cross-sectional cohort-effect bias emphasized earlier in this literature — into a geometric-decay structure. Larger three-or-more-generation linked administrative data would tighten the identification of λ and reduce sampling uncertainty around the noisier conditional-percentile moments. Finally, harmonized cross-country measures of first-generation wage penalties — paralleling the mobility data of Boustan et al. (2025) — would allow the destination-heterogeneity exercise to use country-specific κ values rather than a common value, sharpening the threshold story.

More broadly, the results suggest a methodological lesson that extends beyond the immigration setting: when institutional distortions affect dynamic accumulation decisions — parental investment in children, household wealth formation — and not just current allocations, static wedge accounting can substantially understate their welfare cost. In the present setting, the dynamic channels account for roughly two thirds of the total loss.

Despite these limitations, the model delivers a clear message: the speed at which immigrant families assimilate into the labor market is not merely a social outcome to be observed but an economic variable with first-order consequences for aggregate productivity. Fast assimilation is an efficiency imperative.

References

Abramitzky, R., L. Boustan, E. Jacome, and S. Perez (2021). Intergenerational mobility of immigrants in the US over two centuries. *American Economic Review 111*(2), 580–608.

Abramitzky, R., L. P. Boustan, and K. Eriksson (2014). A nation of immigrants: Assimilation and economic outcomes in the age of mass migration. *Journal of Political Economy 122*(3), 467–506.

Achard, P. (2024). The intergenerational (im)mobility of immigrants. *Journal of Public Economics 238*, 105156.

Adda, J., C. Dustmann, and J.-S. Goerlach (2022). The dynamics of return migration, human capital accumulation, and wage assimilation. *Review of Economic Studies 89*(6), 2841–2871.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics 109*(3), 659–684.

Algan, Y., C. Dustmann, A. Glitz, and A. Manning (2010). The economic situation of first and second-generation immigrants in France, Germany and the United Kingdom. *Economic Journal 120*(542), F4–F30.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *Quarterly Journal of Economics 132*(4), 1553–1592.

Antecol, H., P. Kuhn, and S. J. Trejo (2006). Assimilation via prices or quantities? Sources of immigrant earnings growth in Australia, Canada, and the United States. *Journal of Human Resources 41*(4), 821–840.

Becker, G. S. and N. Tomes (1986). Human capital and the rise and fall of families. *Journal of Labor Economics 4*(3), S1–S39.

Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2019). Who becomes an inventor in America? The importance of exposure to innovation. *Quarterly Journal of Economics 134*(2), 647–713.

Bohn, H. and A. R. Lopez-Velasco (2018). Intergenerational mobility and the political economy of immigration. *Journal of Economic Dynamics and Control 94*, 72–88.

Borjas, G. J. (1993). The intergenerational mobility of immigrants. *Journal of Labor Economics 11*(1), 113–135.

Boustan, L., M. F. Jensen, R. Abramitzky, E. Jácome, A. Manning, S. Pérez, A. Watley, A. Adermon, J. Arellano-Bover, O. Åslund, et al. (2025). Intergenerational mobility of immigrants in 15 destination countries. Technical report, National Bureau of Economic Research.

Card, D. (2005). Is the new immigration really so bad? *Economic Journal 115*(507), F300–F323.

Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2020). Race and economic opportunity in the United States: An intergenerational perspective. *Quarterly Journal of Economics 135*(2), 711–783.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics 129*(4), 1553–1623.

Chiswick, B. R. (1978). The effect of Americanization on the earnings of foreign-born men. *Journal of Political Economy 86*(5), 897–921.

Cunha, F., J. J. Heckman, and S. M. Schennach (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica 78*(3), 883–931.

Eckstein, Z. and Y. Weiss (2004). On the wage growth of immigrants: Israel, 1990–2000. *Journal of the European Economic Association 2*(4), 665–695.

Ehrlich, I. and J. Kim (2015). Immigration, human capital formation, and endogenous economic growth. *Journal of Human Capital 9*(4), 518–563.

Ehrlich, I. and Y. Pei (2021). Endogenous immigration, human and physical capital formation, and the immigration surplus. *Journal of Human Capital 15*(1), 34–85.

Han, J. and A. S. Hermansen (2024). Wage disparities across immigrant generations: Education, segregation, or unequal pay? *ILR Review 77*(4), 598–625.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and U.S. economic growth. *Econometrica 87*(5), 1439–1474.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics 124*(4), 1403–1448.

Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics 38*(3), 469–494.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review 112*(1), 169–212.

Lee, S. Y. and A. Seshadri (2019). On the intergenerational transmission of economic status. *Journal of Political Economy 127*(2), 855–921.

Lessem, R. and C. Sanders (2020). Immigrant wage growth in the United States: The role of occupational upgrading. *International Economic Review 61*(2), 941–972.

Llull, J. (2018). Immigration, wages, and education: A labour market equilibrium structural model. *Review of Economic Studies 85*(3), 1852–1896.

Lubotsky, D. (2007). Chutes or ladders? A longitudinal analysis of immigrant earnings. *Journal of Political Economy 115*(5), 820–867.

Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review 104*(10), 3186–3221.

Oreopoulos, P., M. E. Page, and A. H. Stevens (2006). The intergenerational effects of compulsory schooling. *Journal of Labor Economics 24*(4), 729–760.

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous plants. *Review of Economic Dynamics 11*(4), 707–720.

Solon, G. (1999). Intergenerational mobility in the labor market. *Handbook of Labor Economics 3*, 1761–1800.

A. Robustness Exercises

This appendix reports three robustness exercises that complement the identification analysis in Section 3.4. Section A.1 presents an Andrews-Gentzkow-Shapiro (2017) target sensitivity exercise, the methodologically standard robustness check for SMM calibrations. Section A.2 presents a complementary single-parameter sensitivity check on σ_w. Section A.3 reports an unconstrained-SMM corroboration of the external ρ anchor used in the main specification.

A.1. Target Sensitivity (Andrews-Gentzkow-Shapiro)

A natural concern with any SMM calibration is that the headline numbers may be sensitive to the precise values of the empirical targets. If moving a target by, say, 5% (well within the statistical uncertainty of the underlying empirical estimates) causes the static wedge or the decomposition to shift substantially, the results would rest on a knife edge. Following Andrews et al. (2017), I assess this by perturbing each target individually and re-running the full calibration.

Setup. For each target $T \in \{K/Y, \text{native RR slope}, \text{Gen2 gap @ P25}, \text{wage ratio}, \text{Gen2 RR slope}\}$, I run two perturbed calibrations: one with $T' = 1.05 \cdot T$ and one with $T' = 0.95 \cdot T$, for ten perturbed calibrations in total. All other targets remain at their baseline values, and the full six-parameter optimization is re-run for each perturbation. Each perturbed calibration starts from the baseline optimum with a tight initial Nelder-Mead simplex.

Results. Table 12 reports the headline numbers across all ten perturbations.

The results show strong uniformity. The static wedge varies between 1.58% and 1.64% across all ten perturbations (baseline: 1.57%), and the assimilation share is exactly 80% in every single case. The second-generation gap ranges from +4.50 to +5.21 pp, clustered within one percentage point of the baseline +4.64 pp.

Discussion. The near-invariance of the headline numbers to ±5% target perturbations has a specific structural interpretation. The calibration optimum sits in a wide, flat basin of the objective function: the SMM loss surface is shallow in the neighborhood of the optimum, so perturbations that shift the loss surface by a small amount do not move the minimizer far. In several cases, the perturbed calibration converges back to essentially the same parameter vector as the baseline, producing headline numbers within simulation noise of the baseline. The small residual differences in Table 12 (e.g. +4.73 pp Gen2 gap vs. baseline +4.64 pp) are seed-count artifacts: the perturbed calibrations use 5-seed evaluations with

Table 12: *AGS Target Sensitivity (±5% perturbations)*

Perturbation	Gen2 gap (pp)	Static wedge (%)	Assimilation (%)
K/Y +5%	+4.73	1.58	80
K/Y −5%	+4.73	1.58	80
Native RR slope +5%	+4.73	1.58	80
Native RR slope −5%	+4.73	1.58	80
Gen2 gap @ P25 +5%	+4.98	1.63	80
Gen2 gap @ P25 −5%	+4.50	1.58	80
Wage ratio +5%	+4.80	1.63	80
Wage ratio −5%	+5.21	1.64	80
Gen2 RR slope +5%	+4.73	1.58	80
Gen2 RR slope −5%	+4.73	1.58	80
Baseline	**+4.64**	**1.57**	**80**

Notes: Each row is a complete five-parameter SMM re-optimization (with ρ held fixed at the Lee-Seshadri external anchor) starting from a tight simplex around the baseline optimum with 5 seeds per evaluation and early stopping after 20 stale evaluations. Only the indicated target is perturbed; other targets remain at baseline values. The ten perturbed calibrations ran in parallel on separate CPU cores. The baseline row reports the unperturbed calibration for comparison.

early stopping to keep the AGS exercise computationally tractable, while the baseline row reports the 20-seed production number. A 20-seed re-evaluation of each perturbed vector would move its reported numbers closer to the baseline. This is not an artifact of insufficient optimization (the Nelder-Mead simplex explored 20–40 evaluations per perturbation before early-stopping) but rather a genuine feature of the model's structure.

