ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 027/26

Migration for Marriage

Shadi Farahzadi

www.rfberlin.com JANUARY 2026

Migration for Marriage

Authors

Shadi Farahzadi

Reference

JEL Codes: J12, F22, J15

Keywords: Migration, Integration, Marriage Market, Matching Model

Recommended Citation: Shadi Farahzadi (2026): Migration for Marriage. RFBerlin Discussion Paper No. 027/26

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Migration for Marriage*

Shadi Farahzadi†

January 19, 2026

Abstract

The impact of policies restricting marriage migration depends on whether it is driven by preferences to marry within one's group (endogamy preferences) or gains from residency in a developed country (outside market value). I develop a novel marriage matching model incorporating the choice to marry from one's origin country. I focus on British Muslims, since half of them marry someone from their origin country. I find that 80% of Muslim marriage migration is explained by endogamy preferences. Therefore, raising the costs of marriage migration does not increase their integration through intermarriage; instead, it leads to a larger share of unmarried Muslims.

*I am grateful to Alan Manning, Guy Michaels, Yona Rubinstein, and Catherine Thomas for their priceless guidance and encouragement throughout the project. For very helpful comments, I would like to thank Dan Anderberg, Nava Ashraf, Oriana Bandiera, Philipp Barteska, Patrick Bennet, Jordi Blanes-i-Vidal, Alix Bonargent, Paul Brandily, Jacopo Bregolin, Gharad Bryan, Ian Burn, Matthias Doepke, Nicholas de Roos, Magdalena Domínguez, Jeremy Fox, Alfred Galichon, Alessandro Gavazza, Maitreesh Ghatak, Michael Gmeiner, Cristina Gualdani, Erik Hurst, Ethan Ilzetzki, Osnat Lifshitz, Rocco Macchiavello, Stephen Machin, Kristóf Madarász, Noemi Mantovan, William Matcham, Robert Miller, Sidharth Moktan, Isabela Manelici, Sonia Oreffice, Martin Pesendorfer, Steve Pischke, Imran Rasul, Ricardo Reis, Rebecca Rose, Marcia Schafgans, Pasquale Schiraldi, Daniel Sturm, Christiane Szerman, Veronica Rappoport, John Van Reenen, Jose Vasquez, Leeat Yariv, Alwyn Young and numerous seminar and conference participants, especially at NYU Abu Dhabi, University of Luxembourg, KU Leuven, SOLE 2024, SEHO 2024, Scottish Economic Society 2024, Matching Markets and Inequality Workshop, VI Workshop on Immigration, IÉSEG Workshop on Migration and Family Economics, and City/Royal Holloway Workshop on Family and Health Economics. This work contains statistical data from ONS which is Crown Copyright. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

†University of Liverpool, Centre for Economic Performance, and Rockwool Foundation Berlin. Email: shadi.farahzadi@liverpool.ac.uk

1 Introduction

Spouse visa restrictions have been implemented in many high-income countries with two main objectives: reducing net migration, particularly of low-skilled migrants; and promoting the domestic integration of minority communities. Assessing whether such policies can achieve these objectives requires knowledge of why a country's minority groups choose to import their spouses. I infer motives for marriage migration from data on who marries whom in the United Kingdom using a structural matching model. I find that while the visa restrictions do reduce migration, they neither target the intended low-skilled migrants nor foster integration through intermarriage.

Marriage migration has become a significant component of inflows to many OECD countries (OECD, 2017). In this paper, I define marriage migration as the situation in which a resident or citizen sponsors a foreign spouse who migrates for the purpose of marriage. This is distinct from family reunification, which refers to situations where a migrant is married before leaving the origin country and later sponsors their spouse to join them.[1] In the UK, marriage migration accounts for about 30% of total immigration (Labour Force Survey, 2011).[2] This trend is primarily driven by ethnic minorities seeking spouses from their ancestral countries of origin (Charsley et al., 2020). Policymakers have raised concerns that this pattern may hinder integration, creating "A first generation in every generation" (Goodhart, 2013; Casey, 2016; Charsley et al., 2017). These concerns are reinforced by low intermarriage rates among groups with high marriage migration[3], given that intermarriage is widely viewed as a key indicator of social integration (Gould and Klor, 2016).

The impact of marriage migration policies depends on the motivations driving these unions. If marriage migration is primarily driven by the economic gains from relocating to a higher-income country, then tighter regulations should reduce such migrations and increase intermarriage, consistent with policymakers' objectives. If, instead, the dominant motive is a strong preference for ethnically or religiously similar spouses (endogamy preferences), often more easily found in the country of origin, then restrictive policies will fail to significantly increase intermarriage rates and will instead lead to a higher share of singlehood. These contrasting mechanisms imply different

[1]The analysis focuses on individuals who were born in the host country or who arrived there as children, ensuring that the marriage decision is made after settlement rather than prior to migration.

[2]The long-term stay rate for marriage migrants is 89%, substantially higher than the 18% for migrant students and 57% for migrant workers, contributing significantly more to the overall immigrant population (Hall et al., 2023).

[3]See Figure A1.

welfare implications.

Understanding which mechanism dominates requires a framework capable of separating preferences from market characteristics. A structural matching model is well suited to this task for three reasons. First, it treats marriage as the equilibrium outcome of a two-sided selection process, allowing me to recover the surplus generated by each match net of composition effects. Second, in the context of marriage migration, it enables a clear distinction between endogamy preferences and the gains available in foreign marriage markets. Third, it provides the basis for counterfactual analysis, making it possible to predict how marriage outcomes respond to changes in migration policy.

The majority of existing work on marriage choices focuses on closed marriage markets (Becker, 1973; Choo and Siow, 2006; Hitsch et al., 2010; Chiappori, 2017; Chiappori et al., 2018; Beauchamp et al., 2021; Galichon and Salanié, 2022, 2024; Adda et al., 2025). These models implicitly assume prohibitively high costs for cross-market marriages. This assumption leads to two significant limitations: first, existing models cannot explain and predict marriage migration; and second, these models tend to overestimate preferences for endogamy by neglecting the possibility that some intra-group marriages occur due to gains from migration rather than purely from preferences for within-group marriage.

In this paper, I study the marriage migration of ethnic minorities in high-income countries for the first time in the economics literature. I develop an identification strategy to separately identify determinants of marriage migration. I then estimate these factors, focusing on British Muslims. I define British Muslims as individuals who were born in the UK or who migrated before age 18, ensuring that the marriage decision reflects choices made after settlement. This group has a significantly higher rate of marriage migration compared to other groups; about half of British Muslims marry partners born in their ancestral country of origin[4] compared with rates of 20 percent or less for other major religious groups. Understanding this pattern is essential, as Muslims constitute a large and growing population in Western countries and face documented challenges related to social and economic integration (Adida et al., 2016).

I begin by formulating a general open-economy two-sided transferable utility matching framework in which individuals can form matches either within their own

[4]The ancestral country of origin is identified based on self-reported ethnicity. Individuals with mixed ethnic backgrounds are excluded, as more detailed parental information is not available in the Census.

market or across multiple external markets. In this general setting, equilibrium matching depends on the joint distribution of observable characteristics across all connected markets, and identification requires observing population counts and matches in each market. This framework encompasses a broad class of cross-market matching environments, including marriage migration as well as settings with mobility across regions or countries.

In empirical applications of marriage migration, however, the econometrician typically observes detailed information only on residents of the destination country and on the characteristics of the spouses they marry, while the underlying population distributions in origin markets remain unobserved. I therefore specialize the general framework to a data-restricted case in which only one market is observed. In this setting, individuals choose among three mutually exclusive options: marrying locally, marrying from an external market, or remaining single. Marriage migration is driven by two distinct forces. The first is endogamy preferences, reflecting the desire to marry within one's ethnic or religious group. The second is the outside market value, which captures both access to a larger external pool of potential partners and migration related gains associated with relocating to a higher income country. This structure allows the determinants of marriage migration to be identified despite the absence of data on external markets.

The model treats match formation as the outcome of a random utility framework in which each potential match yields a systematic surplus component and an idiosyncratic unobservable term. I assume that unobserved heterogeneity does not interact with observable traits in generating surplus. Following Choo and Siow (2006), I impose that these unobserved components are independently and identically distributed type I extreme value, which yields a tractable two-sided logit representation of equilibrium matching and implies the Independence of Irrelevant Alternatives (IIA) property. To identify endogamy preferences separately from outside market value, I additionally assume that the migration component of the surplus is separable from the components associated with other observable attributes. These assumptions allow the identification of two distinct determinants of marriage migration. Endogamy preferences are determined from the ratio of intra-group to inter-group local marriages, and outside market values are identified from the ratio of cross-market to local marriages. The model section discusses the implications of these identifying assumptions in detail. In the final section of the paper, I relax the IIA assumption by extending the framework to a nested logit specification that allows for correlation

across alternatives and more flexible substitution patterns in marriage choices. The main results remain robust to this generalization.

To estimate the marriage surplus, I extend the minimum distance estimator developed by Galichon and Salanié (2024), which applies optimal transport methods to recover preference parameters from observed matches and population counts. In addition, I compute the equilibrium matching function using an Iterative Proportional Fitting Procedure (IPFP) adapted from Galichon and Salanié (2022). The model is estimated using microdata from the 2011 Census for England and Wales. The UK's unique practice of collecting information on both ethnicity and religion allows me to identify ethnic and religious preferences separately, an essential feature in analysing the Muslim marriage market where these dimensions are strongly correlated.

The reduced-form patterns indicate that, on average, ethnic minorities at the lower end of the UK education distribution marry spouses who are at the upper end of the education distribution in their countries of origin. This suggests that low-educated UK residents may trade off the value of UK citizenship against the higher educational attainment of partners abroad, contradicting the common policy assumption that marriage migrants are predominantly low-educated (Charsley, 2013).

The structural estimation results align with the familiar patterns of assortative matching on education and age (Siow, 2015). Beyond these findings, results show that couples derive substantially higher surplus from ethnic and religious endogamy compared to age or educational endogamy. Religious minorities in the UK; specifically Muslims, Hindus, and Sikhs; show strong and comparable preferences for marrying within their faith. This stands in sharp contrast to majority groups; Christians and individuals with no religious affiliation; whose preference for religious endogamy is significantly weaker.

Building on these results, the decomposition of marriage migration shows that approximately 80 percent of Muslim marriage migration is attributable to preferences for ethnic and religious endogamy. In a counterfactual scenario without these preferences, the marriage migration rate among Muslims would fall from the observed 50 percent to roughly 10 percent, closely matching the rate among non-Muslim groups. The key distinction between Muslims and other religious minorities such as Hindus and Sikhs therefore lies not in the strength of endogamy preferences, which are similarly strong across these groups, but in outside market value. When strong endogamy preferences interact with a larger outside market, they generate substantial differences in marriage outcomes. A plausible explanation is the relative strength of transnational

networks, which give Muslims access to a larger pool of potential spouses abroad.

My structural estimates allow counterfactual simulations of policy changes. As in many other high-income countries, successive reforms in the UK have raised the cost of marriage migration through policies such as higher income thresholds, language tests, and visa fees. I capture these reforms by introducing a marriage migration tax that reduces the migration component of utility, holding preferences fixed. Higher costs reduce marriage migration among both Muslims and non-Muslims but generate only minimal increases in interreligious marriage. Even a complete ban on marriage migration would raise the interreligious marriage rate among Muslims by less than two percentage points. Instead, individuals adjust along other dimensions. Inter-ethnic marriage rises, while inter-education marriage falls as low-educated Muslims lose access to the relatively inexpensive option of marrying high-educated partners from their ancestral countries. For many individuals, the utility cost of compromising on preferred traits is large, leading to a substantial increase in singlehood. These results imply that restrictive policies reduce migration but are unlikely to increase integration through the marriage market.

The counterfactuals above capture short-run responses to policy changes, holding population composition fixed. Marriage and migration decisions, however, also shape future marriage markets through differential fertility and intergenerational transmission of religion. To capture these feedback effects, I extend the model to a two-generation framework in which current marriage outcomes influence the size and composition of the next generation. A larger co-religionist population reduces reliance on marriage migration by making it easier to find suitable partners locally. Conversely, restrictive marriage migration policies that reduce group size in the short run may increase future incentives to search abroad by making the local marriage market thinner. As a result, policies aimed at limiting marriage migration have offsetting long-run effects that are not captured by static analyses.

The second set of counterfactuals examines how changes in the size of the Muslim population affect marriage market equilibrium. This question is particularly relevant given the higher population growth rates of Muslims in many OECD countries. A larger Muslim population substantially reduces marriage migration by increasing the local availability of ethnically and religiously similar partners. For example, doubling the Muslim population in the UK reduces marriage migration by about 20 percent. Despite this decline, intermarriage rates change little because strong endogamy preferences lead the larger local pool to substitute primarily for spouses who would oth-

erwise be found abroad. I then allow endogamy preferences to depend on group size, consistent with models of cultural transmission (Bisin et al., 2004, 2008). Introducing endogenous preferences amplifies the long-run effects. As the Muslim population grows, both increased local availability and weakened incentives for strict endogamy further reduce marriage migration and generate a modest rise in intermarriage.

My study makes several contributions to the literature. First, I introduce an open economy matching model that extends existing closed-market matching models focused on within-country marriage markets (Becker, 1974; Choo and Siow, 2006; Chiappori, 2017; Galichon and Salanié, 2022; Adda et al., 2025). This open economy matching model can be extended to other domains such as worker-firm matching and student migration across regions or countries. Second, unlike prior studies on marriage migration that focus on brokered cross-market marriages in Asian countries due to unbalanced sex ratios (Kawaguchi and Lee, 2017; Weiss et al., 2018; Ahn et al., 2020), this paper investigates intra-ethnic marriage migration in Western countries driven by preferences for similar cultural backgrounds. Third, it advances the marriage market literature by highlighting the significantly greater importance of religion and ethnicity in marriage decisions, compared to traditionally emphasized characteristics like education and income (Hitsch et al., 2010; Banerjee et al., 2013; Eika et al., 2019; Chiappori et al., 2022; Anderberg and Vickery, 2021). Finally, I contribute to the literature on integration by providing new insights into the impact of religious preferences on the integration of Muslims, a relatively understudied group (Manning and Roy, 2010; Georgiadis and Manning, 2011; Bisin et al., 2004, 2008; Gould and Klor, 2016; Jacquet and Montpetit, 2022).

The paper is organized as follows. Section 2 provides global context on marriage migration and the policies used by governments to regulate it. Section 3 describes the data and presents reduced-form evidence. Section 4 develops the structural model, and Section 5 outlines the estimation strategy. Section 6 reports the main estimation results, with robustness analyses in Section 7. Section 8 presents the counterfactual analyses. Section 9 extends the model by relaxing the IIA assumption and incorporating intergenerational dynamics. Section 10 concludes.

2 Global Context

2.1 Marriage Migration Patterns

Marriage migration (also referred to as transnational or cross-border marriage) refers to international migration undertaken for the purpose of family formation, most commonly involving a citizen or long-term resident sponsoring the immigration of a foreign spouse (Charsley et al., 2012). The phenomenon has long historical roots. For example, in the early twentieth century Japanese men working in the United States arranged marriages with women from their homeland, who migrated as so-called "picture brides" (Glenn and Glenn, 2010).

Although marriage migration is a global phenomenon, its patterns and underlying drivers vary substantially across contexts. In East Asia, countries such as South Korea, Taiwan, and, to a lesser extent, Japan experienced a surge in international marriages beginning in the 1990s. These unions predominantly involved native men, often from rural or lower-status backgrounds, marrying women from economically less-developed neighboring countries such as Vietnam, the Philippines, China, and Cambodia. This pattern emerged largely in response to demographic imbalances, particularly skewed sex ratios and female rural-to-urban migration, which reduced the availability of local marriage partners for men (Constable, 2010; Yang and Lu, 2010). Many of these marriages were facilitated by brokers, raising concerns about commodification, limited spousal agency, and subsequent integration challenges related to language acquisition, cultural adaptation, and socioeconomic disparities (Wang and Chang, 2002; Lindquist et al., 2012; Ahn and Lubotsky, 2024).

By contrast, marriage migration in Western countries often involves both men and women migrating to spouses with shared ethnic, religious, and linguistic backgrounds. A substantial share of children and grandchildren of earlier labor migrants in Europe continue to marry partners from their parents' or grandparents' country of origin, rather than marrying members of the ethnic majority or co-ethnics raised in the same destination country (Beck-Gernsheim, 2007; Charsley, 2013). In the UK, for example, many second-generation British Pakistanis and Bangladeshis marry spouses from their ancestral countries, sustaining transnational marriage ties through repeated family-based migration (Shaw, 2001; Charsley, 2013). Similar patterns are observed in Germany, Belgium, and the Netherlands, where individuals of Turkish and Moroccan descent frequently marry partners from Turkey and Morocco, respectively (Lievens, 1999; Huschek et al., 2012).

These patterns are not solely the result of arranged marriage practices. Drawing on interviews with sibling pairs, one of whom married a migrant spouse and the other a UK-raised partner, Charsley et al. (2020) show that marriage migration is rarely a purely parent-driven arrangement. Participants consistently emphasized their own agency in partner selection. Many highlighted the perceived cultural compatibility of a migrant spouse, particularly when shared dialects and customs were expected to ease household integration. Kinship networks also play an important role. Extended family ties create information channels, reduce search and transaction costs, and often facilitate chain migration, including through consanguineous unions among British Pakistanis and Bangladeshis and similar patterns in Turkish and Moroccan diasporas in Europe. These partner choices are embedded in transnational kin fields in which earlier family migrations shape and enable subsequent cross-market marriages (Shaw, 2001; Charsley, 2013; Huschek et al., 2012).

While these forms of marriage migration may reduce intra-household cultural conflict, they pose challenges for integration into the broader society. Migrant spouses often arrive with limited host-country language proficiency, unrecognized educational credentials, and little familiarity with local labor markets. These barriers constrain both their economic participation and long-term social mobility (Charsley et al., 2012). This is particularly evident in the case of male marriage migrants to the UK from Pakistan. Although they are sometimes viewed as the primary beneficiaries of transnational marriage, many experience downward occupational mobility upon arrival. Despite holding tertiary qualifications or respected professional roles in their origin countries, they frequently encounter significant obstacles to labor market entry, including language barriers, non-transferable qualifications, and structural discrimination. Consequently, many are employed in low-paid and precarious jobs, which may negatively affect their mental health and overall well-being (Charsley, 2005; Charsley and Liversage, 2015). These patterns have generated sustained policy debate over the broader implications of transnational marriage for integration, labor market outcomes, and social cohesion.

2.2 Marriage Migration Policies

Although the right to marry and to found a family is recognized as a fundamental human right under international law, whether this right extends to cross-border mobility remains contested in legal and policy debates (Kofman, 2004). Policy responses to marriage migration vary substantially across regions, reflecting different

priorities regarding integration, demographic change, and immigration control. In East Asia, the prevalence of rapidly arranged marriages involving significant cultural and socioeconomic asymmetries has raised concerns over the integration of migrant spouses. Female marriage migrants, in particular, often face obstacles related to language acquisition, insecure legal status, and childrearing in unfamiliar institutional and cultural environments. In response, governments, most prominently South Korea, have implemented targeted policy interventions. These include income and language requirements for sponsoring spouses, premarital counseling and cultural orientation programs, and tighter regulation of marriage brokerage agencies. Such measures aim to professionalize the transnational matchmaking industry while strengthening the legal and social protection of migrant spouses (Lee, 2008).

In Western countries, policy debates on marriage migration differ from those in East Asia because concerns about integration are closely tied to broader immigration control. Since most marriage migration in Europe occurs within established ethnic and cultural communities, the central issue for policymakers is not the cultural compatibility of the couple but the integration of the migrant spouse into the host society. Governments in Europe and North America therefore treat marriage migration as part of wider immigration regulation, combining measures intended to support integration with restrictions designed to limit inflows (Bonjour and Kraler, 2015). In the UK, these measures include minimum income thresholds, pre-entry language tests, age requirements, higher application fees, and longer processing times (Sumption and Vargas-Silva, 2019). Denmark imposes a minimum age of 24 for both partners and applies integration criteria related to education, employment, and language proficiency, and applicants must also commit to active participation in Danish society (Bonjour and De Hart, 2013). France introduced an integration contract in 2002, making long-term residence permits conditional on language ability and civic engagement, with considerable discretion granted to local authorities (Zappi, 2003). Germany has required pre-entry language acquisition since 2007. Outside Europe, the United States applies a quota-based system that caps the number of family-preference visas for spouses and children of green card holders at 114,200 per year.[5]

Concerns about "sham marriages," defined as unions entered into primarily to circumvent immigration controls, have further shaped the policy landscape. Governments have responded with a range of enforcement measures, including interviews to assess the authenticity of relationships, mandatory cohabitation periods before grant-

[5]Defined as spouses, minor children, and unmarried adult children of green card holders.

ing independent residency, and probationary periods that delay access to permanent status for sponsored spouses (Wray et al., 2014).