The economic reason is that the static wedge and the decomposition are aggregate equilibrium objects that depend on the *joint* configuration of the parameters, not on any single parameter in isolation. When one target shifts by 5%, the optimizer adjusts the parameter vector along the gradient of the loss surface, but because the headline numbers are smooth, slowly-varying functions of the equilibrium, the net effect on the static wedge and the decomposition is negligible. The few perturbations that produce visible shifts in the second-generation gap are the ones most directly tied to it (the second-generation gap target and the wage ratio target), consistent with the dominant-moment identification logic in Section 3.4.

The takeaway for the paper's conclusions is clear: there are no knife-edge results. The 1.57% static wedge, the 80% assimilation share, and the qualitative ordering of the cross-country λ_j distribution are robust features of the model that do not depend on precise knowledge of any individual empirical moment.

A.2. Single-Parameter Sensitivity to σ_w

As a complementary robustness check, I perturb σ_w by ± 5, ± 10, ± 15, and $\pm 20\%$ from its calibrated value, hold all other parameters fixed at their calibrated optima, and recompute the full set of moments and the static wedge at each point. As discussed in Section 3.4, σ_w serves primarily as the identifier for the second-generation rank-rank slope and has comparatively weak influence on the other four moments. It is also the only calibrated parameter that has no direct empirical interpretation in the existing immigrant mobility literature, so an explicit single-parameter robustness check addresses any concern that the calibration is hostage to a hard-to-discipline auxiliary parameter. Table 13 reports the results.

Table 13: *Sensitivity to σ_w ($\pm 20\%$ perturbations, other parameters fixed)*

σ_w	% change	K/Y	Native RR	Gen2 gap	Wage ratio	Gen2 slope	Static wedge
0.266	−20%	2.83	0.307	+4.63	0.949	0.291	1.57%
0.299	−10%	2.83	0.307	+4.70	0.949	0.270	1.57%
0.315	−5%	2.83	0.307	+4.71	0.949	0.261	1.57%
0.332	**base**	**2.83**	**0.307**	**+4.64**	**0.949**	**0.251**	**1.57%**
0.348	+5%	2.83	0.308	+4.55	0.949	0.242	1.57%
0.365	+10%	2.83	0.308	+4.59	0.949	0.234	1.57%
0.398	+20%	2.83	0.308	+4.79	0.949	0.218	1.57%

Notes: Only σ_w is perturbed from its calibrated value; all other parameters remain at their baseline optima. Moments and the static wedge are recomputed at each perturbation without re-optimizing (unlike the AGS exercise in Table 12).

Over the full $\pm 20\%$ range, the static wedge is constant at 1.57%, the capital-output ratio is constant at 2.83, the native rank-rank slope varies by less than 0.3%, and the immigrant-to-native wage ratio is constant at 0.949. The only moment that responds meaningfully is the second-generation rank-rank slope (the moment σ_w is designed to pin), which varies by approximately $\pm 15\%$ across the range. The paper's main conclusions, all of which concern either the output loss or the second-generation gap sign, are therefore invariant to the uncertainty in σ_w.

A.3. Unconstrained-SMM Corroboration of the External ρ Anchor

The main specification pins the ability-persistence parameter ρ externally at Lee and Seshadri's (2019) value of 0.23 (Table 1, Panel B). This delivers a formally just-identified system of five parameters and five moments. As a check that the external anchor is not driving the paper's quantitative results, this subsection reports a complementary calibration in which ρ is instead estimated jointly with the other five parameters by SMM against the same five targets. The resulting variant is formally over-parameterized (six parameters, five moments), but it provides a direct test of whether the data, given the freedom to choose, would select a value of ρ near the external anchor.