Although spouse visa policies are often presented in public discourse as integration-enhancing measures, their practical function frequently lies in restricting family migration. As Wray et al. (2014) argues, these policies act as gatekeeping tools that determine who is able to reunite through marriage by increasing the cost and complexity of sponsorship. Despite extensive regulatory activity and strong normative claims surrounding these measures, rigorous empirical evidence on the underlying motives for marriage migration and on the causal effects of policy constraints remains scarce, a gap this paper seeks to address.

3 Data and Reduced-Form Evidence

3.1 Data Description

This study uses data from the 10% household-level sample of the 2011 Census for England and Wales (Office for National Statistics, 2011). A key advantage of this dataset is that it includes detailed questions on both ethnicity and religion and provides a large number of observations from diverse ethnic and religious groups. The sample is restricted to individuals of marriageable age: women aged 23 to 53 and men aged 25 to 55. These bounds differ by gender because women in the UK marry at younger ages than men and have systematically different age-at-marriage distributions. Using gender-specific age ranges ensures that the age categories capture comparable positions in the marriage market for men and women, avoids heavily unbalanced age cells, and improves the precision of the estimates. The lower bounds also exclude individuals who are likely still in full-time education.

To focus on individuals who made their marital decisions in the UK, I restrict the sample to those who were born in the country or who arrived before age 18. Throughout the paper, I refer to this group as British. This definition ensures that the categories British and marriage migrant are mutually exclusive and exhaustive, increases the sample size for minority groups, and focuses on individuals who were educated and socialized in the UK. For the main analysis, cohabiting couples are treated as single, and separated, divorced, and widowed individuals are excluded.[6] The analysis focuses on heterosexual matches because the dataset contains too few same-sex couples to support reliable estimation.

[6]Robustness checks in Appendix J.1 include cohabiting couples in the married group.

The Census includes an optional question on religious affiliation.[7] Approximately 7 percent of individuals do not report a religion. Non-reporting is lowest among South Asians, who constitute the majority of Muslims in the UK, at around 4 percent. As a result, differential non-response by Muslims is unlikely to affect the findings. Reported religion in the Census does not necessarily reflect active religious practice and is better interpreted as an ethno-religious identity. Evidence from the Citizenship Survey (2010–11) shows that while 97.6 percent of individuals raised as Muslims identify as Muslim, only 76 percent report actively practicing the religion. Although Muslims exhibit higher practice rates than Christians (Figure A3), their rates are comparable to those of other religious groups.

A related concern is the potential endogeneity of religious affiliation in marriage. While this issue cannot be fully addressed using cross-sectional data, evidence from Scotland's 2001 Census is informative. Among married Muslims, 9.5 percent had non-Muslim spouses, and 16.7 percent of these spouses were raised as Muslims. Conversely, among Muslims who married within their group, 5.3 percent had spouses who were not Muslims in childhood. Whether these patterns reflect conversion before or after marriage is unclear. However, even if conversion rates in England and Wales mirror those observed in Scotland, they are too small to account for the large differences in intermarriage and marriage migration between Muslims and non-Muslims.

Measuring marriage migration is challenging because large administrative datasets rarely record individuals' migration intentions. I classify a marriage as involving marriage migration when it occurs between an individual who was born in the UK or arrived before age 18 and a partner who migrated after age 18. This definition may overstate marriage migration by including individuals who migrated for other reasons, such as work or study, and later married a UK resident. Nevertheless, it remains policy-relevant because these unions still involve a foreign-born spouse who can obtain settlement and citizenship rights.[8]

To gauge the extent of possible misclassification, I use data from Understanding Society, which record both the year of marriage and the year of arrival. Restricting marriage migration to cases where the spouse arrived in the same year as or after the year of marriage yields a conservative lower bound of 39.5 percent for men and 43.6 percent for women. These figures are roughly ten percentage points lower than the baseline but preserve the same gender differences and the large gap between Muslims

[7] The question asks "What is your religion?" and is voluntary.

[8] Spouse-visa holders follow a five-year route to settlement, after which they may apply for British citizenship following one additional year, subject to statutory requirements.

and non-Muslims. Appendix I replicates the full analysis using this lower-bound definition, while Appendix J.2 reports additional sensitivity analyses to the marriage migration definition. The main results are robust to these alternative measurement choices.

Given this measurement strategy, the analysis further focuses on marriage migration among ethnic minority groups by restricting attention to same-ethnicity marriages. International marriage migration is extremely rare among the ethnic majority. Only 0.4 percent of White British individuals marry a spouse who migrated to the UK after age 18. Among minorities, same-ethnicity unions account for approximately 93 percent of all observed marriage migration. Restricting the sample in this way removes cases where relationships were likely formed after migration and isolates unions in which partner selection is closely tied to cross-market matching. Accordingly, I define marriage migration as a same-ethnicity marriage involving an ethnic minority individual who was born in the UK or arrived before age 18 and a spouse who migrated after age 18.

Table 1. Summary Statistics by Religion, Ethnicity, and Gender

	Non-Muslim					
	White British		Ethnic Minorities		Muslim	
	Male	Female	Male	Female	Male	Female
Age	39.1	36.9	37.3	35.0	35.2	32.4
College education (%)	34.4	36.8	46.8	51.7	38.5	36.9
UK-born (%)	98.4	98.3	73.7	74.0	54.2	62.0
Married (%)	52.1	53.1	41.5	37.7	64.6	66.5
Marriage						
Marriage migration (%)	0.0	0.0	22.8	19.5	49.6	51.5
Inter-religious^{+} (%)	20.2	20.5	18.2	20.7	6.0	4.7
Inter-ethnic^{++} (%)	1.9	1.3	19.3	19.9	12.4	12.0
Number of observations	666,377	646,837	57,328	60,178	20,181	20,498

Notes. Marriage migration, inter-religious, and inter-ethnic marriage rates are computed conditional on being married. Marriage migration is defined only for ethnic minority individuals and is zero for White British by construction. $^{+}$Inter-religious marriage is measured based on the following religious groups: No religion, Christian, Buddhist, Hindu, Jewish, Muslim, Sikh, and others. $^{++}$Inter-ethnic marriage is measured based on the following ethnic groups: White, Black, Indian, Pakistani, Bangladeshi, Chinese, Other Asians, and Others. Mixed ethnicities are excluded. *Source.* Census for England and Wales, 2011.

Table 1 reports summary statistics for the sample. Muslims in the UK differ demographically from other groups. They are younger on average, reflecting more recent immigration, and have lower educational attainment than other minorities. Muslims marry at younger ages, and by age forty about 90 percent are married compared with 69 percent among non-Muslims.[9]

Muslims exhibit substantially higher rates of marriage migration, roughly twice those of non-Muslim ethnic minorities. This pattern persists even after conditioning on ethnicity (Figure 1), which motivates treating religion and ethnicity as distinct components of the marital surplus rather than combining them into a single group-identity term. On average, about half of Muslims who were born in the UK or arrived before age 18 are married to a spouse who migrated after age 18. This magnitude is consistent with sociological evidence indicating that 40 to 60 percent of British Pakistanis and Bangladeshis marry partners from their ancestral countries of origin (Charsley, 2018; Dale, 2008; Kibria, 2012; Charsley et al., 2020).

Research on marriage-related migration has focused primarily on migrant wives, with relatively little attention to male marriage migration (Charsley and Liversage, 2015). However, marriage migration rates are similar for Muslim men and women (Table 1). Several factors contribute to this fact. Arranged marriages reduce search costs and provide assurances through extended family networks. Women may also benefit from lower in-law influence when marrying someone raised abroad, a feature noted in some Muslim communities. As more individuals marry from abroad, those with strong endogamy preferences may increasingly search overseas to avoid remaining single, which can reinforce the pattern over time.

A key determinant of marriage migration is the migration utility, which is primarily influenced by the income disparity between the host and origin countries. Larger income gaps imply greater migration utility and therefore stronger incentives for marriage migration. This pattern is evident in Table 2, which reports the share of individuals who marry abroad by country of origin.[10] Lower income levels in the origin country are associated with higher marriage migration rates. However, these descriptive patterns cannot determine whether the higher rates reflect stronger preferences for marrying abroad, better availability of partners in origin countries, or

[9]Muslim men marry spouses who are on average 4.5 years younger, while non-Muslim men marry spouses who are 2.4 years younger.

[10]Country of origin is defined using Understanding Society (2010–11) as follows: the individual's birth country for those born outside the UK, the father's birth country for those born in the UK, and the grandfather's birth country when both the father and the individual were born in the UK.

Figure 1. Marriage Migration Rates by Ethnicity and Religion

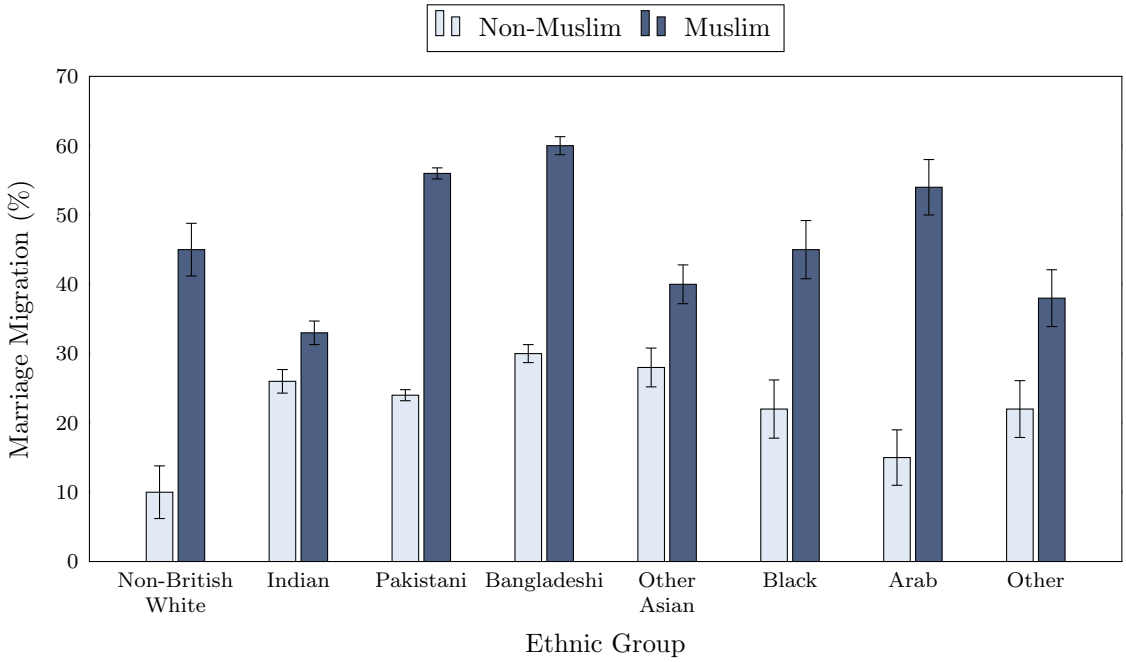

Notes. Each bar represents the percentage of married British individuals within each ethnic and religious group who married a spouse that migrated to the UK after age 18. Error bars denote 95 percent confidence intervals. Source. Census for England and Wales, 2011.

both.

Table 2. Marriage Migration Rates by Income Level of Country of Origin

Income level of ancestral country of origin	Marriage migration (% of married)
Low income	49%
Lower middle income	41%
Upper middle income	27%
High income	12%

Notes. The table reports marriage migration rates conditional on being married among British individuals with a non-UK country of origin. The country of origin is determined by the first generation in the patrilineal line (individual, father, or grandfather) born outside the UK. *Source.* Understanding Society, 2010-11, and World Bank Data, 2011.

Muslims not only have higher rates of marriage migration but also lower rates of intermarriage. As shown in Table 1, the marriage market exhibits strong endogamy,

with a high prevalence of marriages within the same ethnic and religious groups. This pattern remains even after conditioning on ethnicity (Figure A4a), and similar trends appear for inter-ethnic marriages (Figure A4b).

For the remainder of the analysis, I aggregate several ethnic groups to increase statistical power. Most Muslims in the UK are of Pakistani, Bangladeshi, or Indian origin. I combine the Pakistani and Bangladeshi populations into a single group because they constitute similar shares of the Muslim population, share Islamic practices, and have had comparable income levels over recent decades. The resulting ethnic categories used in the analysis are White British, Other White, Indian, Pakistani/Bangladeshi, and Others.

3.2 Sorting into Marriage Migration

Selection into marriage migration is not random but systematically related to individual characteristics. To examine how observable traits differ between those who engage in marriage migration and those who do not, I estimate the following reduced-form linear probability model:

$$\text{CollegeEducation}_i = \alpha + \beta\,\text{MarriageMigration}_i + \mathbf{X}_i'\theta + \delta_{r(i)} + \gamma_{c(i)} + \varepsilon_i, \qquad (1)$$

where $\text{CollegeEducation}_i$ equals 1 if individual i holds a college degree. The indicator $\text{MarriageMigration}_i$ equals 1 if the individual married a spouse who migrated to the UK after age 18. The vector $\mathbf{X}_i$ includes individual-level controls, while $\delta_{r(i)}$ and $\gamma_{c(i)}$ denote region and 10-year birth-cohort fixed effects, respectively. Region fixed effects control for unobserved differences across local marriage markets, ensuring that the estimates reflect within-region variation in marriage and education outcomes.

Table 3 reports the results of the regression. Across all specifications, individuals who marry abroad are significantly less likely to hold a college degree than those who marry locally. This negative association appears for both Muslims and non-Muslims, and for men and women alike. The education gap is largest among Muslim women, who are about 14 percentage points less likely to have a college degree relative to comparable women who marry locally, conditional on controls. Muslim men also exhibit a sizeable gap of roughly 10 percentage points. Among non-Muslims, the education gap is smaller but remains substantial, particularly for women, where the coefficient reaches -0.11.

These results indicate that selection into marrying abroad is negatively related to

Table 3. Educational Sorting in Marriage Migration (Linear Probability Model)

	Dependant variable: College education			
	Muslim		Non-Muslim	
	Male	Female	Male	Female
Baseline specification				
Marriage migration	-0.111***	-0.139***	-0.038***	-0.098***
	(-13.1)	(-17.4)	(-4.6)	(-11.0)
R^2	0.014	0.023	0.001	0.006
Adjusted specification				
Marriage migration	-0.097***	-0.137***	-0.052***	-0.114***
	(-11.4)	(-17.2)	(-6.3)	(-12.8)
Region & cohort FE	Y	Y	Y	Y
R^2	0.038	0.056	0.038	0.047
N	12,458	12,857	21,060	19,687

Notes. The table reports linear probability models where the dependent variable is an indicator equal to one if the individual has a college degree. Baseline specifications include ethnic group, region, and 10-year birth-cohort fixed effects. Adjusted specifications additionally control for age and spouse's religion and ethnicity. Sample is limited to ethnic minorities. t statistics in parentheses. *p<0.10, **p<0.05, ***p<0.001. *Source.* Census for England and Wales, 2011.

education, and that this negative selection is more pronounced among Muslims, especially women. Including controls and fixed effects increases precision without altering the substantive conclusions. As a robustness check, I re-estimate the specification using a logit model. The results, reported in Table A1, are quantitatively similar, confirming that the negative association between marrying abroad and education is robust to nonlinear specifications.

Marriage migrants are not a random draw from the populations in their countries of origin. To assess selection on education, I construct country-level educational distributions using the Indian Demographic and Health Survey (1998–99), the 2001 Bangladesh Census, and the 1998 Pakistan Census, which provide comparable measures of education by gender for the relevant marriage cohorts. Figure 2 compares the educational attainment of marriage migrants in the UK to that of the origin-country population, separately by country and gender. Panels (a) and (b) show that Indian male and female marriage migrants are substantially more educated than the

Figure 2. Educational attainment of marriage migrants and origin-country population

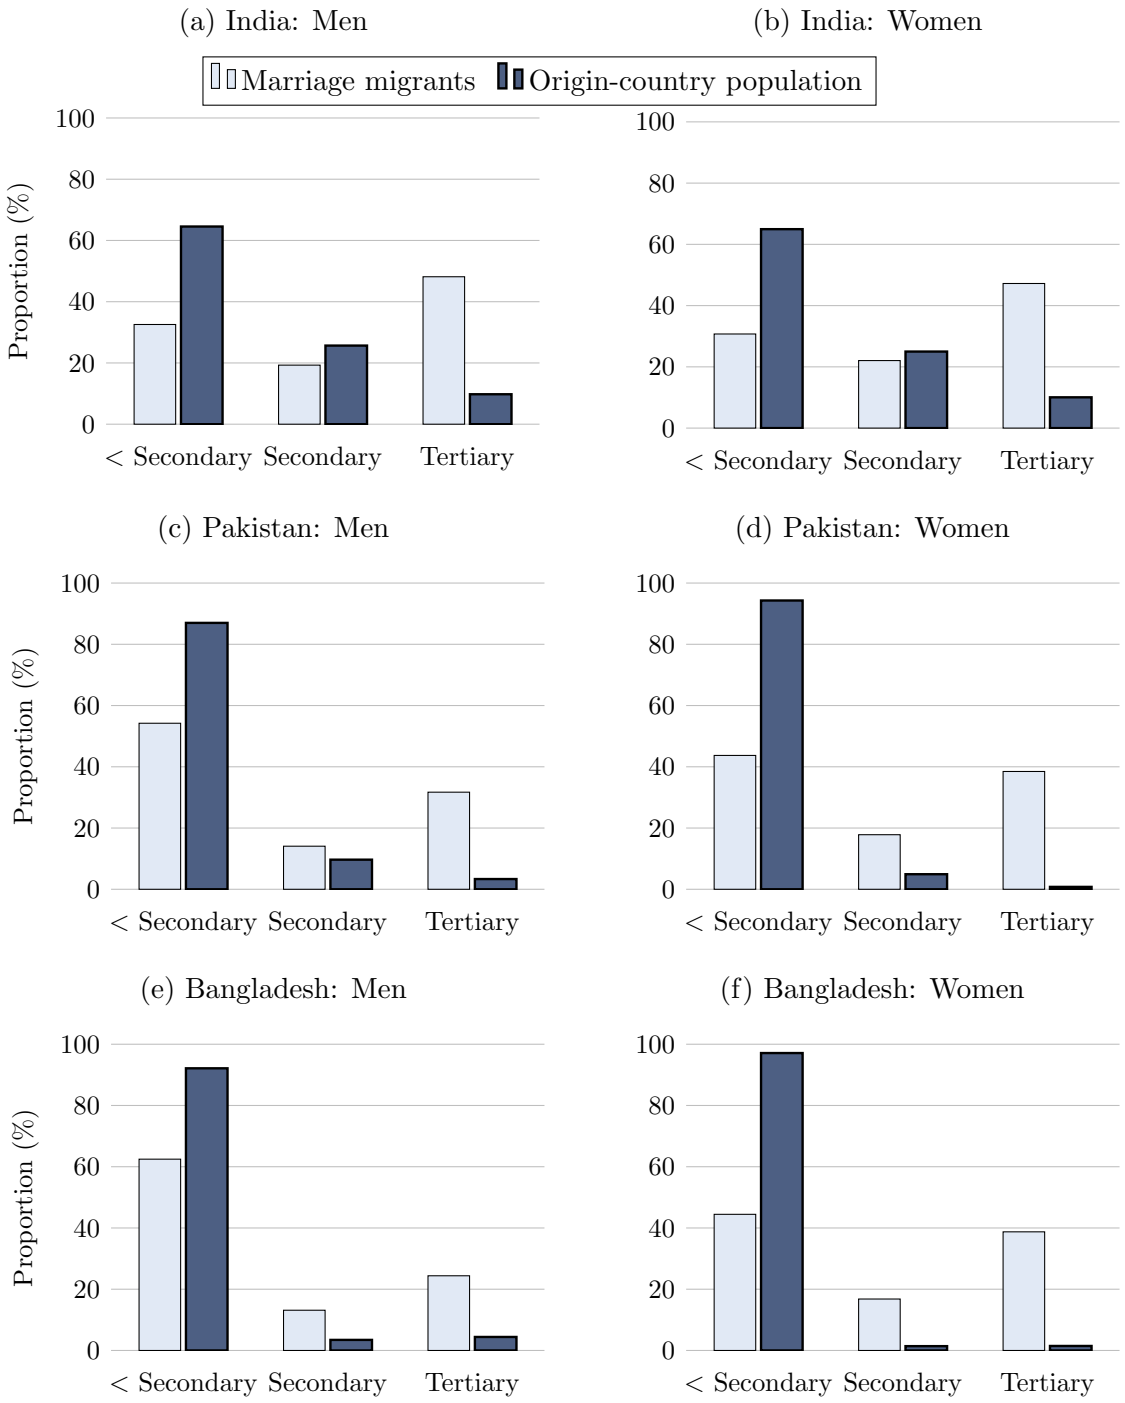

Notes. Each panel compares the educational attainment of marriage migrants in the UK to the educational distribution in the origin country. *Sources.* Census for England and Wales (2011); India DHS (1998–99); Pakistan Census (1998); Bangladesh Census (2001).

population in India. Nearly half of migrants of both genders hold tertiary education, compared to roughly 10 percent in the origin population. Panels (c)–(f) document even stronger positive selection for migrants from Pakistan and Bangladesh. In both countries, the education gap between migrants and non-migrants is large for men and particularly pronounced for women. Approximately 38 percent of Pakistani female marriage migrants and 39 percent of Bangladeshi female marriage migrants have tertiary education, compared to less than 1 percent among non-migrant women in the origin populations.