Setup. I perform a complete six-parameter Nelder-Mead optimization with ρ included in the simplex, using the same grid (50×75 state space, 80×50 choice space), agent count (400,000), seed count (20 for final evaluation), targets, and weights as the main calibration. Starting point and convergence tolerances are identical. The unconstrained variant therefore differs from the main specification in exactly one respect: whether ρ is externally pinned or SMM-selected.

Results. Table 14 reports the calibrated parameters and moment fit for both specifications.

Table 14: *Main Specification vs. Unconstrained ρ: Parameters and Moment Fit*

	Main (externally-anchored ρ)	Unconstrained (SMM ρ)
Parameters:		
β	0.967	0.963
ρ	0.230 (fixed, Lee and Seshadri 2019)	0.255 (SMM)
λ	0.175	0.184
$\bar{h}_{\text{mig}}$	0.759	0.696
σ_w	0.332	0.353
κ	0.296	0.264
Moments vs. targets (% error):		
K/Y (3.0)	2.83 (−5.5%)	2.71 (−9.5%)
Native RR slope (0.33)	0.308 (−6.8%)	0.332 (+0.7%)
Gen2 gap @ P25 (5.4)	+4.64 (−14.0%)	+4.86 (−10.1%)
Imm/nat wage ratio (0.92)	0.949 (+3.2%)	0.938 (+1.9%)
Gen2 RR slope (0.25)	0.251 (+0.5%)	0.260 (+4.1%)
SMM objective	0.018	0.012

Notes: Both specifications use identical grids, agent counts, seed counts, targets, and weights. The unconstrained variant calibrates all six parameters by SMM; the main specification fixes ρ externally at the Lee-Seshadri value. Objective is the weighted sum of squared relative errors in the five targeted moments.

Discussion. Two observations. First, the SMM optimum in the unconstrained variant selects $\rho = 0.255$, within 11% of the Lee-Seshadri external anchor of 0.23. This proximity is evidence that the external anchor is not a convenient choice but rather reflects where the data themselves point when given the freedom to choose. Second, the unconstrained variant fits the moments slightly better on balance (objective 0.012 vs. 0.018), as expected when the optimizer has an additional degree of freedom; but the differences in fit are small (largest gap: 3 percentage points on the capital-output ratio), and no moment shifts qualitatively.

The paper's headline numbers are quantitatively similar across the two specifications: the static wedge is 1.34% in the unconstrained variant versus 1.57% in the main specification, the second-generation gap at P25 is +4.86 pp versus +4.64 pp, and the selection-assimilation decomposition is 23/77 versus 20/80. All qualitative conclusions — a positive and economically large second-generation mobility advantage, a sign-flip at intermediate λ, and first-generation dominance of the static loss — are robust across both specifications.

The main paper uses the externally-anchored ρ as its lead specification because the formal just-identification is cleaner and because a literature-anchored value is more transparent to the reader than an internally-selected one. This appendix establishes that the choice does not drive the quantitative conclusions.

A.4. A Diagnostic on the Identifying Content of κ

Appendix A.1 probes whether the calibration's headline numbers are sensitive to the precise empirical target values; Appendix A.3 probes whether the external anchor on ρ is driving the results. This subsection closes the loop with an analogous check on κ: I fix the first-generation friction at zero and freely re-optimize the remaining five parameters (β, ρ, λ, $\bar{h}_{\text{mig}}$, σ_w) on the same five moments, using the same SMM objective and moment weights as the main calibration. The purpose is to isolate whether the paper's other structural channels — positive selection on latent ability via $\bar{h}_{\text{mig}}$, Becker-Tomes skill persistence via ρ, and the household investment problem — can jointly reproduce the targeted moments without the wedge. The exercise is a diagnostic on the identifying content of κ *within the framework of this paper*; it is not a reimplementation of any specific prior model, and its interpretation in relation to Borjas (1993) is discussed below.

Table 15 reports the comparison. The baseline (main specification with ρ anchored externally) attains an SMM objective of 0.018; the $\kappa = 0$ variant settles at 0.098, roughly 5.4 times worse. The failure concentrates on a single moment: the immigrant-to-native median wage ratio crosses above unity, delivering 1.054 against a target of 0.92 — a reversal in direction, not merely in magnitude. With no wedge separating latent ability from observed earnings, the only route to a positive Gen2 advantage within this parameter class runs through positive selection on latent ability, and in the absence of a friction wedge that selection translates directly into higher first-generation observed wages. The optimizer attempts to compensate by raising ρ to 0.287 and lowering $\bar{h}_{\text{mig}}$ to 0.588, but the wage-ratio miss cannot be closed: the remaining parameters cannot simultaneously deliver a below-unity first-generation wage ratio and a positive second-generation advantage. (Note that λ has no effect on the model when $\kappa = 0$, since $\theta_g = \kappa\lambda^{g-1} \equiv 0$ for all g; its reported calibrated value of 0.224 in Table 15 is an optimization artifact over a flat objective surface, not a meaningful structural choice.)

Table 15: *Baseline vs. $\kappa = 0$: Moment Fit and Parameters*

	Baseline	$\kappa = 0$	Target
Moment fit:			
Capital-output ratio (K/Y)	2.83	2.98	3.00
Native rank-rank slope	0.307	0.368	0.33
Gen2 gap at P25 (pp)	+4.64	+4.71	+5.4
Imm/nat median wage ratio	0.949	1.054	0.92
Gen2 rank-rank slope	0.251	0.271	0.25
SMM objective	*0.018*	*0.098*	—
Calibrated parameters:			
κ	0.296	0 (fixed)	—
ρ	0.23 (anchored)	0.287	—
λ	0.175	0.224	—
$\bar{h}_{\text{mig}}$	0.759	0.588	—
β	0.967	0.973	—
σ_w	0.332	0.310	—

Notes: The baseline fixes ρ at the Lee-Seshadri (2019) value; the $\kappa = 0$ variant fixes κ at zero and re-optimizes ρ alongside the other four parameters. Same SMM objective and moment weights as the main calibration; the $\kappa = 0$ optimization uses 5 seeds per evaluation with a 10-seed evaluation at the optimum and 200,000 agents per seed (vs. 20 seeds and 400,000 agents in the main calibration), identical random seeds across specifications.

Relationship to Borjas (1993). This exercise is not an implementation of Borjas's own preferred specification, which is richer than a pure no-wedge model. His framework contains the closest conceptual analogues of ρ and $\bar{h}_{\text{mig}}$ — a country-specific Markov skill-transmission parameter δ_y and Roy-style selection on source-country skills — and his equations (6)–(7) explicitly allow mean incomes μ_{yt} to shift across generations to accommodate "the assimilation process [that] could affect future earnings opportunities," though this shift is not formalized as a structural parameter distinct from skill persistence. Empirically, his Table 4 Row 1 recovers a 7.0% common intergenerational intercept from cross-ethnic-group regressions, which he attributes expositionally to schooling quality, English proficiency, and reduced ties to ethnic enclaves — the object that the structural $\theta_g = \kappa\lambda^{g-1}$ decaying wedge in this paper is designed to formalize. The $\kappa = 0$ exercise in this appendix should therefore be read as a diagnostic on whether the structural channels other than the wedge, within this paper's framework, can jointly recover the targeted moments — not as a test of any of Borjas's own preferred specifications. The answer is that they cannot, and the failure concentrates on the wage-ratio moment: this is the object that the wedge is load-bearing for within this framework.

This appendix is a diagnostic rather than a formal specification test: the $\kappa = 0$ variant has free parameters of its own and is not nested in a standard statistical testing framework. But the wage-ratio crossing unity does establish that κ adds identifying content beyond what is available from the remaining parameters — the wedge is not nested in the other channels, and positive selection combined with skill persistence alone cannot replicate the data within this framework.

A.5. Local Identification Jacobian

As a direct check on the identification logic of Section 3.4, I compute the local Jacobian of the model's five SMM moments with respect to the five free parameters, evaluated at the v6.2 calibration. For each parameter $p \in \{\beta, \lambda, \bar{h}_{\text{mig}}, \sigma_w, \kappa\}$, I perturb p by $\pm 2\%$ of its calibrated value (with a minimum absolute step of 0.005), re-solve the model to steady state, and compute the finite-difference derivative of each moment with respect to each parameter. Table 16 reports the resulting elasticity matrix $d \log m / d \log p$.