3.3 Trade-offs in Marriage Migration

To assess how individuals trade off partner characteristics against the option of marrying abroad, it is necessary to compare the attributes of spouses drawn from the outside market with those of otherwise similar spouses selected locally. I implement this comparison by exploiting variation in marriage choices among individuals with comparable observable characteristics, some of whom marry abroad while others marry locally. To formalize this comparison, I estimate the following linear probability model:

$$\begin{aligned}\text{CollegeEducation}_{s(i)} = \alpha + \beta_1 \,\text{MarriageMigration}_i + \beta_2 \,\text{CollegeEducation}_i \\ + \beta_3 \left(\text{MarriageMigration}_i \times \text{CollegeEducation}_i\right) + \mathbf{X}_i'\theta + \delta_{r(i)} + \gamma_{c(i)} + \varepsilon_i, \quad (2)\end{aligned}$$

where the dependent variable $\text{CollegeEducation}_{s(i)}$ equals 1 if the spouse of individual i holds a college degree. All other variables and fixed effects are defined as in Equation (1). The interaction term captures how the association between marriage migration and spousal education varies with the individual's own education level.

Table 4 reports the regression results. The coefficient β_2 captures assortative matching by education. Across all groups, individuals with a college degree are substantially more likely to have a college-educated spouse, consistent with strong positive sorting on education. The positive and statistically significant coefficient β_1 indicates that marriage migration allows less-educated individuals to marry spouses with higher educational attainment, suggesting that cross-border matching expands opportunities for upward educational pairing. In contrast, the interaction coefficient β_3 is negative and significant, implying that this advantage reverses for highly educated individuals. Those who marry abroad are less likely to have an equally educated spouse than comparable individuals who marry locally, indicating a tendency toward

Table 4. Educational trade-offs in marriage migration

	Dependent variable: Spouse's college education			
	Muslim		Non-Muslim	
	Male	Female	Male	Female
Baseline specification				
Marriage migration	0.019	0.052***	0.075***	0.060***
	(1.9)	(5.3)	(7.1)	(5.2)
College education	0.381***	0.405***	0.409***	0.408***
	(32.9)	(32.5)	(56.6)	(55.5)
Marriage migration	-0.122***	-0.112***	-0.107***	-0.097***
× College education	(-7.2)	(-6.2)	(-7.0)	(-5.8)
R^2	0.114	0.108	0.150	0.152
Adjusted specification				
Marriage migration	0.019	0.048***	0.064***	0.047***
	(1.9)	(4.9)	(6.0)	(4.1)
College education	0.356***	0.381***	0.382***	0.392***
	(30.9)	(30.4)	(52.5)	(52.6)
Marriage migration	-0.115***	-0.101***	-0.099***	-0.087***
× College education	(-6.9)	(-5.5)	(-6.5)	(-5.2)
R^2	0.136	0.120	0.171	0.162
N	12,450	12,847	21,042	19,668

Notes. The table reports linear probability models where the dependent variable is an indicator equal to one if the spouse has a college degree. Baseline specifications include controls for own education, ethnic group, region, and 10-year birth-cohort fixed effects. Adjusted specifications additionally control for spouse's religion, spouse's ethnicity, and the age gap between spouses. t statistics in parentheses. *p<0.10, **p<0.05, ***p<0.001. *Source.* Census for England and Wales, 2011.

downward educational matching.

To assess robustness, I re-estimate the specification using a logit model. The marginal effects of marriage migration, reported in Appendix Figures A5–A6, closely mirror the linear probability results. The direction, magnitude, and statistical significance of the coefficients remain stable, confirming that the estimated trade-offs are robust to nonlinear specifications.

Marriages involving migration display systematically larger age gaps than local unions (Table A2). Among Muslims, husbands in marriage-migrant unions are on average about 0.5 years older than their wives compared with locally matched couples.

Among non-Muslims, the pattern depends on which spouse is the migrant. When the wife is the migrant, the spousal age gap is roughly 1.5 years larger than in local marriages, whereas when the husband is the migrant, the gap is about 0.7 years. These patterns are consistent with well-documented gender asymmetries in age preferences, whereby men tend to seek younger partners and women prefer older partners (Buss, 1989; Kenrick and Keefe, 1992; Bech-Sørensen and Pollet, 2016).

To summarize the reduced-form findings, individuals in the UK with below-average education are more likely to marry spouses from their countries of origin who have above-average education relative to the local population there. Thus, low-educated individuals tend to marry up when choosing a spouse from abroad, accessing more highly educated partners than they would in the local market. By contrast, among highly educated individuals, those who marry internationally are less likely to match with an equally educated spouse than comparable individuals who marry locally.

These patterns reflect the joint influence of two forces. First, selection into the international option varies with own education because migration-related gains and network access differ across groups, which shifts the set of feasible and attractive international opportunities. Second, there is an equilibrium composition effect. When low-educated individuals disproportionately marry up abroad, the remaining local market becomes more positively assortative in education, so the expected education of a local spouse for a highly educated individual rises simply because the local pool tilts toward high education.

Overall, the descriptive evidence documents substantial differences across ethnic and religious groups in marriage migration, intermarriage, and partner characteristics. These correlations, however, do not reveal the mechanisms underlying the observed sorting. In particular, the reduced-form patterns cannot distinguish whether high rates of intra-group marriage and marriage migration reflect strong endogamy preferences, the attractiveness of partner pools in origin countries, or equilibrium effects arising from population composition. To separate these forces, I now turn to a structural matching framework that links observed matches to their underlying surplus components.

4 The Model

Existing work shows that marriage and location choices are often jointly determined even within national borders, as individuals negotiate both partner selection and

residence (Beauchamp et al., 2021; Dupuy, 2021). I extend this insight to an international setting by allowing matches to form across markets linked through marriage migration. I introduce an open-economy matching model in which individuals are distributed across multiple locations and may form matches either within the same location or across locations. This feature departs from standard marriage market models, which typically assume a closed economy with no cross-market matching. Location is modeled as a trait in the type space, but it is distinct from other traits because cross-market matches require one partner to relocate. Unlike traits such as education or religion, location affects match surplus through a co-location constraint: partners initially in different locations cannot consume the match without one partner relocating, which introduces a bilateral migration component that has no analogue for other characteristics.

I then specialize to an inside–outside variant that collapses geography into a resident market and a set of external partner pools. I introduce this restricted model to obtain identification of the contributions of the main drivers of marriage migration, namely outside market value and endogamy preferences, using data observed for a single destination country. The restricted formulation preserves the core comparative statics of the general framework and maps directly to the information typically available in microdata on marriage migration.

This model centers on one-to-one matching in marriage markets, but it is more general. It can be employed to examine various types of two-sided matching, such as between CEOs and firms. Moreover, the model's framework can be adapted to analyze one-to-many matching scenarios (Corblet, 2022), enabling the study of matches between students and schools or workers and firms.

4.1 General Marriage Migration Model

I model heterosexual marriages within a two-sided transferable utility matching framework that allows individuals to form matches across locations.[11] Under the standard transferable utility assumption, observed matching patterns summarize preferences and all matching frictions in reduced-form utilities (Becker, 1973; Choo and Siow, 2006; Galichon and Salanié, 2022).

Individuals need not have perfect information, nor is matching required to be fric-

[11] Extending the framework to same-sex unions would require a unipartite version of the model in which both partners are drawn from the same population and the equilibrium matching distribution satisfies a symmetry constraint (Ciscato et al., 2020).

tionless. Systematic differences in the ease of meeting partners, including information frictions, transaction costs, and institutional barriers associated with cross-market matching, enter the model through match-specific utilities rather than through an explicit search technology. Equilibrium matching outcomes therefore reflect the joint influence of preferences and frictions, bundled into the deterministic surplus components that are identified from the data.

Figure 3. Local Marriage Markets and Cross-market Matching

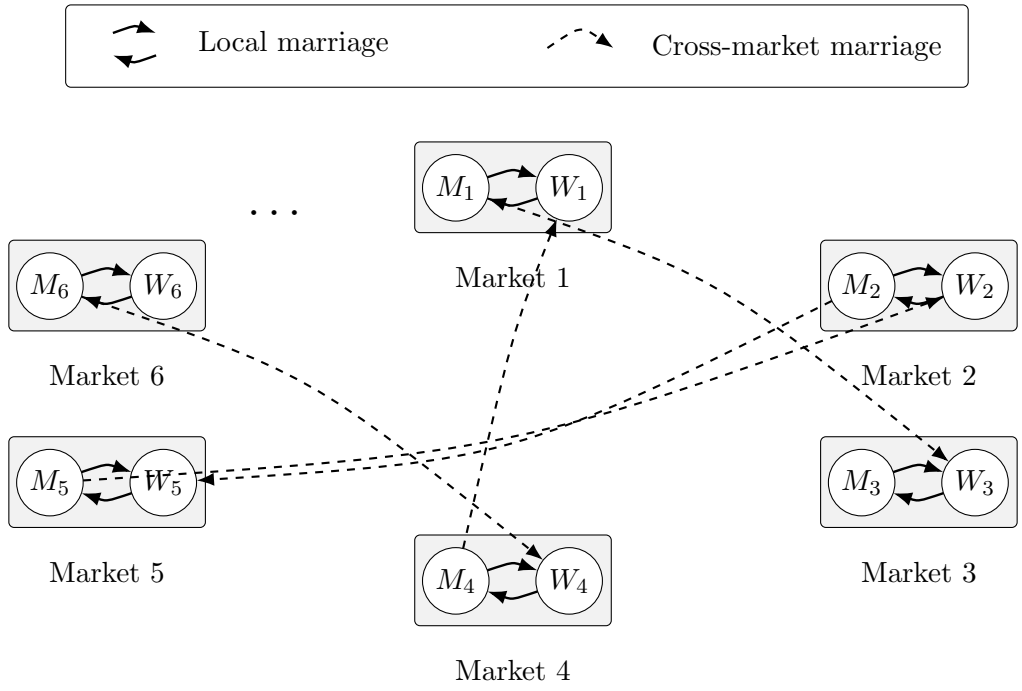

The model comprises L disjoint local marriage markets, each corresponding to a country, a state, or any geographically or institutionally defined area. Individuals are distinguished by observable characteristics. Let $x \in \mathcal{X} = \{1, \ldots, X\}$ index women's types and $y \in \mathcal{Y} = \{1, \ldots, Y\}$ index men's types, where each type x and y corresponds to a vector of observable traits (such as age, education, religion, and ethnicity). In market $\ell \in \{1, \ldots, L\}$, $n_{x,\ell}$ and $n_{y,\ell}$ denote the mass of women of type x and men of type y, respectively. Hence, the distribution of (vector-valued) types can vary across locations (Figure 3).

4.1.1 Local Matches

First, I restrict attention to matches that occur within a local market ℓ. Marital utility is specified in a random utility framework with a deterministic component and

an idiosyncratic taste shock, following McFadden (1974). Consistent with standard practice, I assume that these components are additively separable.[12] This implies that observable characteristics affect utility only through the deterministic term, and there is no interaction between observed and unobserved components. The utility of a woman i of type x who matches with a man j of type y in market ℓ is

$$U_{ij} = U(x, y, \ell, \ell) + \varepsilon_{iy,\ell\ell}, \tag{3}$$

and the utility of a man j of type y who matches with a woman i of type x in the same market is

$$V_{ji} = V(y, x, \ell, \ell) + \eta_{jx,\ell\ell}. \tag{4}$$

where $U(x, y, \ell, \ell)$ and $V(y, x, \ell, \ell)$ denote the deterministic components of utility, which depend only on observed characteristics of both partners and the market in which the match occurs. The notation (ℓ, ℓ) indicates that both partners belong to the same local marriage market. The deterministic utility of remaining single is normalized to zero for all types. The terms $\varepsilon_{iy,\ell\ell}$ and $\eta_{jx,\ell\ell}$ are idiosyncratic taste shocks capturing unobserved preferences of woman i for men of type y and of man j for women of type x, respectively. These shocks are assumed to be identically and independently distributed and independent of observable characteristics. Under this structure, the deterministic components $U(\cdot)$ and $V(\cdot)$ summarize the systematic contribution of observable traits relative to unobserved heterogeneity.

Demand and supply for a particular match type do not necessarily align. Mismatches can arise because partners derive asymmetric utilities from the same match, or because imbalances in the population distribution across types affect outside options on each side of the market. In a transferable-utility setting, such imbalances are absorbed through intra-household transfers. Individuals who face stronger competition for their preferred partners compensate by accepting a smaller share of the surplus, effectively raising the partner's payoff. Transfers may be monetary, such as bride price or dowry, or non-monetary, such as post-marital commitments of time, effort, or resources (Chiappori, 2017). I assume perfectly transferable utility, so that the total surplus generated by a match can be redistributed between spouses without

[12] This assumption may be challenged if individuals have idiosyncratic preferences for unobserved characteristics of their potential partners. However, Chiappori et al. (2019) demonstrates that even when the assumption is only approximately correct, it generates only small biases.

loss. Under this assumption, transfers adjust utilities so as to support a stable allocation, while the sum of utilities within each couple is pinned down by the underlying surplus function. Equilibrium matchings therefore maximize aggregate surplus across all couples in the market.

Therefore, the deterministic component of the marriage payoff has two parts: (i) the utility generated by the partners' observable characteristics, and (ii) an equilibrium transfer that reallocates surplus within the couple. Transfers can be positive or negative and may depend on both partner types and location. For a woman of type x and a man of type y who match in location ℓ, deterministic utilities are given by

$$U(x, y, \ell, \ell) = f(x, y) - \tau_{xy,\ell\ell}, \qquad V(y, x, \ell, \ell) = g(y, x) + \tau_{xy,\ell\ell}.$$

where $f(x, y)$ denotes the deterministic utility that a woman of type x obtains from matching with a man of type y based solely on their observable characteristics, and $g(y, x)$ analogously denotes the deterministic utility that a man of type y obtains from matching with a woman of type x. Since transfers occur within the couple, the transfer term enters with opposite sign in the two utilities.[13] The joint deterministic surplus from the match is

$$\Phi(x, y, \ell, \ell) = f(x, y) + g(y, x),$$

which depends only on partner characteristics and is independent of $\tau_{xy,\ell\ell}$. The transfer reallocates surplus between spouses but does not affect its level. Consequently, the total value of each match can be identified from equilibrium matching patterns without observing or estimating the transfers themselves.

4.1.2 Cross-market Matches

I now extend the framework to allow for matches between partners residing in different locations. Consider a potential match between a woman of type x in location ℓ and a man of type y in location $k \neq \ell$. The deterministic surplus from such a match has two components. The first is the intrinsic match surplus $f(x, y) + g(y, x)$, which depends only on the partners' observable characteristics and is common to both spouses, as in the case of local matches. The second component is a *migration utility*, which arises because one partner must relocate for the match to be realized. If the man moves

[13] The sign convention is without loss of generality: a positive $\tau_{xy,\ell\ell}$ can be interpreted as a transfer from the woman to the man, and a negative value as a transfer in the opposite direction.

from k to ℓ, he receives a migration utility $b^m_{yx,k\ell}$; if instead the woman moves from ℓ to k, she receives a migration utility $b^f_{xy,\ell k}$. These migration terms capture location- and type-specific gains or costs from mobility and are what distinguish cross-market matches from purely local ones.

A key distinction between this setup and one in which location is treated as a fixed individual characteristic is that partners who start in different locations cannot remain geographically separated once matched. One partner must relocate to the other's location.[14] The couple is assumed to choose the location that maximizes their joint surplus, taking into account intrinsic match surplus and migration utilities. In equilibrium, the partner for whom migration is relatively more attractive (or less costly) relocates, and the resulting transfer schedule reallocates surplus to support this location choice. The aggregate migration utility for a cross-market couple is defined as

$$b_{\ell k}(x,y) = \max\left\{b^f_{xy,\ell k},\ b^m_{yx,k\ell}\right\},$$

The term $b_{\ell k}(x,y)$ represents the *migration utility* associated with relocating from location ℓ to k in the context of a cross-market match between a woman of type x and a man of type y. It captures the net contribution of the location choice to the couple's deterministic surplus, encompassing economic, social, legal, and psychological factors that influence the attractiveness of moving. In addition, $b_{\ell k}(x,y)$ incorporates any systematic preference for the partner's location or origin that is not already captured by the intrinsic match surplus $\Phi(x,y)$. For example, if an individual derives utility from having a spouse who resides in, or has grown up in, a particular country or region, this preference is reflected in $b_{\ell k}(x,y)$. A positive value of $b_{\ell k}(x,y)$ indicates that relocation yields a net gain in utility, while a negative value corresponds to a net migration cost that reduces the joint surplus.

The total deterministic surplus from a cross-market match between a woman of type x in ℓ and a man of type y in k is therefore

$$\Phi(x,y,\ell,k) = f(x,y) + g(y,x) + b_{\ell k}(x,y). \tag{5}$$

The specification nests local matches as a special case where $\ell = k$ and $b_{\ell k}(x,y) = 0$. The additive form assumes separability between the intrinsic component of match surplus, given by $f(x,y)+g(y,x)$, and the migration utility. If data were observed for all markets, interaction effects between migration utilities and partner characteristics

[14] I abstract from the possibility of stable long-distance relationships.

could be identified directly. In the present setting, the assumption is imposed solely to achieve identification for the model in the next subsection.

In principle, $b_{\ell k}(x, y)$ can be parametrized analogously to a gravity equation, where migration utilities depend on both economic incentives and separation between locations. For illustration,

$$b_{\ell k} = \alpha \cdot \log\left(\frac{\text{GDP}_k}{\text{GDP}_\ell}\right) - \delta \cdot \text{Distance}_{\ell k},$$

where GDP_k and GDP_ℓ capture differences in economic opportunities and $\text{Distance}_{\ell k}$ reflects physical, cultural, or institutional distance between the two locations (Beine et al., 2016).

4.1.3 Matching Equilibrium

Each individual can either marry locally, marry across markets (with one partner relocating), or remain single. Let μ denote the equilibrium matching distribution, representing the number (or mass) of matches of each type. The notation distinguishes between local and cross-market matches and identifies the direction of migration in the latter case. Specifically, for all locations ℓ and k (with $\ell \neq k$ in the cross-market case):

1. $\mu_{xy,\ell,\ell}$: matches between women of type x and men of type y, both residing in location ℓ;
2. $\mu_{xy,k\to\ell,\ell}$: matches between women of type x residing in location k and men of type y in location ℓ, where the woman migrates to ℓ after marriage;
3. $\mu_{xy,\ell,k\to\ell}$: matches between women of type x residing in location ℓ and men of type y in location k, where the man migrates to ℓ after marriage;
4. $\mu_{x0,\ell}$: single women of type x in location ℓ;
5. $\mu_{0y,\ell}$: single men of type y in location ℓ.

Feasibility requires that each individual is involved in at most one match. Hence, for each type and origin location, the total number of individuals must equal the sum of those who marry (locally or across borders) and those who remain single.