Table 16: *Local Identification Jacobian: Moment Elasticities w.r.t. Parameters*

Moment	β	λ	$\bar{h}_{\text{mig}}$	σ_w	κ
K/Y	**+8.14**	−0.01	+0.07	+0.00	−0.00
Native RR slope	+0.62	+0.01	+0.01	+0.01	−0.14
Gen2 gap at P25 (pp)	−13.67	**−1.12**	+5.91	−0.36	**−0.97**
Imm/Nat wage ratio (P50)	−3.18	+0.00	**+0.99**	+0.01	−0.42
Gen2 RR slope	−1.17	−0.02	−0.91	−0.72	+0.88

Notes: Elasticity of each moment (rows) with respect to each parameter (columns), $d \log m / d \log p$, evaluated at the v6.2 calibration by two-sided finite difference with step $h = \max(0.005, 0.02|p|)$. Bold entries are the parameter's dominant moment (largest |elasticity| in that parameter's column), with κ and λ both loading primarily on the Gen2 gap. Each perturbation is a full GE re-solve (400,000 agents, 5-seed moment averaging). Replication: `model/jacobian_v6_2.py`.

The dominant-moment pattern supports the just-identification argument in Section 3.4.

Each parameter has a distinct identifying moment: β pins K/Y (elasticity $+8.14$, two orders of magnitude above the next-largest column entry for this moment); $\bar{h}_{\text{mig}}$ separately moves the immigrant-native wage ratio (elasticity $+0.99$, the only parameter with non-trivial loading on this moment); σ_w pins the Gen2 rank-rank slope (-0.72); and λ and κ share the Gen2 gap at P25 as a common primary moment, but are separately identified through their loadings on other moments — κ additionally moves the wage ratio (-0.42) and the Gen2 rank-rank slope ($+0.88$), while λ's elasticities on those moments are effectively zero. The native rank-rank slope has its largest loading on β ($+0.62$), which is already pinned by K/Y; in the main specification ρ is the parameter disciplined by this moment, fixed externally at the Lee-Seshadri (2019) value (see Appendix A.3 for the unconstrained variant, in which ρ is estimated jointly with the other five parameters).

Two properties of the Jacobian deserve comment. First, β's elasticities on the rank-based moments are large (Gen2 gap -13.67; wage ratio -3.18; Gen2 RR slope -1.17). This reflects the equilibrium propagation of β through wealth, investment, and the human-capital distribution: moving β shifts aggregate savings, which shifts the capital-labor ratio and the wage level, which in turn shifts parental investment decisions and children's outcomes. The magnitude of these cross-elasticities is not a failure of identification — β is tightly pinned by K/Y (on which no other parameter loads non-trivially), so once β is calibrated, the elasticity of other moments with respect to β is a property of the equilibrium, not a degree of freedom. Second, the singular-value decomposition of the Jacobian yields singular values $\{100.1, 14.3, 1.5, 0.5, 0.07\}$, a condition number of 1495, and full effective rank 5/5. The system is therefore locally identified. The smallest singular value (0.07) corresponds to a right singular vector that places nontrivial weight on λ, $\bar{h}_{\text{mig}}$, σ_w, and κ simultaneously — a weak direction in parameter space that would be hard to pin down with data substantially noisier than ours. In practice, this weak direction is the empirical object that the target-sensitivity exercise in Appendix A.1 probes: moving each target by $\pm 5\%$ shifts the calibrated parameters along this weak direction but leaves the headline moments and decomposition outcomes quantitatively unchanged, because the weak direction is approximately orthogonal to the model outputs we care about.

B. First-Generation Wage Gap Estimates

This appendix reports the Mincerian wage regressions used to calibrate the first-generation wage penalty κ in the main text (Section 3.1) and the origin-specific penalties κ_g in the heterogeneity exercise (Section 5.1).

B.1. Aggregate Immigrant Wage Gap

Table 17 reports the results of estimating

$$\ln(\text{hourly wage})_i = a + b \cdot \text{ForeignBorn}_i + f(\text{age}_i) + \psi(\text{education}_i) + u_i$$

by weighted least squares on IPUMS American Community Survey data (2015–2019), using IPUMS person weights. The sample consists of prime-age workers (25–55) with positive wages and usual hours, after trimming hourly wages below $3 or above $500. The age controls are a quadratic in age; education controls (Specification 1 only) are dummies for high school, some college, bachelor's, and graduate degree, with less-than-high-school as the omitted category. The foreign-born indicator equals one for individuals with birthplace outside the 50 US states (IPUMS NATIVITY = 5).