For women of type x originating in location ℓ,

$$n_{x,\ell} = \mu_{x0,\ell} + \sum_y \left[\mu_{xy,\ell,\ell} + \sum_{k\neq\ell} (\mu_{xy,\ell,k\to\ell} + \mu_{xy,\ell\to k,k})\right], \quad \forall x, \ell, \tag{6}$$

where $\mu_{xy,\ell,\ell}$ are local matches in ℓ, $\mu_{xy,\ell,k\to\ell}$ are matches in which the man migrates from k to ℓ and the woman remains in ℓ, and $\mu_{xy,\ell\to k,k}$ are matches in which the woman migrates from ℓ to k.

Similarly, for men of type y originating in location ℓ,

$$n_{y,\ell} = \mu_{0y,\ell} + \sum_x \left[\mu_{xy,\ell,\ell} + \sum_{k\neq\ell} \left(\mu_{xy,k\to\ell,\ell} + \mu_{xy,\ell\to k,k} \right) \right], \quad \forall y, \ell, \tag{7}$$

where $\mu_{xy,k\to\ell,\ell}$ are matches in which the woman migrates from k to ℓ and the man remains in ℓ, and $\mu_{xy,\ell\to k,k}$ are matches in which the man migrates from ℓ to k. These constraints ensure that the matching distribution μ exhausts the available populations $n_{x,\ell}$ and $n_{y,\ell}$ without double-counting any individual.

If the distribution of idiosyncratic taste shocks is left unrestricted, the one-to-one transferable-utility model is underidentified (Galichon and Salanié, 2022), because any observed matching pattern can be rationalized by a suitable specification of unobserved heterogeneity (Gualdani and Sinha, 2023). To address this issue and obtain a tractable empirical model, I follow Choo and Siow (2006) and assume that taste shocks are independently and identically distributed according to an extreme value type I distribution.[15] This assumption transforms the model into a tractable two-sided just-identified logit model, suitable for studying the main mechanisms in the marriage market.

Under a logit specification with separable utilities, a modified Independence of Irrelevant Alternatives (IIA) property holds in matching models: the double odds ratios $\mu_{xy}\mu_{zt}/(\mu_{xt}\mu_{zy})$ are invariant to changes in subpopulation sizes (Galichon and Salanié, 2017).[16,17]

In equilibrium, each individual chooses the option that maximizes her or his expected utility. A feasible matching configuration is *stable* if (i) no unmatched pair would both prefer to match with one another rather than accept their current outcomes, and (ii) no matched individual would prefer to remain single rather than stay with their current partner. The probability that a woman of type x residing in

[15] If a random variable ε follows an extreme value type I (Gumbel) distribution, its cumulative distribution function is $F(\varepsilon) = \exp(-\exp(-\varepsilon))$ and its probability density function is $f(\varepsilon) = \exp(-\varepsilon)\exp(-\exp(-\varepsilon))$.

[16] In a one-sided logit model, the standard IIA property implies that single odds ratios are invariant to population scale.

[17] Appendix C provides formal statements and proofs.

location ℓ matches with a man of type y in location k is given by:[18]

$$P_{x,\ell}(y,k) = \Pr\Big\{(y,k) = \underset{z\in\mathcal{Y},\, m\in\{0,1,\ldots,L\}}{\arg\max} \big[U(x,z,\ell,m) + \varepsilon_{iz,\ell m}\big]\Big\}$$
$$= \frac{\exp\big[U(x,y,\ell,k)\big]}{1 + \sum_{z\in\mathcal{Y}}\sum_{m=1}^{L} \exp\big[U(x,z,\ell,m)\big]} \tag{8}$$

This expression shows that the probability of observing a match depends on the relative magnitude of the deterministic utility associated with that specific match, normalized by the sum of the exponentiated utilities of all possible alternatives, including remaining single. Intuitively, people choose among alternative possible spouses based on their relative attractiveness. The probability of choosing a spouse is determined by comparing the systematic utility of that choice to the sum of the systematic utilities of all available options.

In equilibrium, transfers (or equivalently, the implicit prices) adjust to clear the market such that aggregate marriage flows are consistent with these choice probabilities. Accordingly, the deterministic joint surplus satisfies:

$$\Phi(x,y,\ell,k) = f(x,y) + g(y,x) + b_{\ell k}(x,y) = 2\ln\left[\frac{\mu_{xy,k\to\ell,\ell}}{\sqrt{\mu_{x0,\ell}\,\mu_{0y,k}}}\right], \quad \forall\, \ell \neq k. \tag{9}$$

For local matches, where both partners reside in the same location, the corresponding joint surplus is:

$$\Phi(x,y,\ell,\ell) = f(x,y) + g(y,x) = 2\ln\left[\frac{\mu_{xy,\ell,\ell}}{\sqrt{\mu_{x0,\ell}\,\mu_{0y,\ell}}}\right], \quad \forall\, \ell. \tag{10}$$

Taken together, the feasibility constraints, which guarantee that every individual appears in at most one match and that population balances hold in each type–location cell, and the stability conditions, which rule out any blocking pairs or profitable deviations to singlehood, completely determine the equilibrium matching vector μ^* for a general market with possibility of migration for marriage.

[18] See Appendix B for the full derivation.

4.2 Restricted Model with Single-Destination Data

A central empirical challenge in studying international marriage migration is the lack of harmonized micro-level data spanning both sending and receiving countries. National statistical agencies typically collect detailed demographic and marital information only for residents within their own borders, while systematic data on emigrants are often sparse or entirely unavailable. Consequently, in an international setting, the full set of information required to estimate the structural parameters of the general model is rarely observed jointly for all sides of the marriage market.

I therefore develop a restricted version of the model that can be estimated using microdata observed exclusively for a single destination country. Although linked origin–destination data would strengthen identification, a destination-only framework remains informative in settings where international marriages are predominantly formed through the immigration of spouses from lower-income origin countries to high-income destinations, as is the case for ethnic minorities in many advanced economies.

I abstract from competition among alternative destinations. This mirrors the institutional structure of spouse-visa sponsorship, in which the resident sponsor determines the destination subject to legal and financial requirements. A prospective spouse abroad cannot independently migrate without sponsorship and rarely faces multiple destination offers.[19] Accordingly, once a resident chooses to marry abroad, the relevant outside option for the potential migrant spouse is to forgo migration rather than to seek entry into another host country. Uneven diaspora geographies reinforce this focus. For many origins, one high-income destination hosts a numerically dominant and well-established community, with dense networks, language familiarity, and extended kin that make it the natural focal point for marriage search. Flows to secondary destinations are small in absolute terms and do not meaningfully alter the pool of potential partners. Cross-market marriages are also largely unidirectional from lower-income origins to higher-income hosts.[20]

Therefore, in the restricted model, I summarize the net attractiveness of marrying abroad using a single reduced-form outside market term. This object captures the incentives associated with forming an international marriage that are not mediated

[19] Qualitative evidence on Pakistani and Bangladeshi transnational marriages is consistent with this arrangement: the UK-based partner or family typically drives the decision, while the overseas fiancé(e) has limited agency over the destination (Charsley et al., 2012).

[20] In the UK, only about 2,000 residents per year emigrate to join a partner, 1.3 percent of long-term emigrants and less than 0.02 percent of the population, and such moves overwhelmingly target other high-income countries such as Australia (Office for National Statistics, 2024; Australian Department of Home Affairs, 2023).

by local marriage-market conditions. Any residual influence of alternative destinations for a small subset of couples does not affect the identification of endogamy preferences and is absorbed into this term. If competition from multiple destinations were quantitatively important, the outside market term should be interpreted as an expectation over potential destinations, and the policy counterfactuals that vary visa costs should be read as partial-equilibrium exercises holding other destinations fixed.

Figure 4. Inside Market with Group-specific Outside Options

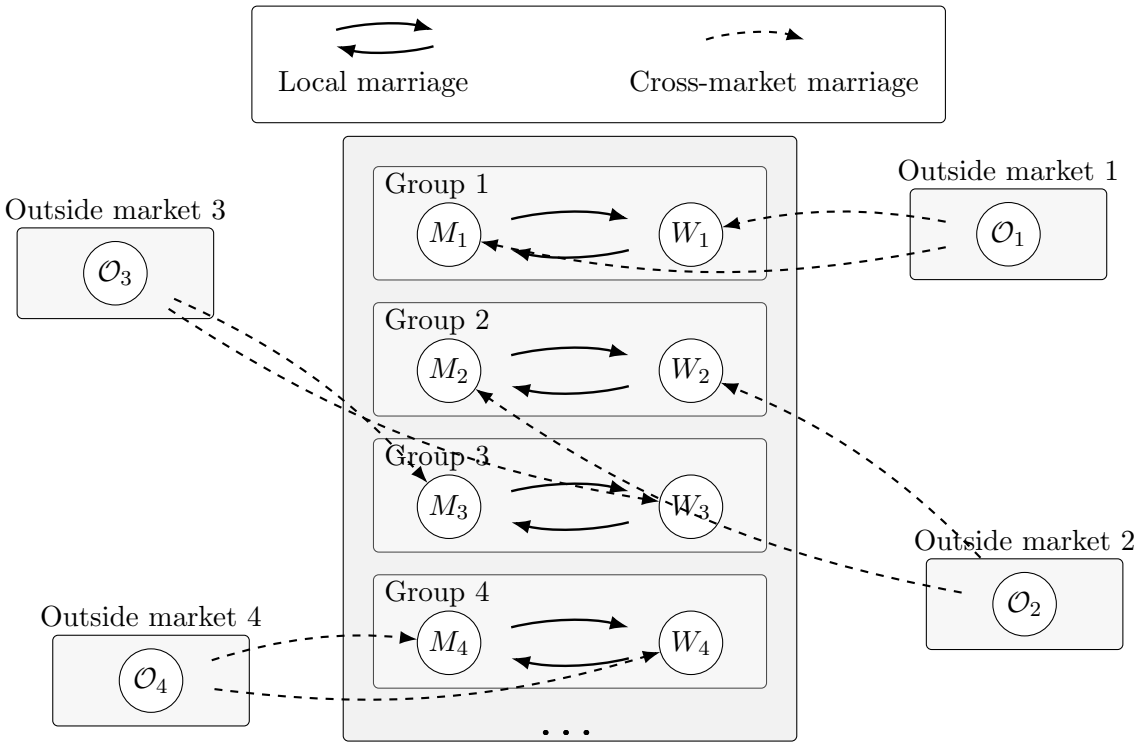

In the restricted version of the model, every individual can match in the "inside" market, and members of ethnic minority groups additionally have access to a group-specific "outside" market located in their country of origin (see Figure 4). For example, a UK resident of Pakistani origin draws potential partners from the Pakistani outside pool, a Bangladeshi resident from the Bangladeshi pool, and so on. Formally, an individual of group g faces the option set $S_g = \mathcal{L} \cup O_g$, where $\mathcal{L}$ is the resident market and O_g is the co-ethnic outside pool, with $O_0 = \varnothing$ for the majority.[21]

[21]This assumption aligns with the data. First, 93% of marriage migration among UK ethnic minorities is within ethnicity. Second, for the majority group, only 0.4% of marriages involve marrying abroad, implying that an effective foreign pool is absent.

I distinguish match types using the following notation. Superscripts indicate the locus of the match: 'in' for matches formed entirely within the local market, 'out' for matches formed entirely outside, and 'm' for cross-market matches in which one spouse enters the local market. Arrows denote the direction of the migrant's movement, and the subscript 0 denotes single individuals:

1. μ_{xy}^{in}: Number of matches between women of type x and men of type y inside the country
2. μ_{xy}^{out}: Number of matches between women of type x and men of type y outside the country
3. $\mu_{x\to y}^{\text{m}}$: Number of matches where women of type x from outside the country marry men of type y within the country (woman is the migrant spouse)
4. $\mu_{y\to x}^{\text{m}}$: Number of matches where men of type y from outside the country marry women of type x within the country (man is the migrant spouse)
5. μ_{x0}^{in}: Number of single women of type x inside the country
6. μ_{0y}^{in}: Number of single men of type y inside the country
7. μ_{x0}^{out}: Number of single women of type x outside the country
8. μ_{0y}^{out}: Number of single men of type y outside the country

The econometrician observes only outcomes in the local market and the cross-market matches that bring a spouse into that market. The observable objects are $\mu_{x0}^{\text{in}}, \mu_{0y}^{\text{in}}, \mu_{xy}^{\text{in}}, \mu_{x\to y}^{\text{m}}, \mu_{y\to x}^{\text{m}}, n_x^{\text{in}}, n_y^{\text{in}}$ for each x and y. Matches and singles formed entirely in outside markets, such as $\mu_{\cdot\cdot}^{\text{out}}$ and $\mu_{\cdot 0}^{\text{out}}$, are not observed. Feasibility requires that each individual is matched with at most one partner. The feasibility conditions are therefore:

$$n_x^{\text{in}} = \sum_{y=1}^{Y} \mu_{xy}^{\text{in}} + \sum_{y=1}^{Y} \mu_{y\to x}^{\text{m}} + \mu_{x0}^{\text{in}}, \quad x = 1, ..., X \tag{11}$$

$$n_y^{\text{in}} = \sum_{x=1}^{X} \mu_{xy}^{\text{in}} + \sum_{x=1}^{X} \mu_{x\to y}^{\text{m}} + \mu_{0y}^{\text{in}}, \quad y = 1, ..., Y \tag{12}$$

I adopt the same payoff specification as in the general model. For marriages

formed in the inside market, the marriage surplus is given by:

$$\Phi_{xy} \equiv \Phi(x, y, \text{in}, \text{in}) = f(x, y) + g(y, x) = 2 \ln \left[\frac{\mu^{\text{in}}_{xy}}{\sqrt{\mu^{\text{in}}_{x0} \mu^{\text{in}}_{0y}}} \right] \tag{13}$$

For marriages that involve migration:

$$\Phi(x, y, \text{in}, \text{out}) = \Phi_{xy} + b_{y \to x} = 2 \ln \left[\frac{\mu^{\text{m}}_{y \to x}}{\sqrt{\mu^{\text{in}}_{x0} \mu^{\text{out}}_{0y}}} \right]$$

$$\Phi(x, y, \text{out}, \text{in}) = \Phi_{xy} + b_{x \to y} = 2 \ln \left[\frac{\mu^{\text{m}}_{x \to y}}{\sqrt{\mu^{\text{out}}_{x0} \mu^{\text{in}}_{0y}}} \right]$$

where $\Phi(x, y, \text{in}, \text{out})$ denotes the deterministic match value when a husband of type y migrates to join a wife of type x, and $\Phi(x, y, \text{out}, \text{in})$ the converse. The terms $b_{y \to x}$ and $b_{x \to y}$ capture the migration utility generated in each of these scenarios, respectively. The separability assumption between migration utility and rest of marriage surplus is required for identification in the restricted version of the model. It allows the migration component of surplus to be distinguished from endogamy preferences using information observed solely in the destination country.

Because the outside pool is not observed in the data, the migration utility b cannot be separately identified from the composition of the outside pool. What is identified is their combined effect on the attractiveness of the outside option. Combining the expressions above with Equation (13) yields:

$$B_{y \to x} \equiv b_{y \to x} + \ln[\mu^{\text{out}}_{0y}] = 2 \ln \left[\frac{\mu^{\text{m}}_{y \to x}}{\mu^{\text{in}}_{xy}} \right] + \ln[\mu^{\text{in}}_{0y}], \tag{14}$$

$$B_{x \to y} \equiv b_{x \to y} + \ln[\mu^{\text{out}}_{x0}] = 2 \ln \left[\frac{\mu^{\text{m}}_{x \to y}}{\mu^{\text{in}}_{xy}} \right] + \ln[\mu^{\text{in}}_{x0}] \tag{15}$$

I refer to $B_{y \to x}$ and $B_{x \to y}$ as outside market values. These values are and gender-specific, reflecting potentially asymmetric migration incentives for men and women. These objects summarize the net attractiveness of marrying abroad, combining the utility from migration with the effective size of the outside market. While the two components cannot be separately identified in a single-destination setting, their composite effect is directly identified from observed matching patterns.

In practice, the outside market value B absorbs two additional components: un-

observed traits that differ between migrants and non-migrants, and all costs and frictions of cross-market search and matching. On the trait side, B captures unmeasured religiosity or preferences for traditional family arrangements. On the friction side, it incorporates cultural and information asymmetries, distance, administrative and legal hurdles, and the effort required to search abroad. Kinship ties and migrant networks shift B by lowering search and processing costs or by reducing uncertainty about cross-market partners. These forces do not confound the estimation of endogamy preferences unless network intensity is systematically correlated with inside market intermarriage propensities in a way that alters local matching proportions.

The parameters entering the surplus function, including the endogamy terms and the migration components, are best viewed as *indirect* utilities rather than primitive taste parameters. In cultural–transmission models such as Bisin and Verdier (2001); Bisin et al. (2004), the probability that children adopt their parents' traits depends on group size and the surrounding social environment. That is, the intensity of endogamy captured in the model can be interpreted as reflecting expected returns to fertility and within–group socialization. Moreover, the migration utilities summarize the attractiveness of the foreign marriage market as shaped by its income distribution, religious and ethnic composition, and related local characteristics. The Φ and B parameters therefore represent the equilibrium valuation of different partner types, conditional on the prevailing demographic and economic environment, rather than fixed, exogenous preference primitives. This reduced–form interpretation is consistent with the standard approach in additive random utility matching models (e.g., Choo and Siow (2006)), in which observed sorting patterns identify surplus terms that bundle together preference and environmental determinants of marital choices.

The utility gap between local marriage and marriage abroad for the resident partner can be expressed by the following equations:

$$\Delta U(x,y) = U(x,y,\text{in}) - U(x,y,\text{out}) = \ln\left[\frac{\mu_{xy}^{\text{in}}}{\mu_{y\to x}^{\text{m}}}\right], \tag{16}$$

$$\Delta V(y,x) = V(y,x,\text{in}) - V(y,x,\text{out}) = \ln\left[\frac{\mu_{xy}^{\text{in}}}{\mu_{x\to y}^{\text{m}}}\right] \tag{17}$$

This gap arises from the interaction of both outside market value and endogamy preferences.

It is important to clarify the key assumptions underlying the model. The first is the implicit Independence of Irrelevant Alternatives (IIA) property that arises from

the logit structure. Economically, this implies that exogenous changes in the size of a group do not alter the relative allocation of additional members across marriage migration, intermarriage, and intramarriage: marginal entrants split across these options in fixed proportions determined by payoffs rather than by group size. Likewise, if the option of marriage migration were removed, those who would otherwise have married abroad would reallocate between local intermarriage and intramarriage in the same proportions as existing local matches. This assumption could be problematic if, for example, individuals who select marriage migration would otherwise have disproportionately chosen intragroup marriage, or if endogamy preferences are themselves endogenous to group size. To assess the empirical relevance of this restriction, I re-estimate the main specifications across regions with sharply different Muslim population shares and test for systematic differences in estimated parameters and counterfactual responses. Section 6.3.4 reports that the results are stable across these partitions, suggesting that any violations of IIA are not first-order for the main conclusions. In addition, Section 9.1 relaxes the restriction by adopting a nested-logit structure, where agents first decide between intragroup and intergroup marriage, and conditional on intragroup marriage, choose between forming the match locally or importing a spouse. Finally, to allow for the possibility that endogamy preferences depend on market thickness, Section 9.3 extends the model to incorporate group-size–dependent endogamy preferences and examines the identification and counterfactual implications.

A second assumption is the separability of migration utility from intrinsic match surplus. The gain or cost associated with migration depends only on the direction of relocation and enters additively into the total surplus, without interacting with the observed characteristics of the spouses. This implies that migration utilities apply uniformly across all cross-market matches. If, however, this assumption were violated and migration utilities interacted with partner characteristics, the estimated strength of endogamy preferences would be affected: a positive interaction would make them appear weaker than they are, while a negative interaction would make them appear stronger.

To conduct simulation and counterfactual analyses, it is crucial to compute the equilibrium based on various model parameters. For this purpose, I employ the Iterative Proportional Fitting Procedure (IPFP). The details of this algorithm, as applied to the model in this section, are provided in Appendix D. The primary advantage of the IPFP algorithm is its computational efficiency compared to alternative solv-

ing methods. Appendix E presents simulation results from a simplified illustrative version of the model, which highlights the separate roles of outside market value and endogamy preferences in generating marriage migration.

5 Estimation

The distribution of observed matches across types reflects the relative attractiveness of alternative marriage options in equilibrium. Under the logit-based transferable-utility structure, observed match frequencies correspond to choice probabilities implied by systematic surplus components. By inverting this relationship, the underlying preference parameters can be recovered from the realized matching patterns. This section outlines how I use the observed allocation of matches and the availability of each type to identify preferences and the value of the outside market.

I estimate the model using the entropic regularization approach developed in the optimal transport literature (Cuturi, 2013) and adapted to transferable-utility matching by Galichon and Salanié (2022). This framework delivers substantial computational advantages in settings with many types and guarantees a unique stable matching under the logit assumption on unobserved heterogeneity. In this environment, the equilibrium matching pattern maximizes total systematic surplus plus a generalized entropy term that summarizes the contribution of unobserved taste shocks. As shown in Galichon and Salanié (2022), the equilibrium matching pattern solves:

$$\max_{\boldsymbol{\mu}} \left(\sum_{x,y} \mu_{xy} \Phi^{\boldsymbol{\beta}}_{xy} + \mathcal{E}(\boldsymbol{\mu}, \boldsymbol{n}) \right)$$

where $\boldsymbol{\beta}$ represents preference parameters, and $\mathcal{E}(\boldsymbol{\mu}, \boldsymbol{n})$ is the generalized entropy term. This entropy depends only on the matching pattern $\boldsymbol{\mu}$ and the type counts $\boldsymbol{n} = (\boldsymbol{n}_x, \boldsymbol{n}_y)$. It captures the contribution of unobserved heterogeneity to total surplus, with its precise form determined by the assumed distribution of taste shocks.

The first-order conditions of this problem imply

$$\Phi^{\boldsymbol{\beta}}_{xy} = -\frac{\partial \mathcal{E}(\boldsymbol{\mu}, \boldsymbol{n})}{\partial \mu_{xy}}, \quad \forall x, y,$$

which form a system of equilibrium restrictions linking surplus parameters to observed matching patterns. I estimate the surplus parameters using Minimum Distance Estimation (MDE). Relative to moment-based Poisson approaches (Galichon and Salanié,

2024), MDE delivers a tighter finite-sample fit in settings with many types and sparse match cells, which is a salient feature of marriage market data. Conditional on the observed availabilities $\boldsymbol{n}$, MDE selects the parameter vector that best satisfies the system of first-order conditions:

$$\boldsymbol{D}^{\boldsymbol{\beta}}(\boldsymbol{\mu}, \boldsymbol{n}) \equiv \boldsymbol{\Phi}^{\boldsymbol{\beta}} + \frac{\partial \mathcal{E}(\boldsymbol{\mu}, \boldsymbol{n})}{\partial \boldsymbol{\mu}} = \mathbf{0}$$

The sample analogues of $\boldsymbol{\mu}$ and $\boldsymbol{n}$ are denoted by $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{n}}$. Given an assumed distribution for the unobservables, the parameters $\boldsymbol{\beta}$ can be consistently estimated following Galichon and Salanié (2024).[22]

1. Choose an initial positive definite weighting matrix $\boldsymbol{S}$ and minimize, over $\boldsymbol{\beta} \in \mathbb{R}^d$ (where d is the number of parameters),

$$\left\| \boldsymbol{D}^{\boldsymbol{\beta}}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{n}}) \right\|_{\boldsymbol{S}}^2 = \sum_{x,y,z,t} S_{xy,zt} \left(\Phi_{xy}^{\boldsymbol{\beta}} + \frac{\partial \mathcal{E}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{n}})}{\partial \mu_{xy}} \right) \left(\Phi_{zt}^{\boldsymbol{\beta}} + \frac{\partial \mathcal{E}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{n}})}{\partial \mu_{zt}} \right).$$

   where $\|\mathbf{v}\|_{\mathbf{S}}^2 \equiv \mathbf{v}' \, \mathbf{S} \, \mathbf{v}$ denotes the quadratic form induced by the positive definite weighting matrix $\mathbf{S}$. This yields a consistent estimator $\hat{\boldsymbol{\beta}}$.

2. Apply the delta method to obtain the variance estimator $\hat{\boldsymbol{\Omega}}$ at $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}$, and set $\hat{\boldsymbol{S}} = \hat{\boldsymbol{\Omega}}^{-1}$.

3. Re-estimate $\boldsymbol{\beta}$ by repeating step 1 with $\hat{\boldsymbol{S}}$. The variance-covariance matrix of this estimator is

$$\left(\hat{\boldsymbol{F}}' \hat{\boldsymbol{S}} \hat{\boldsymbol{F}} \right)^{-1},$$

   where $\hat{\boldsymbol{F}}$ is the Jacobian of $\boldsymbol{D}^{\boldsymbol{\beta}}$ with respect to $\boldsymbol{\beta}$, evaluated at the updated estimate $\hat{\boldsymbol{\beta}}$.

To ensure a transparent mapping between surplus parameters and observed match frequencies, I specify the joint surplus as linear in parameters, $\boldsymbol{\Phi}^{\boldsymbol{\beta}} = \boldsymbol{\beta}\boldsymbol{\phi}$. This specification is standard in empirical transferable-utility matching models, delivers point identification under the logit entropy structure, and yields coefficients that are directly interpretable as marginal contributions to systematic surplus. Under this restriction, the estimation problem reduces to a quasi–generalized least squares procedure. Under the logit assumption on unobserved heterogeneity, the generalized

[22] See Galichon and Salanié (2024) for the formal proof.

entropy takes a closed form, implying:

$$\frac{\partial \mathcal{E}(\boldsymbol{\mu}, \boldsymbol{n})}{\partial \mu_{xy}} = -\ln\left[\frac{\mu_{xy}}{\mu_{x0}}\right] - \ln\left[\frac{\mu_{xy}}{\mu_{0y}}\right]$$

This expression links the marginal contribution of each match type to the log odds of matching relative to the option of remaining single on each side of the market. Because the generalized entropy under the logit specification is invariant to the scale of underlying subpopulations, estimation can be carried out using only matching patterns observed within the destination country. In this case, the minimum distance estimator simplifies to a least squares regression of the log odds of observed matches on the surplus shifters:

$$2\ln\left[\hat{\mu}^{\text{in}}_{xy}/\sqrt{\hat{\mu}^{\text{in}}_{x0}\hat{\mu}^{\text{in}}_{0y}}\right] \text{ on } \phi_{xy}$$

where observations are the complete set of combinations of x and y. Outside market values are then recovered directly from Equations (14) and (15). This approach yields a streamlined estimation procedure: the logit structure delivers a closed-form inversion for the systematic surplus, and the model's separability allows preferences and outside market values to be estimated jointly but in a tractable manner.

A practical complication in estimating the model is the presence of type cells with zero observed matches, which arises frequently when the set of observable characteristics is large. In such cases, the partial derivatives of the generalized entropy become unbounded, making the parameter vector $\boldsymbol{\beta}$ infeasible to recover. To address this, I apply a standard small-cell correction and add a strictly positive constant δ to each empirical match count. The adjusted frequencies are constructed as:

$$\tilde{\boldsymbol{\mu}} = \frac{\hat{\boldsymbol{\mu}} + \delta}{N + \delta} N$$

where N is the sample size. This adjustment is a standard regularization device in entropy-based and multinomial-logit settings. It removes numerical singularities without altering the identifying variation. The correction leaves the total population mass unchanged, preserves the asymptotic properties of the estimator, and corrects for finite-sample bias while leaving the economic content of the model intact (Galichon and Salanié, 2024).

6 Results

This section reports three sets of results. I first present the estimated surplus parameters. I then decompose observed marriage migration into endogamy preferences and the outside market value. Finally, I examine heterogeneity by re-estimating the model across age groups, religious affiliations, nativity, and geographic areas.

To implement the model empirically, I specify the set of observable characteristics that enter the systematic surplus and determine access to outside markets. In the UK Census data, I define these characteristics as follows:

1. *Age group*, coded in 10-year intervals:[23]
   - Young: women aged 23–32, men aged 25–33
   - Middle-aged: women aged 33–42, men aged 35–43
   - Old: women aged 43–52, men aged 45–53
2. *Educational level*: a binary indicator for having a college degree or higher. College degree is a significant attribute in the marriage market and is positively associated with match surplus (Chiappori et al., 2017).
3. *Ethnicity*: White British (the majority group), Other White, Indian, Pakistani/Bangladeshi, or Other. The latter four groups are treated as ethnic minorities.
4. *Religion*: A binary variable equal to one for Muslim respondents and zero otherwise.

Each individual is therefore characterized by a combination of age, education, ethnicity, and religion. Ethnicity is particularly important because it determines whether individuals have access to an external partner pool. For the White British majority, marriage migration is essentially absent in the Census, consistent with limited co-ethnic networks abroad, so I treat them as participating only in the domestic marriage market. Minority ethnic groups, by contrast, retain social and familial links to their countries of origin, which makes matching with a co-ethnic partner abroad a relevant option.

6.1 Estimated Preference Parameters

The estimated surplus parameters ($\boldsymbol{\beta}$) describe how observable attributes contribute to the systematic component of the joint surplus relative to the idiosyncratic taste

[23]The cutoff points differ for men and women because women in the UK marry at younger ages than men and have systematically different age-at-marriage distributions.

shocks, which capture unobserved match-specific factors such as emotional affinity or compatibility. Hence, these coefficients measure the marginal substitution between each observable trait and the unobserved component of utility. A larger coefficient indicates that a deficit in that attribute must be compensated by a higher draw of unobserved utility for the match to remain equally attractive.

The primary object of interest in the surplus function is the surplus gains from endogamy. To quantify this, I compare the surplus generated by endogamous matches between types z and z' with the surplus from the corresponding exogamous matches. The relevant measure is the double difference

$$\mathcal{D}_{zz'} \equiv \Phi(x = z, y = z) + \Phi(x = z', y = z') - \Phi(x = z, y = z') - \Phi(x = z', y = z),$$

Here, z and z' denote group labels (such as religion or ethnicity) that apply symmetrically to men and women, so that $x = z$ and $y = z$ represent same-group matches on both sides of the market. A positive $\mathcal{D}_{zz'}$ indicates that the two endogamous matches collectively create more surplus than the two exogamous matches, consistent with a preference for endogamy (or, equivalently, an aversion to exogamy). For example, if z denotes Muslims and z' denotes non-Muslims, $\mathcal{D}_{zz'}$ measures how much more surplus religiously homogamous couples generate compared with mixed Muslim–non-Muslim couples. Because the data contain only realized matches, the model identifies the total surplus of each observed union rather than each partner's preferences. In particular, I cannot separately recover the willingness of Muslims to marry non-Muslims from the willingness of non-Muslims to marry Muslims; the estimated interreligious terms capture a joint outcome. This interpretation is consistent with research that views intermarriage patterns as the result of an equilibrium in which the attitudes and preferences of both groups mutually shape each other (Adida et al., 2014).

Table 5 reports estimated endogamy gains across characteristics. For each characteristic, I summarize the evidence by averaging the type-specific endogamy measures $\mathcal{D}_{zz'}$ across relevant pairs, weighting by the observed match distribution. The coefficients represent the contribution of each observable attribute to the deterministic component of the joint surplus. A positive coefficient indicates that matches with that attribute combination generate higher systematic surplus relative to the omitted category, whereas a negative coefficient indicates lower systematic surplus. These estimates therefore quantify how observed characteristics shift the surplus of realized matches, holding the unobserved match-specific component fixed. Muslims exhibit a larger penalty for large age gaps than non-Muslims. Across characteristics, the sur-

Table 5. Estimated Endogamy Preferences

	$\mathcal{D}_{zz'} = \Phi_{zz} + \Phi_{z'z'} - \Phi_{zz'} - \Phi_{z'z}$	
Groups (z & z')	Non-Muslim	Muslim
Age		
Young & Middle-aged	6.50	7.93
	(0.47)	(0.47)
Young & Old	9.89	9.73
	(0.47)	(0.47)
Middle-aged & Old	5.59	7.49
	(0.47)	(0.47)
Education		
Less than college & College or more	4.04	3.40
	(0.21)	(0.21)
Ethnicity		
Other White & White British	5.85	4.18
	(1.31)	(1.31)
Indian & White British	14.93	13.41
	(1.31)	(1.31)
Pakistani/Bangladeshi & White British	14.05	7.67
	(1.31)	(1.31)
Other & White British	10.11	2.82
	(1.31)	(1.31)
Religion		
Muslim & Non-Muslim	22.67	22.67
	(1.15)	(1.15)

Notes. The table reports estimated endogamy surplus parameters derived from the structural matching model. $\mathcal{D}_{zz'}$ measures the relative surplus of endogamous matches compared to exogamous matches between groups z and z', holding all other characteristics constant. Standard errors are reported in parentheses.

plus gain from religious endogamy is largest on average, followed by ethnic and then educational endogamy. The surplus associated with ethnic endogamy is substantially larger for non-Muslims than for Muslims, implying weaker resistance to ethnic mixing among Muslims. Finally, among all ethnic groups, White British and Other White individuals exhibit the lowest penalties for ethnic exogamy, consistent with relatively small cultural and social distances between these groups.

Utility gaps, defined as the difference between the surplus from marrying locally and from marrying abroad conditional on observable characteristics, are estimated using Equations (16) and (17) and reported in Figure 5. These estimates reflect differences in relative match surplus rather than differences in marriage rates, market thickness, or access to partners. Muslims exhibit substantially lower utility gaps

Figure 5. Utility Gap Between Local Marriage and Marriage Migration by Religion

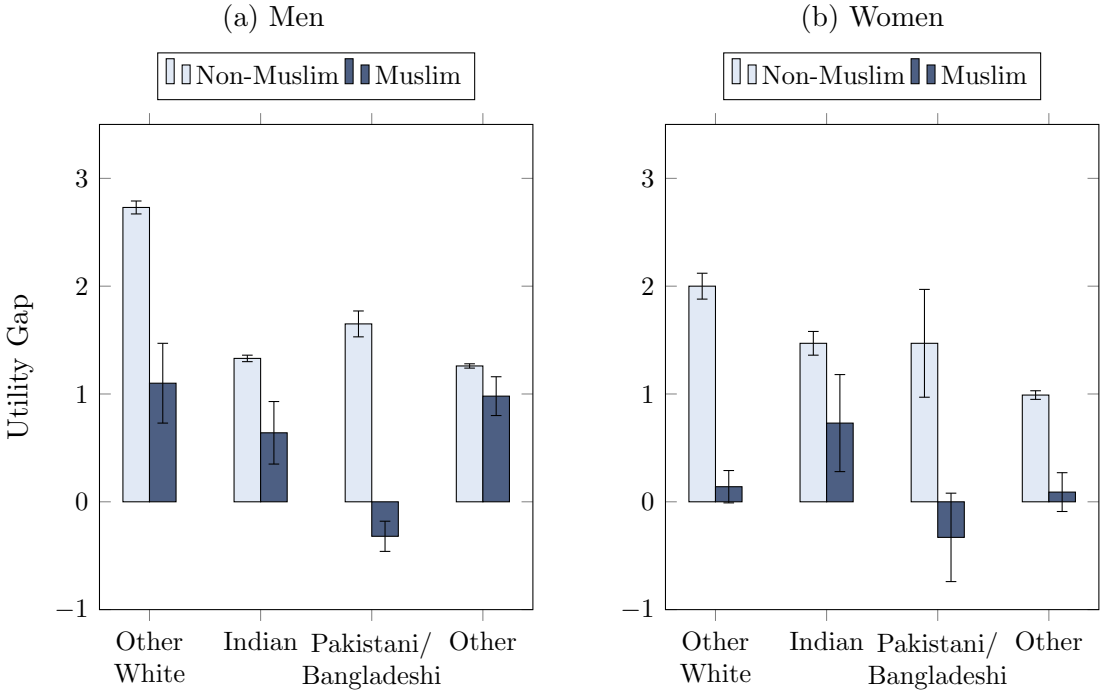

Notes. Utility gap is defined as $U(\text{local}) - U(\text{marriage migration})$, holding other characteristics constant. Error bars show 95% confidence intervals.

than non-Muslims, indicating that, conditional on observables, marriage migration is relatively more attractive for Muslims than for other groups. For Pakistani and Bangladeshi Muslims, the estimated gap is negative, implying that matches formed with someone in the origin country generate higher surplus than local matches in the UK.

These patterns are consistent with qualitative evidence on marriage migration among British South Asian. Ethnographic research documents the central role of extended kin networks among British Pakistanis, through which spouses frequently migrate from Pakistan to join a partner in the UK (Shaw, 2001, 2014). Such kin-based arrangements remain a routine and socially valued route into marriage for second-generation British Pakistanis (Shaw and Charsley, 2006), with similar patterns documented for British Bangladeshis (Gardner, 2006). This qualitative evidence aligns with the negative utility gaps estimated for these groups, suggesting that origin-country matches offer systematically higher surplus than local alternatives, conditional on observed characteristics.

Gender differences in the estimated utility gaps also accord with sociological find-

ings on the organization of transnational marriages. Among British Pakistanis, cross-border marriages often involve British-resident women marrying male relatives from Pakistan within extended kin networks. These arrangements are commonly framed as strategies for managing risk and securing trustworthy partners (Charsley, 2007; Charsley and Ersanilli, 2019). Such marriages may also reshape post-marital living arrangements by reducing the likelihood of extended co-residence with in-laws and facilitating more independent households in the UK (Charsley, 2007). Related work further shows that marriage decisions in these contexts draw on family ties and shared cultural frameworks, including religious practice, expectations around gender roles, and norms governing marital behavior (Phillips et al., 2020).

The estimated utility gaps suggest that policies regulating spousal migration are more likely to bind for groups and genders for whom origin-country matches are particularly valuable. For Pakistani and Bangladeshi Muslims, tighter visa requirements may limit access to the type of match that the model identifies as yielding higher surplus, potentially shifting marriage behavior toward lower-valued local matches or increasing delays to marriage.

It is important to note that differences in utility gaps across ethnic groups are partly driven by variation in educational composition. Figure 6 shows that individuals with lower educational attainment experience larger gains from marrying abroad, reflecting the fact that international matches allow them to pair with more highly educated spouses, as shown in Section 3. Consistent with this mechanism, Muslim Indians, who have a higher tertiary-education rate (around 40 percent) than Muslim Pakistanis and Bangladeshis (approximately 35 percent), exhibit correspondingly lower rates of marriage migration.

6.2 Decomposition of Marriage Migration Determinants

To disentangle the roles of endogamy preferences and outside market value, I construct a counterfactual marriage market equilibrium in which preferences for ethnic and religious homogamy are set to zero. In this counterfactual, all marriage migration arises from outside market value, with no contribution from endogamy preferences. Comparing this counterfactual migration rate with the observed rate yields a two-part decomposition. The counterfactual captures the portion explained by outside market value alone, while the residual is attributed to endogamy preferences.

Figure 7 illustrates the contributions of outside market value and endogamy preferences to marriage migration among ethnic minorities. For Muslims, endogamy

Figure 6. Utility Gap Between Local Marriage and Marriage Migration by Education

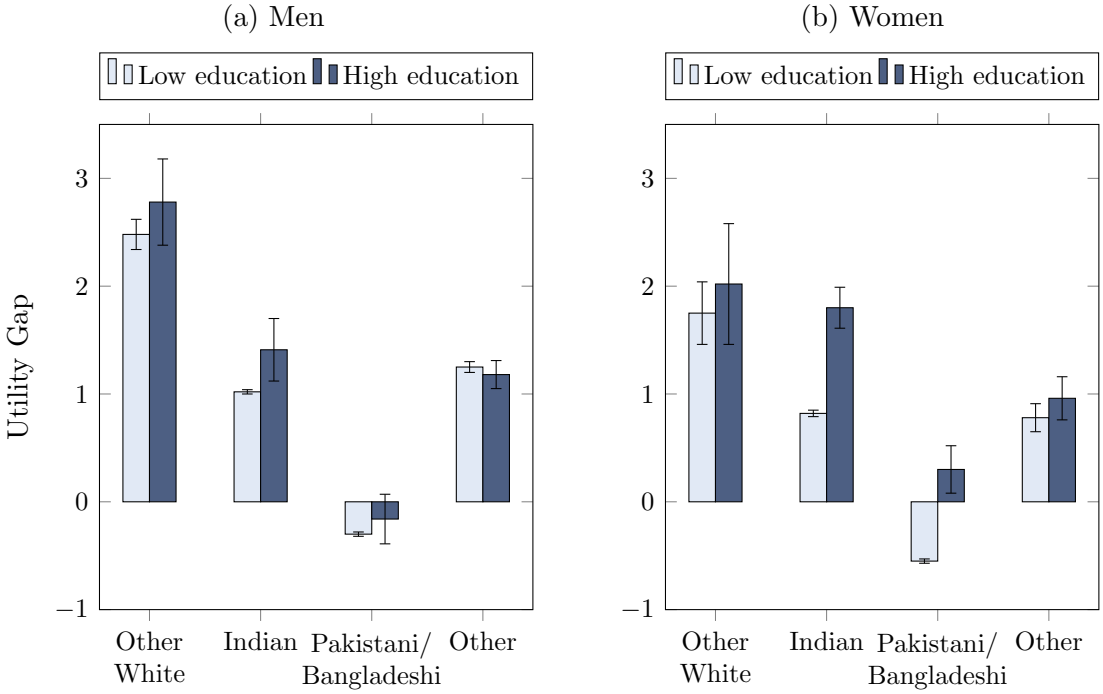

Notes. Utility gap is defined as $U(\text{local}) - U(\text{marriage migration})$, holding other characteristics constant. Error bars show 95% confidence intervals. High education is defined as a college education or more, while low education is defined as less than a college education. Error bars show 95% confidence intervals.

preferences account for the bulk of observed marriage migration, with endogamy preferences accounting for roughly 80% of the observed marriage migration rate and only the 20% remaining share explained by outside market value. In a counterfactual without endogamy preferences, marriage migration rates for both Muslims and non-Muslims fall to about 11%. In addition, the utility gap for Muslims increases to a level comparable to that of non-Muslims (Table A7). In this counterfactual equilibrium, the model implies that approximately 86% of Muslims would engage in interreligious marriages, and the rate of interethnic marriages among ethnic minorities would surge to between 80% and 90%. These findings strongly suggest that the low incidence of intermarriages and the high rate of marriage migration among Muslims are primarily attributed to their preference for endogamy, rather than to gains derived from migration.