Table 17: *Aggregate First-Generation Wage Gap: Mincerian Regressions*

	(1) Conditional	(2) Unconditional
Foreign-born	−0.0437	−0.0979
	(0.0008)	(0.0008)
Implied wage gap (%)	−4.3	−9.3
Age controls	Yes	Yes
Education controls	Yes	No
Observations	4,690,977	
of which foreign-born	821,441	

Notes: Weighted least squares using IPUMS person weights. Sample: ACS 2015–2019, workers aged 25–55 with hourly wage $3–$500. Dependent variable: log hourly wage. Specification (1) includes age quadratic and education dummies. Specification (2) includes age quadratic only. Standard errors in parentheses. Implied wage gap computed as $(\exp(\hat{b}) - 1) \times 100$. The regression coefficients a and b are local to this appendix and should not be confused with the capital share (α) and discount factor (β) of Section 2.

The conditional specification (column 1) yields a foreign-born coefficient of −0.044, corresponding to a 4.3% residual wage gap after controlling for education — the same-credential wage penalty facing immigrants, capturing credential non-recognition, discrimination, and other sources of within-education-level wage differences. The unconditional specification (column 2) yields −0.098, or a 9.3% gap. The difference of approximately 5 percentage points is the composition effect: immigrants have different average educational attainment than natives (in this sample, on net lower), so omitting the education control attributes to the immigrant dummy the earnings consequences of that composition difference. The credential non-recognition channel itself is embedded in the conditional residual, not in the difference between specifications.

Both cross-sectional estimates are informative about the order of magnitude of the friction but are not used directly as the calibration target; the baseline calibration treats κ as a freely estimated parameter jointly identified with $\bar{h}_{\text{mig}}$ by the immigrant-to-native median wage ratio and the second-generation mobility gap. The calibrated value $\kappa = 0.296$ lies near the upper end of the entry-cohort penalty range (15–30%) documented by Borjas (1993) and Lubotsky (2007), and well above the cross-sectional Mincerian gap (4–9%). As discussed in Section 4.4, the model's one-period-per-generation structure makes κ a parsimonious reduced form for within-cohort dynamics; it is identified here by the Gen2 mobility gap, and the comparison to empirical wage-gap ranges is a plausibility check on magnitude rather than an identification argument.

B.2. Origin-Specific Wage Gaps

Table 18 reports origin-specific wage gaps from the same IPUMS data. For each origin group g, I regress log hourly wages on an indicator for being born in country g, with the comparison group restricted to US-born natives. Both conditional (with education) and unconditional (age only) specifications are reported.

The conditional specification is the appropriate mapping to the model's κ_g, which represents the wedge between latent ability and realized earnings for a given level of human capital. Using the conditional gap ensures that κ_g captures institutional friction (credential non-recognition, language barriers, employer discrimination) rather than differences in the quantity of human capital across origin groups.

The gap between the conditional and unconditional columns is informative about the role of education. For Mexico, the unconditional gap (-33.6%) is more than double the conditional gap (-13.1%), indicating that the majority of the raw wage gap is accounted for by differences in educational attainment. For Cuba, the gap narrows much less (-23.0% conditional vs. -29.8% unconditional), suggesting that Cuban immigrants have educational attainment closer to natives but face larger residual frictions, consistent with the refugee-cohort composition of Cuban immigration.

Table 18: *Origin-Specific First-Generation Wage Gaps*

		Conditional		Unconditional		
Origin	N	$\hat{b}_g$	Gap (%)	$\hat{b}_g$	Gap (%)	κ_g
Mexico	198,507	−0.1404	−13.1	−0.4098	−33.6	0.14
		(0.0014)		(0.0014)		
Dominican Rep.	9,924	−0.1759	−16.1	−0.2846	−24.8	0.18
		(0.0052)		(0.0056)		
Cuba	17,980	−0.2612	−23.0	−0.3540	−29.8	0.26
		(0.0042)		(0.0045)		
Vietnam	26,932	−0.0616	−6.0	−0.1203	−11.3	0.06
		(0.0037)		(0.0041)		

Notes: Weighted least squares using IPUMS person weights. Sample: ACS 2015–2019, prime-age workers (25–55) with hourly wage $3–$500. N reports the origin-group sample size; the comparison group in each regression is 3,843,595 US-born natives (birthplace codes 100–5600). Conditional specification includes age quadratic and education dummies; unconditional includes age quadratic only. Standard errors in parentheses. κ_g is the calibration input used in Section 5, set equal to $|\hat{b}_g|$ from the conditional specification, which isolates institutional friction net of human capital composition. Sample sizes differ from Table 17 because the native comparison group is restricted to IPUMS birthplace codes 100–5600 (states and federal districts), which drops approximately 0.7% of the aggregate native sample.