A caveat to this decomposition concerns the measurement of marriage migration. As discussed in Section 3, the baseline definition may overstate marriage migration by including individuals who migrated for other reasons and subsequently married a UK

Figure 7. Determinants of Marriage Migration

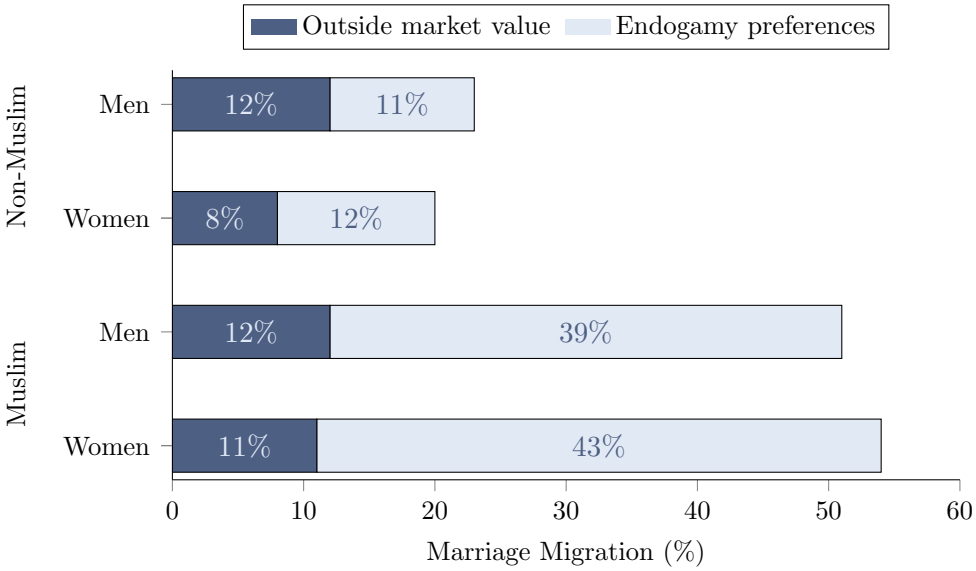

Notes. Each bar shows the contribution of outside market value and endogamy preferences to the observed marriage migration rate, based on the estimated structural model. The sample is limited to ethnic minorities.

resident, thereby overstating the migration-related surplus. To assess the magnitude of the measurement error, I construct a conservative lower-bound measure of marriage migration using the first wave of Understanding Society, classifying a case as marriage migration only when the migrant spouse's reported year of arrival is strictly after the year of marriage. This definition captures unambiguous instances of post-marital migration but necessarily understates true marriage migration by excluding couples who marry and migrate within the same calendar year or shortly thereafter. This alternative measure yields lower marriage migration rates and correspondingly reduces the contribution attributed to outside market value (Appendix I). As a result, the outside market value contribution estimated in this section should be interpreted as an upper bound.

6.3 Heterogeneity Analysis

I examine heterogeneity in marriage migration by re-estimating the model across key subgroups. Specifically, I assess whether the relative roles of endogamy preferences and outside market value vary across cohorts and time, religious affiliations, geographic regions, and nativity status (UK-born versus first-generation).

6.3.1 Heterogeneity by Cohort and Period

In this subsection, I explore how preferences for endogamy vary between different cohorts and periods.

Table 6. Estimated Endogamy Preferences by Age Cohort

	$\mathcal{D}_{zz'} = \Phi_{zz} + \Phi_{z'z'} - \Phi_{zz'} - \Phi_{z'z}$					
	Muslim			Non-Muslim		
Groups (z & z')	Young	Middle-aged	Old	Young	Middle-aged	Old
Education						
Less than college & College or more	5.74	5.93	5.30	5.10	1.94	4.66
	(0.59)	(0.59)	(0.59)	(0.59)	(0.59)	(0.59)
Ethnicity						
Other White & White British	8.03	7.04	7.26	6.35	5.37	5.59
	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)
Indian & White British	18.20	18.50	18.36	16.68	16.97	16.84
	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)
Pak/Bng & White British	17.48	14.07	15.68	11.10	7.70	9.30
	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)
Other & White British	13.07	13.78	12.47	5.78	6.49	5.18
	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)	(3.7)
Religion						
Muslim & Non-Muslim	24.46	23.62	22.45	24.46	23.62	22.45
	(1.43)	(1.43)	(1.43)	(1.43)	(1.43)	(1.43)

Notes. The table reports estimated endogamy surplus parameters derived from the structural matching model. $\mathcal{D}_{zz'}$ measures the relative surplus of endogamous matches compared to exogamous matches between groups z and z', holding all other characteristics constant. Age groups are defined as follows: Young: women aged 23–32 and men aged 25–34; Middle-aged: women aged 33–42 and men aged 35–44; Old: women aged 43–52 and men aged 45–54. Standard errors are reported in parentheses.

I estimate age-specific preference parameters by interacting endogamy indicators with age cohort dummies. The resulting estimates, reported in Table 6, show a clear negative gradient between age and both ethnic and religious endogamy preferences. Younger cohorts display stronger estimated preferences for ethnic and religious endogamy than older cohorts. This pattern can arise through two observationally equivalent mechanisms; either endogamy preferences are genuinely stronger at younger ages, or individuals with stronger preferences tend to marry earlier, leaving a selected group with weaker preferences at older ages. The model identifies only average preferences within age groups and does not capture within-group heterogeneity in religious

attachment.

Table 7. Estimated Endogamy Preferences by Census Year

	$\mathcal{D}_{zz'} = \Phi_{zz} + \Phi_{z'z'} - \Phi_{zz'} - \Phi_{z'z}$			
	2011		2001	
Groups (z & z')	Muslim	Non-Muslim	Muslim	Non-Muslim
Education				
Less than college & College or more	3.21 (0.14)	2.12 (0.15)	4.09 (0.21)	2.81 (0.21)
Religion				
Muslim & Non-Muslim	22.31 (0.76)	22.31 (0.76)	20.9 (1.19)	20.9 (1.19)

Notes. The table reports estimated endogamy surplus parameters derived from the structural matching model. $\mathcal{D}_{zz'}$ measures the relative surplus of endogamous matches compared to exogamous matches between groups z and z', holding all other characteristics constant. Due to data limitations in earlier waves, marriage migration is defined here as marriage between an individual born in the UK and a spouse born outside the UK. Ethnic endogamy parameters are not reported because of limited sample size. Standard errors are reported in parentheses.

To examine changes in endogamy preferences over time, I compare estimates from the 2011 and 2001 Censuses. Because the 2001 Census does not record year of arrival, I redefine marriage migration for this exercise as marriages between UK-born ethnic minorities and partners born abroad. The resulting estimates, reported in Table 7, show that gains from religious endogamy are smaller in 2001 than in 2011, indicating an increase in the value placed on religious endogamy over time. This temporal pattern is consistent with the stronger endogamy preferences observed among younger cohorts in the cross-section. Over the same period, the importance of educational endogamy declines.

A potential concern is that the observed cohort and period gradients reflect selective outmigration of more traditional individuals at older ages rather than changes in preferences. However, evidence on return migration in the UK, while not religion-specific, indicates that outmigration is concentrated in the early years after arrival rather than among long-settled first- or second-generation populations (Dustmann et al., 2011; Dustmann and Görlach, 2016). At the same time, census-based analyses document sustained growth of Muslim and South Asian populations in England and Wales, alongside a rising share of UK-born Muslims (Office for National Statistics, 2012, 2022, 2015; Muslim Council of Britain, 2015, 2024). Combined with the fact

that the analysis is restricted to individuals who were either born in the UK or arrived before age 18, these patterns make selective outmigration an unlikely explanation for the cohort and time differences documented here. The increase in religious endogamy preferences across cohorts and over time may instead reflect broader socio-political dynamics, including rising Islamophobia in the UK and other Western countries (Gould and Klor, 2016; Allen, 2016).

6.3.2 Heterogeneity by Religion

To examine heterogeneity by religion, I re-estimate the model separately for Christians, Hindus, Sikhs, and individuals with no religious affiliation. In each specification, the religion indicator is redefined to correspond to the group under consideration.[24] These groups constitute the largest religious populations in the UK.

Figure 8. Estimated Ethnic and Religious Endogamy Preferences by Religion

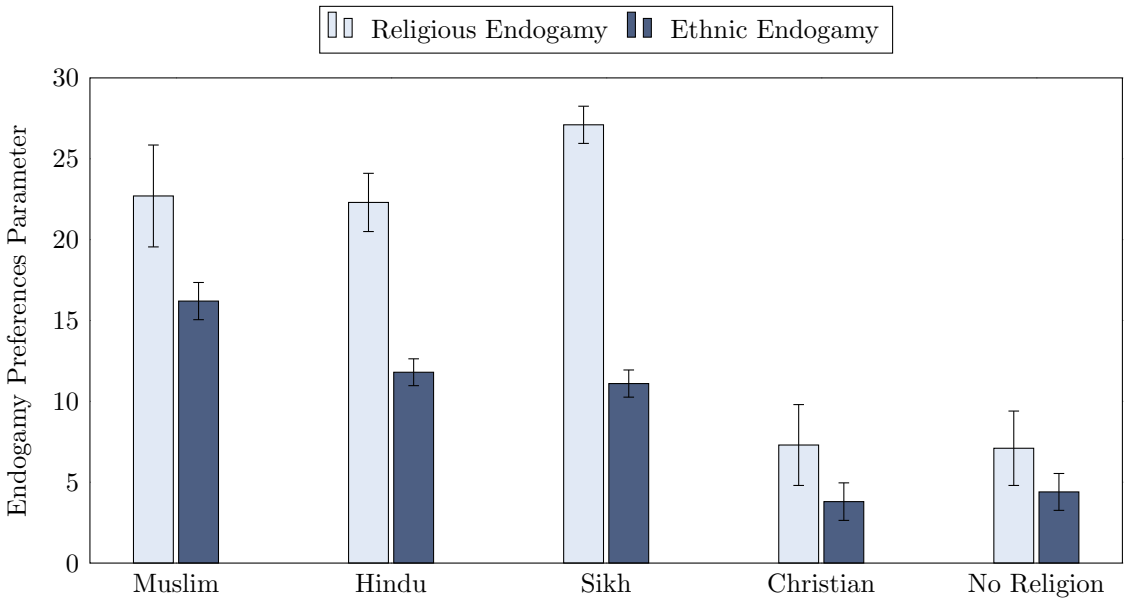

Notes. Error bars denote 95 percent confidence intervals.

Figure 8 reports the estimated endogamy preferences by religious group, with corresponding regression results in Table A8. Religious majorities, namely Christians and individuals with no religious affiliation, exhibit the weakest preferences for re-

[24]Rather than including multiple religious categories simultaneously, I estimate one model per group to preserve statistical power. In the matching framework, introducing multiple religious indicators substantially increases the dimensionality of types and leads to sparse cells, which weakens identification and inflates standard errors.

ligious endogamy. By contrast, religious minorities (Muslims, Hindus, and Sikhs) derive substantially higher surplus from marrying within their religion. This pattern is consistent with models of cultural transmission (Bisin et al., 2004), in which minority groups face stronger incentives to sustain group identity across generations, resulting in higher effective preferences for religiously homogamous marriages relative to majority groups.

Gains from ethnic endogamy exhibit a similar but more muted pattern than those for religious endogamy. Differences between minority and majority groups are smaller, yet religious minorities still derive higher surplus from marrying within ethnicity than religious majorities. This pattern is intuitive given that religious minorities predominantly belong to non-White ethnic groups, whereas most Christians and individuals with no religious affiliation are White and therefore share greater similarity with the White British majority. Among minorities, Muslims exhibit slightly higher gains from intra-ethnic marriage, while Hindus and Sikhs (who are largely drawn from the Indian ethnic group) display comparable preferences for ethnic endogamy.

Figure 9. Determinants of Marriage Migration by Religion

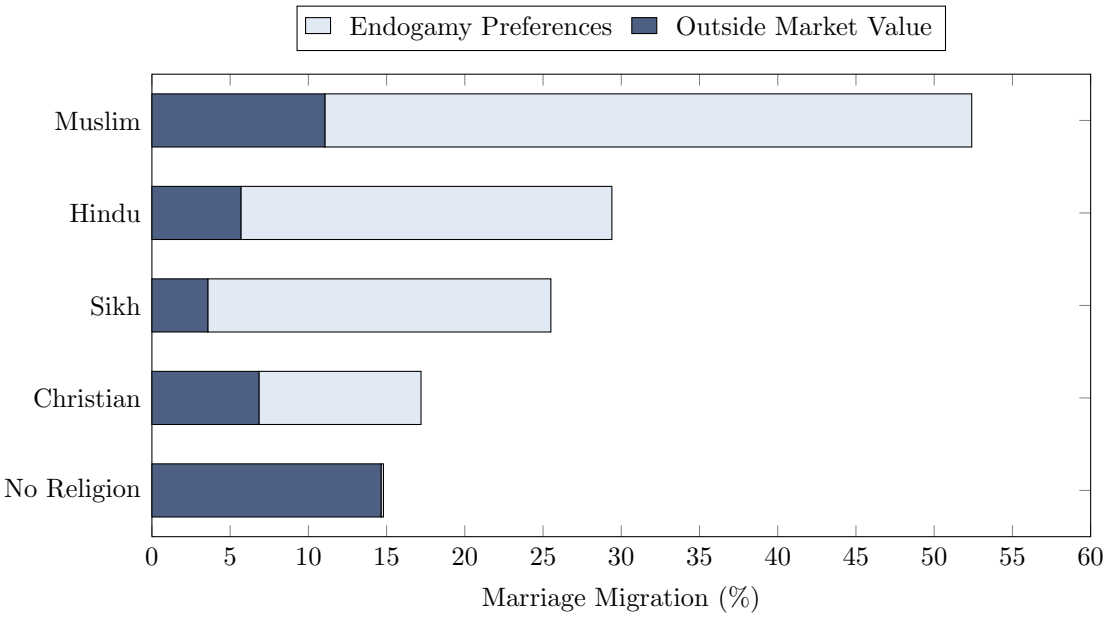

Notes. Each bar shows the contribution of outside market value and endogamy preferences to the observed marriage migration rate, based on the estimated structural model. The sample is limited to ethnic minorities.

Figure 9 shows that the contribution of outside market value to observed marriage migration varies substantially across religious groups. Among individuals with a re-

ligious affiliation, outside market value accounts for only a small share of marriage migration, with endogamy preferences playing the dominant role. By contrast, for individuals with no religious affiliation, outside market value is the primary driver of marriage migration. Importantly, higher marriage migration among Muslims does not mean that they have stronger endogamy preferences than those of other religious minorities. Marriage migration arises from the interaction between endogamy preferences and outside market value; hence, even when preferences for within-group marriage are similar across groups, higher outside market value increases the likelihood that endogamy is realized through marriage migration rather than a local one. As a result, comparable preferences can generate substantially higher marriage migration when the outside market is more attractive.

Figure 10. Marriage Migration by Age for Religious Minorities

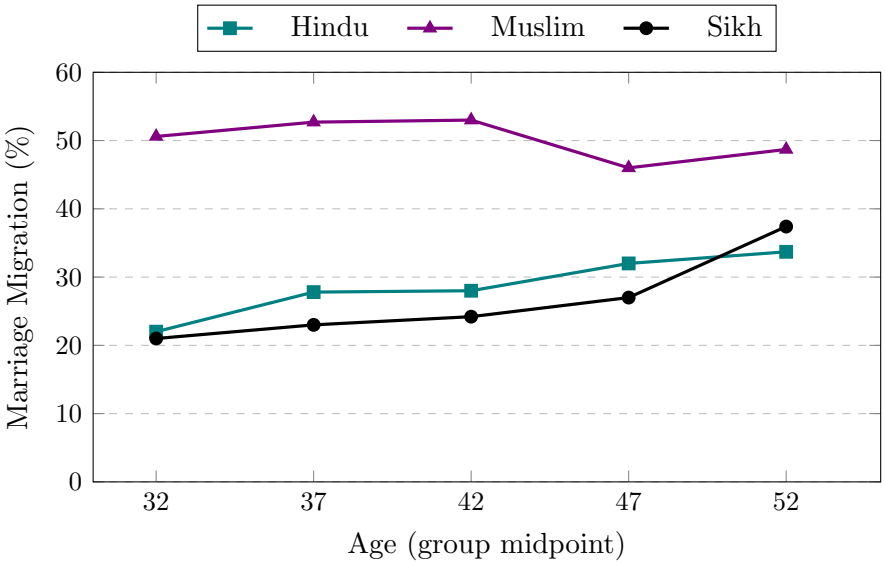

Source. Census for England and Wales, 2011.

Outside market value reflects two components: migration utility and the availability of potential partners in the outside market. The latter depends not only on the size of the origin-country population but also on the strength and persistence of connections to the country of origin. Muslims, as relatively recent immigrants to the UK, have consistently engaged in cross-market marriage with partners from their countries of origin. This repeated pattern sustains strong transnational ties across generations, effectively creating what Goodhart (2013) describes as a first-generation within every generation. This interpretation is consistent with cohort patterns in

marriage migration (Figure 10). Marriage migration rates among Muslims remain relatively stable across age cohorts, whereas for Hindus and Sikhs they decline in younger generations.

6.3.3 Heterogeneity by Nativity Status

In the main analysis, the sample includes two groups: UK-born individuals and those who migrated to the UK before age 18, whom I refer to as UK-raised. A potential concern is that marriage migration behavior and preference parameters may differ across these groups, with UK-born minorities expected to be more assimilated and therefore less likely to rely on marriage migration. The descriptive evidence points in the opposite direction. As shown in Figure 11, marriage migration rates are higher for UK-born Muslims than for UK-raised Muslims across all age groups. This pattern runs counter to a simple assimilation hypothesis but is consistent with earlier findings that younger cohorts exhibit both higher marriage migration rates and stronger religious and ethnic endogamy preferences. It also aligns with qualitative evidence on British Muslims. Theories of segmented assimilation emphasize that children of immigrants may respond to discrimination and exclusion by strengthening ethnic and religious identities and maintaining clear social boundaries in domains such as marriage, even in the presence of substantial structural integration (Portes and Rumbaut, 2001; Zhou, 1997; Rumbaut, 2008).

Figure 11. Muslim Marriage Migration by Nativity Definition

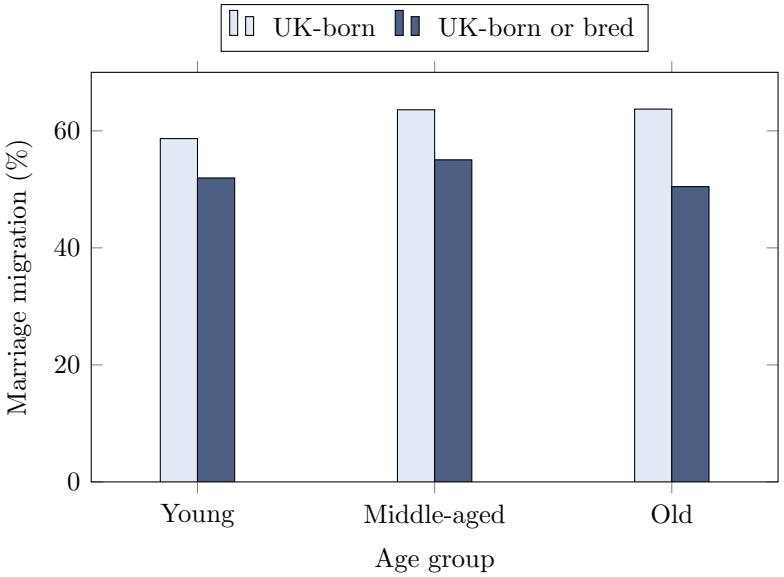

To assess whether preference parameters differ by nativity status, I re-estimate the model using an alternative definition of marriage migration that restricts the resident side to UK-born individuals only, treating all first-generation immigrants, including those who arrived during childhood, as potential foreign spouses. This approach preserves the full set of feasible matches for each resident type, which is required by the structural matching model's feasibility constraints.[25] This redefinition mechanically raises measured marriage migration rates, particularly for Muslims, because marriages between UK-born and UK-raised minorities are now classified as marriage migration.

Figure 12. Determinants of Marriage Migration by Nativity

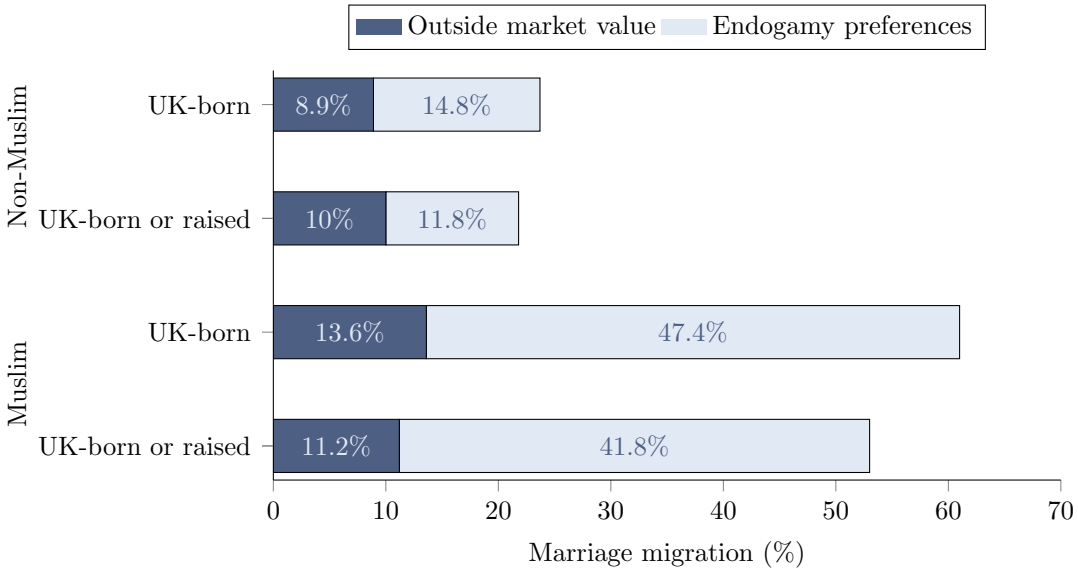

Notes. Each bar shows the contribution of outside market value and endogamy preferences to the observed marriage migration rate, based on the estimated structural model. Each bar shows the contribution of migration gains and endogamy preferences to the overall observed marriage-migration rate for the corresponding group. "UK-born" restricts the resident sample to UK-born minorities; "UK-born or raised" corresponds to the original sample including individuals who arrived before age 18. The underlying values are reported in Table A6. The sample is limited to ethnic minorities.

The estimated parameters (Table A6) show that religious endogamy preferences remain highly stable across specifications, indicating that the core preference patterns are not sensitive to whether UK-raised immigrants are classified as resident or foreign spouses. Figure 12 summarizes the decomposition of marriage migration under the alternative definition and compares it with the baseline specification that includes

[25]Conditioning instead on spouses having arrived after age 18 would truncate the choice set faced by UK-born and UK-raised residents and violate the requirement that all matches and singlehood outcomes exhaust the observed population, leading to biased estimates.

both UK-born and UK-raised residents. For Muslims, endogamy preferences account for the majority of observed marriage migration in both cases. Restricting the resident sample to UK-born individuals increases the overall marriage-migration rate, but the relative importance of endogamy preferences remains dominant. Overall, the figure indicates that early-arrival immigrants are integrated into the UK-born minority marriage market in terms of the underlying preference structure, and that the central mechanism driving marriage migration (strong endogamy preferences among Muslims) is robust to alternative sample definitions.

6.3.4 Heterogeneity by Local Market Population

The main analysis treats the UK as a single marriage market, abstracting from geographic heterogeneity in population composition. In practice, Muslims are unevenly distributed across regions, with some areas having substantially larger Muslim populations than others (Figure A7). A potential concern is that ignoring this heterogeneity could bias the interpretation of marriage migration if local market thinness mechanically pushes individuals toward marriage migration.

Several pieces of evidence suggest that this is not the case. First, regional variation in Muslim population share does not predict marriage migration rates. As shown in Figure 13, areas with larger Muslim populations do not exhibit systematically lower marriage migration, despite offering thicker co-ethnic markets.

Figure 13. Muslim population share and marriage migration across regions

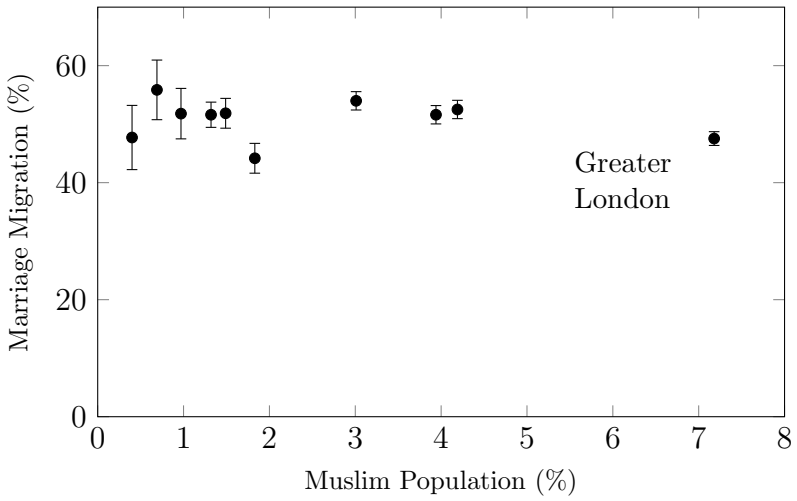

Second, I explicitly restrict the sample to a set of geographically proximate regions with relatively large Muslim populations: the East Midlands, West Midlands, East

of England, Greater London, and South East of England. Focusing on these regions substantially reduces geographic heterogeneity while preserving thick local co-ethnic markets. The resulting estimates are very similar to those obtained in the full sample (Table A3), further indicating that regional variation in market thickness does not drive the main findings.

Finally, evidence on internal mobility for marriage reinforces this results. Using detailed marriage biography data from a Muslim matching platform, I show that Muslim women in particular exhibit substantial willingness to relocate across regions in order to marry, implying that marriage search is not confined to the local area (Appendix H). This internal mobility weakens the link between local population shares and effective partner availability, making the national-market approximation empirically plausible.

Building on this evidence, I examine heterogeneity in marriage migration across local market composition by re-estimating the model separately in areas with low and high shares of Muslim residents. Specifically, I split local marriage markets by whether the Muslim population share lies below or above the median and repeat the decomposition of observed marriage migration into outside market value and endogamy preferences within each stratum. This exercise provides a diagnostic check on the IIA structure. If IIA were strongly violated, changes in local availability or competitive pressure would be expected to reallocate mass between local and international options in a way that depends on market composition.

Figure 14 shows that, for Muslims, overall marriage migration rates are very similar across low- and high-density areas. Within this stable aggregate level, the decomposition changes modestly, with outside market value accounting for a slightly larger share in higher-density markets while endogamy preferences remain quantitatively dominant. Importantly, variation in local Muslim population share does not meaningfully affect whether marriage migration occurs, nor does it generate a qualitative reordering of the underlying mechanisms.

Therefore, these results indicate that marriage migration among Muslims is not primarily driven by local market thinness. They also suggest that potential violations of IIA related to changes in local market composition are unlikely to be first-order for the main conclusions. These results are consistent with the nested-logit extension in Section 9.1, which relaxes IIA directly and yields qualitatively similar findings.

Figure 14. Determinants of Marriage Migration by Muslim Population Density

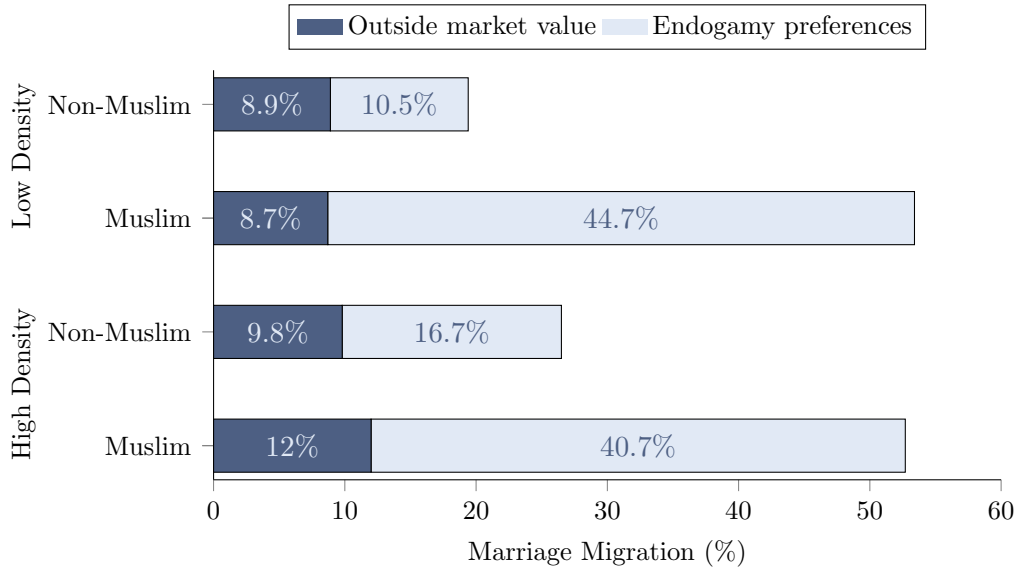

Notes. Each bar shows the contribution of outside market value and endogamy preferences to the observed marriage migration rate, based on the estimated structural model. The sample is limited to ethnic minorities.

6.4 Model Fit

The model's fit is evaluated by comparing predicted and observed matching patterns in Table A4. The estimated parameters closely replicate the empirical distribution of local marriages, marriage migration, and singles across ethnic and religious groups. In particular, the model matches both the high prevalence of intra-group marriage among minorities and the substantial share of marriages involving a spouse who migrates through marriage. This fit indicates that the estimated surplus components replicate the key margins in the data across groups.

Table A5 shows that excluding cross-border matching leads to substantially higher estimated surplus from ethnic endogamy. In a closed-market specification, the model has no margin through which individuals can satisfy endogamy preferences via marriage migration, so intra-ethnic marriage is mechanically attributed to stronger endogamy surplus. As a result, the estimated endogamy surplus conflates preferences with the absence of outside options.

Allowing for cross-border matching separates endogamy preferences from the value of the outside option. Once access to an external partner pool is admitted, some within-group marriage that would be attributed to preferences in a closed model is

instead explained by the ability to realize endogamy through marriage migration. Accounting for this margin is therefore essential for identifying endogamy preferences without upward bias.

7 Robustness Checks

This section briefly summarizes the main robustness checks. Full descriptions, additional results, and supporting figures and tables are reported in Appendix J. Across all exercises, the substantive conclusions remain unchanged.

Including cohabiting couples. The baseline specification treats only formal marriages as matches. As an additional check, I reclassify cohabiting couples as married in the local market. Because spouse visas require formal marriage, this reclassification applies only to local matches. The results show a modest decline in the estimated incidence of religious intermarriage, but the main findings on the determinants of marriage migration remain virtually unchanged (Appendix J.1). This indicates that the results are not sensitive to the treatment of cohabitation.

Refining the definition of marriage migration. The baseline definition classifies a marriage as involving marriage migration when the spouse arrived in the UK after age 18. A potential source of overestimation arises when individuals migrate as adults for education, obtain UK degrees, and subsequently marry someone in the UK. To address this concern, I exclude individuals who arrived after age 18 and subsequently completed a UK university degree (Appendix J.2). This refinement reduces the estimated contribution of outside market value to marriage migration by about two percentage points, but the overall patterns remain consistent. This confirms that the main findings are not driven by misclassification of international students.

Varying the age cutoff. The definition of marriage migration requires setting an age threshold for distinguishing child versus adult migrants. I test alternative thresholds of 16 and 20 instead of 18 (Appendix J.2). The results are stable across thresholds, indicating that the findings are not sensitive to the precise cutoff.

8 Policy Analysis

This section examines two policy-relevant counterfactuals. The first studies how increases in the cost of marriage migration affect marriage market equilibrium outcomes and integration. The second analyzes how an increase in the Muslim population alters marriage migration incentives and equilibrium sorting within the UK marriage market.

Because the data cover only the UK, the analysis adopts a partial-equilibrium perspective with respect to origin countries. I assume that the population residing outside the UK is large relative to the flow of marriage migrants, so that changes in UK marriage migration rates do not affect population composition or marriage market conditions abroad.[26] Under this assumption, policy-induced changes in marriage migration operate through incentives and matching within the UK, while conditions in origin-country marriage markets are held fixed.

These counterfactuals speak to short- and medium-run policy effects on marriage migration and integration in the UK. Longer-run intergenerational impacts are modeled and discussed in Sections 9.2 and 9.3.

8.1 Effect of a Marriage Migration Tax

This section studies the equilibrium effects of policies that increase the cost of marriage migration. These policies, described in detail in Section 2, are widely used across high-income countries and include income thresholds, language requirements, age restrictions, administrative fees, and quantitative limits. Rather than focusing on any single institutional instrument, the analysis models these interventions as an increase in the cost of marrying a spouse from abroad.

Hence, I estimate the equilibrium matching under a counterfactual in which the government introduces a marriage migration tax. This tax directly reduces the utility of migration while leaving endogamy preferences unchanged. From Equations (14) and (15) we have:

$$\tilde{B}_{y\to x} = b_{y\to x} - T + \ln[\mu_{0y}^{\text{out}}] = B_{y\to x} - T$$
$$\tilde{B}_{x\to y} = b_{x\to y} - T + \ln[\mu_{x0}^{\text{out}}] = B_{x\to y} - T$$

[26] Annual out-migration from Pakistan and Bangladesh, which exhibit the highest marriage migration rates in the data, is below 100,000 individuals, a negligible share relative to their population sizes.

where T denotes a lump-sum tax that reduces the marriage surplus for all marriages involving a migrant partner. The tax is subtracted symmetrically from all cross-border matches. $\tilde{B}_{y\to x}$ and $\tilde{B}_{x\to y}$ show outside market value after implementation of the tax. The tax is expressed in the same units as the marriage surplus and therefore has no direct monetary interpretation. This reflects the nature of marriage migration policies, which primarily operate through non-pecuniary channels such as administrative delays, legal uncertainty, and compliance requirements rather than explicit fees. Intuitively, by lowering the surplus from international matches, it reduces the effective attractiveness of the outside option, generating equilibrium effects equivalent to an increase in search frictions or a contraction in the effective size of the outside marriage market.

Figure 15 illustrates how higher marriage migration costs reshape marriage market outcomes along multiple margins. Increases in the migration tax lead to sharp declines in marriage migration for both Muslims and non-Muslims, alongside a rise in singlehood, particularly among Muslims. These aggregate responses reflect the contraction of the outside marriage market and the limited willingness of individuals with strong endogamy preferences to substitute toward local inter-group matches.

Beyond these aggregate effects, higher migration costs also alter sorting within the domestic marriage market. Education-based mixing declines as marrying more highly educated partners from abroad becomes increasingly costly for low-educated individuals, who previously traded access to migration for higher match quality. As these individuals are pushed toward local matches, opportunities for upward educational pairing diminish, leading to lower overall educational mixing.

By contrast, higher migration costs increase mixing along ethnic and religious dimensions, though the magnitudes remain modest. Interreligious marriage among Muslims rises from 3.8% to about 5%, a modest increase that remains far below the 94% rate implied by random matching on religion. The increase in intereth-nic marriage is larger, particularly among Muslims, reflecting weaker preferences for ethnic endogamy relative to religious endogamy. When marriage migration becomes more costly, some individuals are willing to relax ethnic matching constraints, while remaining less willing to sacrifice partner religion.

These results indicate that marriage migration taxes are unlikely to generate substantial gains in integration through intermarriage. As marriage migration becomes more costly, the increase in interreligious marriage among Muslims is limited, while a sizable share instead remains single or compromise on other valuable attributes. This

Figure 15. Effect of Migration Tax on Marriage Migration Equilibrium

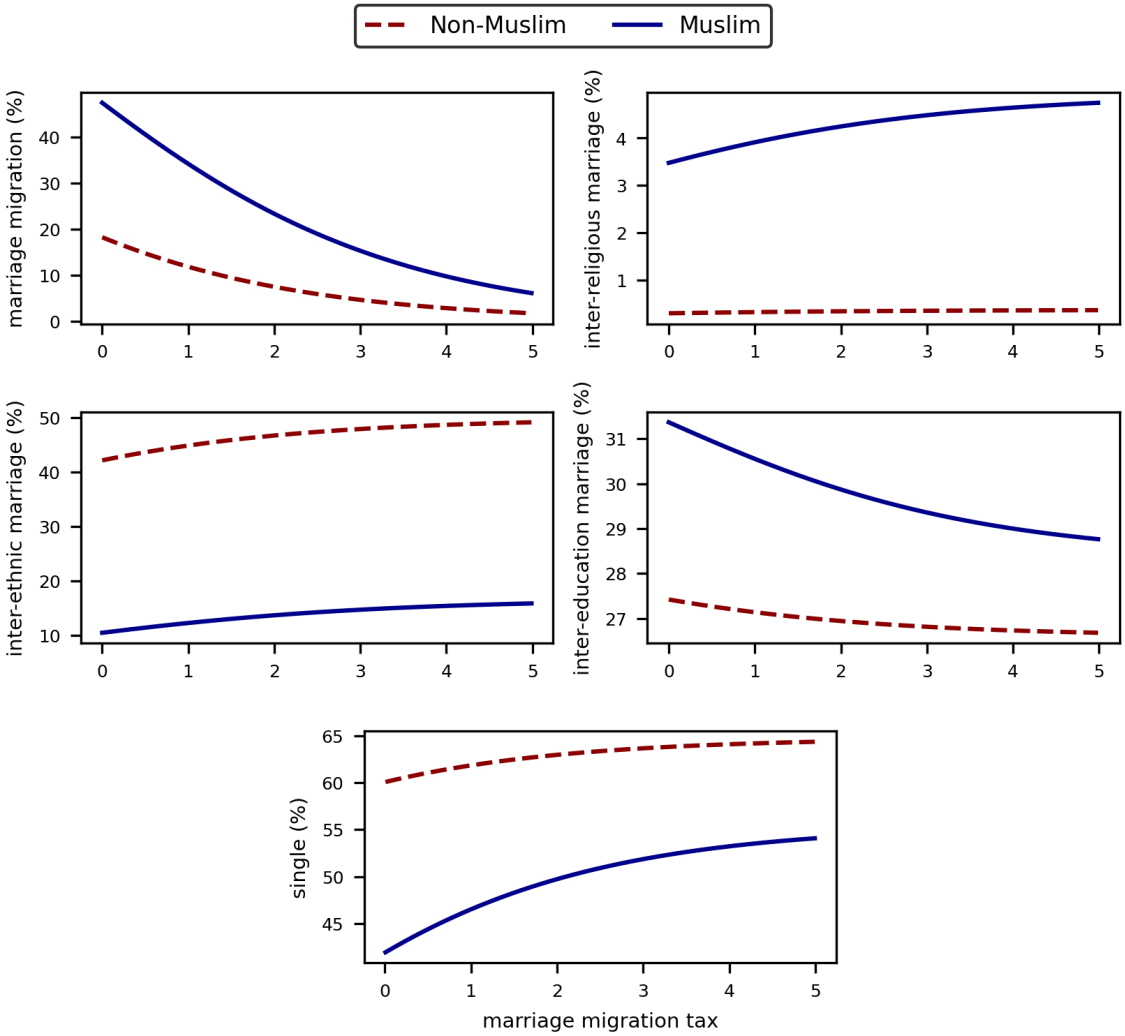

Notes. Figures show the effect of an increase in migration tax on marriage market equilibrium. The horizontal axis reports the level of the lump-sum migration tax, expressed in the same units as the marriage surplus.

reflects the strength of religious endogamy preferences and the limited willingness to substitute away from religious assortativeness. As a result, policies that restrict marriage migration primarily reduce marriage formation rather than meaningfully increasing integration.

The similarity of counterfactual outcomes across the baseline logit and nested logit specifications indicates that the policy results do not hinge on fine substitution patterns implied by IIA. Allowing for correlated unobserved tastes primarily affects within-nest substitution but leaves the aggregate response of marriage, migration,

and singlehood largely unchanged.

The counterfactual exercises should be interpreted as short-run comparative statics within the marriage market. The model abstracts from endogenous skill acquisition, labor-market responses, and long-run socio-economic adaptation. The migration tax therefore isolates the mechanical effect of altering migration utilities in the marriage market, holding all other behavioral margins fixed. In the Section 9.2, I extend the framework to a two-generation setting in which fertility, religious transmission, and group size feed back into the composition of future cohorts; this allows me to show how the long-run impact of the same policy interventions can differ from the short-run effects quantified here.

8.2 Effect of an Increase in Muslim Population

The Muslim population in the UK has experienced substantial growth over the past decade, increasing by 44% between 2011 and 2021 (Office for National Statistics, 2022). This expansion reflects the relatively young median age of Muslims, higher fertility rates compared with other religious groups, and continued immigration from Muslim-majority countries. To assess the implications of this population increase, I conduct a counterfactual analysis that maintains constant preferences and outside market value while adjusting the proportion of Muslims in the UK population.

Figure 16. Effect of Increase in Muslim Population on Marriage Market Equilibrium

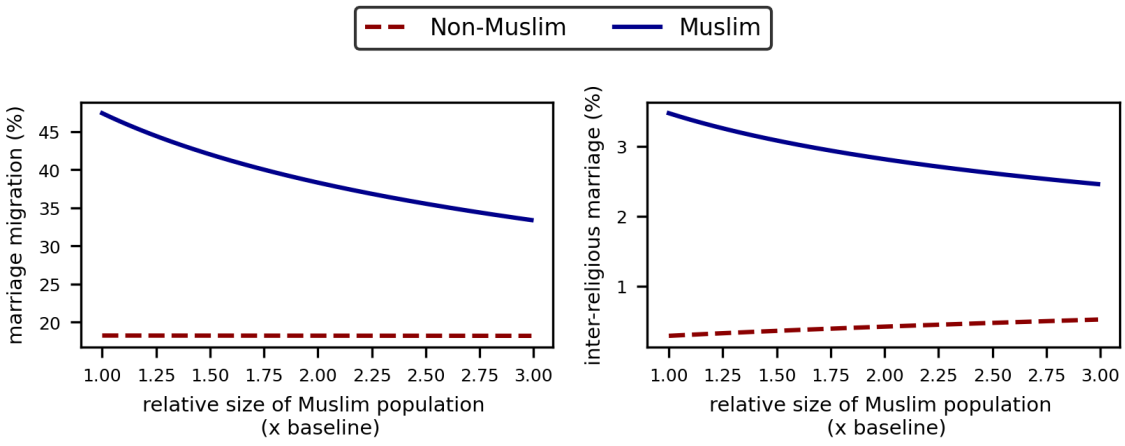

Notes. The horizontal axis reports the Muslim population scaled relative to its baseline level.

Figure 16 shows the effects of an exogenous increase in the Muslim population on marriage migration and interreligious marriage. As the Muslim population grows,

finding a spouse within the same religious group becomes easier within the UK, reducing reliance on marriage migration. If the Muslim population were to double relative to its 2011 level, marriage migration among Muslims declines by around 10 percentage points, corresponding to a reduction of roughly 20 percent relative to baseline levels. At the same time, interreligious marriages fall, as the expanded availability of same-religion partners makes religious endogamy easier to sustain.

The analysis above holds preferences fixed. Allowing endogamy preferences to vary with group size alters the quantitative implications. In Section 9.3, I embed this mechanism in a two-generation model with endogenous endogamy preferences and show that the static comparative statics presented here provide an upper bound on the long-run decline in marriage migration and a lower bound on the long-run reduction in interreligious marriage once demographic feedbacks are taken into account.

9 Model Extensions

The baseline model provides a tractable framework to separate endogamy preferences from the outside market value using cross-sectional data. While deliberately parsimonious, this structure relies on assumptions regarding heterogeneities and preferences. This section extends the framework along several dimensions to assess the robustness of the core findings and analyze the long-run implications of marriage migration and policies. Specifically, the first extension relaxes the Independence of Irrelevant Alternatives assumption by allowing for correlated unobserved tastes, the second introduces intergenerational dynamics and endogenous population composition, and the third allows endogamy preferences to depend on group size.

9.1 Nested Matching: Relaxing the IIA Assumption

In the baseline model, individuals choose among all available options simultaneously, and unobserved utility shocks follow a Type I extreme-value distribution. This structure implies the Independence of Irrelevant Alternatives (IIA) property, under which the relative odds of any two choices are invariant to the presence or attractiveness of other alternatives. IIA implies that if the utility of marrying abroad changes, for example due to a policy shock, individuals who switch from marriage migration to local marriage would reallocate between intermarriage and intramarriage in fixed proportions. This restriction may be empirically strong if domestic and foreign intragroup marriages are closer substitutes than either is with intergroup marriage.

I relax the IIA restriction by modeling the marriage market using a nested logit structure. For simplicity, I allow heterogeneity along a single dimension, group status, distinguishing minority individuals M from majority individuals N. Majority members choose among remaining single, intragroup marriage, and intergroup marriage. Minority members face the same first-stage choice, but conditional on choosing intragroup marriage they make a second-stage decision between marrying a same-group partner locally or importing a partner from abroad. This structure allows shocks to the attractiveness of marriage migration to reallocate individuals between local and cross-border intragroup matches without inducing proportional changes in intermarriage. Substitution patterns within the intragroup nest are governed by dissimilarity parameters $\rho \in (0,1)$ for women and $\delta \in (0,1)$ for men. Close to zero values of ρ and δ imply that domestic and foreign intragroup marriages behave as near-perfect substitutes along unobserved dimensions, while remaining distinct from intergroup marriage options.

Let μ_{ij} denote the share of marriages between individuals from groups i and j, and let μ_{i0} and μ_{0i} denote the shares of single individuals on each side of the market. For minority individuals, I distinguish between intragroup marriages formed locally (μ_{MM}), intragroup marriages formed abroad (μ^x_{MM} for women and μ^y_{MM} for men), and intergroup marriages with majority individuals (μ_{MN} and μ_{NM}). Outside market values for minority women and men, B_x and B_y, are defined analogously to the baseline model.

Within the nested logit framework the systematic utility of each match type must be consistent with observed match shares, taking account of the correlation structure of taste shocks. The resulting equilibrium conditions are:

$$\Phi_{MM} = -\rho \ln\left[\frac{\mu_{MM}}{\mu_{M0}}\right] - (1-\rho)\ln\left[\frac{\mu_{MM}+\mu^x_{MM}}{\mu_{M0}}\right] - \delta\ln\left[\frac{\mu_{MM}}{\mu_{0M}}\right] - (1-\delta)\ln\left[\frac{\mu_{MM}+\mu^y_{MM}}{\mu_{0M}}\right];$$

$$\Phi_{MM} + \underbrace{b_x - \ln[\mu^{\text{out}}_{0M}]}_{B_x} = -\rho\ln\left[\frac{\mu^x_{MM}}{\mu_{M0}}\right] - (1-\rho)\ln\left[\frac{\mu_{MM}+\mu^x_{MM}}{\mu_{M0}}\right] - \ln[\mu^x_{MM}];$$

$$\Phi_{MM} + \underbrace{b_y - \ln[\mu^{\text{out}}_{M0}]}_{B_y} = -\ln[\mu^y_{MM}] - \delta\ln\left[\frac{\mu^y_{MM}}{\mu_{0M}}\right] - (1-\delta)\ln\left[\frac{\mu_{MM}+\mu^y_{MM}}{\mu_{0M}}\right];$$

$$\Phi_{MN} = -\ln\left[\frac{\mu_{MN}}{\mu_{M0}}\right] - \ln\left[\frac{\mu_{MN}}{\mu_{0N}}\right]; \qquad \Phi_{NM} = -\ln\left[\frac{\mu_{NM}}{\mu_{N0}}\right] - \ln\left[\frac{\mu_{NM}}{\mu_{0M}}\right];$$

$$\Phi_{NN} = -\ln\left[\frac{\mu_{NN}}{\mu_{N0}}\right] - \ln\left[\frac{\mu_{NN}}{\mu_{0N}}\right]$$

To ensure that all match shares sum to population totals, we impose the following feasibility constraints:

$$\begin{aligned} n_M &= \mu_{M0} + \mu_{MM} + \mu_{MN} + \mu^x_{MM} \\ m_M &= \mu_{0M} + \mu_{MM} + \mu_{NM} + \mu^y_{MM} \\ n_N &= \mu_{N0} + \mu_{NN} + \mu_{NM} \\ m_N &= \mu_{0N} + \mu_{NN} + \mu_{MN} \end{aligned}$$

I estimate the nested logit model for the case in which the minority group consists of Muslims in the UK and the majority group consists of non-Muslims. Table 8 reports parameter estimates from the baseline and nested logit specifications. Relative to the baseline model, estimated outside market values are slightly lower under the nested structure, while estimated endogamy preferences are modestly higher. This reflects the fact that once domestic and foreign intragroup marriages are allowed to share correlated unobserved components, part of what is attributed to the outside market value in the baseline logit is instead captured by endogamy preferences. The dissimilarity parameters ρ and δ are estimated close to zero, indicating that domestic and foreign intragroup marriages are highly correlated in unobserved tastes and therefore behave as close substitutes.

Table 8. Estimated Parameters of Baseline and Nested Model

	Baseline	Nested
Φ_{MM}	-0.24	1.12
Φ_{MN}	-9.99	-9.99
Φ_{NM}	-7.36	-7.36
Φ_{NN}	0.04	0.04
B_x	9.08	8.35
B_y	8.89	8.10
ρ		0.01
δ		0.16

Table 9 presents a counterfactual in which endogamy preferences are set to zero. In the data, roughly half of Muslim men and women marry a spouse from abroad. In the absence of endogamy preferences, marriage migration falls to around 12 percent for both genders, while intermarriage rates rise above 80 percent. These counterfactual

outcomes are quantitatively very similar under the baseline and nested specifications.

Table 9. Counterfactual without Endogamy Preferences

		Original	Baseline	Nested
Intermarriage rate	Women	2.5	83.4	87.2
	Men	9.0	84.3	88.3
Marriage migration rate	Women	51.9	12.6	12.8
	Men	46.7	11.8	11.7

Notes. The table reports predicted intermarriage and marriage migration rates under a counterfactual in which all endogamy preference parameters are set to zero. "Original" reports observed rates in the data, while "Baseline" and "Nested" report equilibrium outcomes implied by the estimated models.

Taken together, these results indicate that marriage migration among Muslims in the UK is primarily sustained by strong preferences for endogamy rather than by outside market value. Relaxing the IIA assumption allows for more realistic substitution patterns between domestic and foreign intragroup marriages, but it does not alter the central quantitative conclusions.

9.2 Two-Generation Matching: Long-run Effects

The baseline model is static and estimated from a single cross section, and therefore abstracts from how marriage migration reshapes population composition over time. In reality, marriage migration introduces new individuals into the local population, and differences in fertility and religious transmission across marriage types imply that these inflows affect the size and religious composition of subsequent cohorts. As a result, policies that influence marriage migration alter not only contemporaneous marriage-market outcomes but also the structure of future marriage markets. To study these long-run effects, I embed the static matching equilibrium in a two-period, two-generation framework. Estimated surplus parameters and outside market values are held fixed, while marriage and migration choices in Generation 1 determine population stocks and feasible matching outcomes in Generation 2. For tractability, I restrict attention to a single matching dimension, religion, abstracting from age, education, and ethnicity in the intergenerational analysis.

To model intergenerational transmission, I introduce fertility rates and religious transmission probabilities that vary by parental religious affiliation and marriage type. Fertility parameters determine the expected number of children per couple, while

transmission probabilities govern the likelihood that children identify as Muslim. Both sets of parameters are estimated using the Understanding Society longitudinal survey (University of Essex, 2024), drawing on fertility histories and the religious affiliation of parents and their co-residing children.

The two-generation framework is implemented in three steps. First, endogamy preferences, outside market values, and the observed population distribution determine the matching equilibrium in Generation 1, including both marriage patterns and the number of spouses entering the country through marriage migration. Second, each marriage type generates children at a different average rate. I set:

$$f_{MM} = 3.07, \quad f_{MN} = 2.78, \quad f_{NM} = 2.46, \quad f_{NN} = 2.36,$$

where the first index denotes the mother's religion and the second index the father's, so that, for example, f_{MN} is the average number of children born to a Muslim mother married to a non-Muslim father. Fertility is highest among Muslim–Muslim couples and lowest among non-Muslim–non-Muslim couples.

Third, children's religious affiliation is not mechanically inherited from parents and varies systematically across parental religious combinations:

$$p_{MM} = 0.978, \quad p_{MN} = 0.976, \quad p_{NM} = 0.238, \quad p_{NN} = 0.001,$$

where p_{ij} is the probability that a child becomes Muslim given a marriage between a mother of type $i \in \{M, N\}$ and a father of type $j \in \{M, N\}$. Religious transmission is nearly complete among Muslim couples and almost absent among non-Muslim couples, while mixed marriages exhibit strong asymmetry depending on whether the father is Muslim.

Given these parameters, marriages of type i, j in Generation 1 generate the following expected number of Muslim children:

$$f_{ij} \times (\text{number of } (i,j) \text{ couples}) \times p_{ij}.$$

and the number of non-Muslim children is:

$$f_{ij} \times (\text{number of } (i,j) \text{ couples}) \times \big(1 - p_{ij}\big).$$

These expressions determine the size and religious composition of Generation 2 implied by any given Generation 1 matching equilibrium. I assume an equal sex ratio at

birth, so half of the Muslim children become Generation 2 Muslim sons $n_M^{(2)}$ and half become Muslim daughters $m_M^{(2)}$; the same split produces non-Muslim sons $n_N^{(2)}$ and daughters $m_N^{(2)}$. The resulting population stocks $n_M^{(2)}$, $m_M^{(2)}$, $n_N^{(2)}$, $m_N^{(2)}$ replace their Generation 1 counterparts in the second-period matching equilibrium.

Preference parameters are held fixed across generations, so that only population composition evolves over time.[27] Fertility, religious transmission, intermarriage, and marriage migration jointly determine group sizes, implying that Generation 2 outcomes reflect the demographic consequences of Generation 1 behavior rather than changes in underlying preferences.

I use the estimated structural parameters from the single-characteristic version of the model, in which religion is the only matching dimension, and conduct comparative statics to examine how intermarriage and marriage-migration rates respond to changes in key parameters across generations. Restricting attention to religion simplifies the intergenerational mapping and focuses on the role of religious endogamy in shaping long-run outcomes. Figures 17 and 18 report the resulting equilibrium outcomes as endogamy preferences and outside market values are varied.

Figure 17. Sensitivity of Matching Outcomes to Endogamy Preferences

(a) Inter-marriage rate

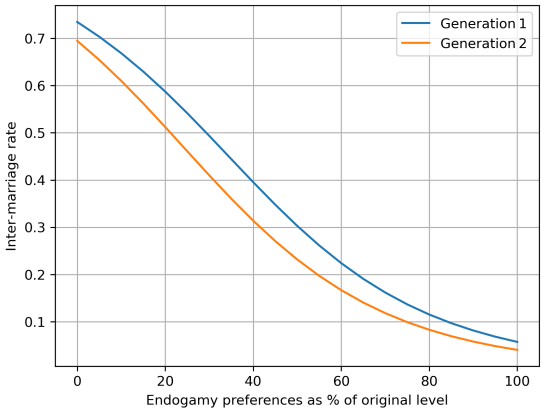

(b) Marriage-migration rate

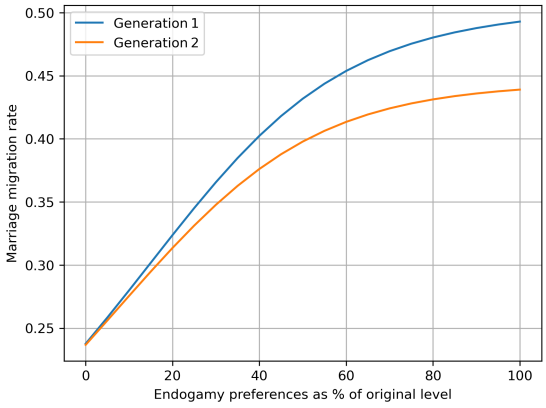

Notes. The figure shows how equilibrium intermarriage and marriage migration rates respond to changes in endogamy preferences. The horizontal axis reports endogamy preferences scaled as a percentage of their baseline level.

For any given level of endogamy preferences, both intermarriage and marriage migration are lower in Generation 2 than in Generation 1. Stronger endogamy reduces

[27] The next section relaxes this assumption by allowing endogamy preferences to depend on group size.

intermarriage in Generation 1 and, combined with higher fertility and near-complete religious transmission among Muslim couples, increases the relative size of the Muslim population in Generation 2. This demographic expansion relaxes same-religion scarcity in the domestic marriage market, reducing reliance on marriage migration. As a result, marriage migration declines mechanically across generations, with the intergenerational gap widening as endogamy preferences strengthen.

Figure 18. Sensitivity of Matching Outcomes to Higher Outside Market Values

(a) Inter-marriage rate

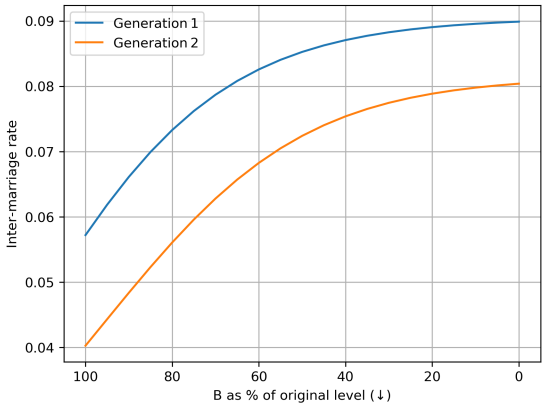

(b) Marriage-migration rate

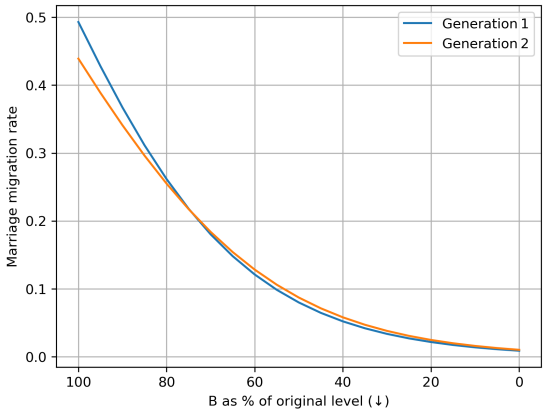

Notes. The figure shows how equilibrium intermarriage and marriage migration rates respond to increases in the outside market values B_x and B_y. The horizontal axis reports outside market values scaled relative to their baseline levels.

For any level of the outside market value, intermarriage rates are substantially lower in Generation 2 than in Generation 1. When intermarriage is low in Generation 1, higher fertility and strong religious transmission among Muslim couples mechanically increase the availability of same-religion partners in Generation 2, further reducing intermarriage even when preferences are unchanged. By contrast, marriage-migration rates are very similar across generations for all values of the outside market value. Changes in B primarily affect incentives to marry abroad within each generation, while population composition plays a limited role in shaping the intergenerational response of marriage migration.

The effects of marriage migration restrictions differ sharply between the short run and the long run. In the short run, a reduction in the outside market value lowers marriage migration and increases intermarriage, as individuals who would otherwise marry abroad substitute toward domestic matches. Over time, however, higher fertility and strong religious transmission among Muslim couples expand the Muslim pop-

ulation, increasing the availability of same-religion partners in the domestic market. This demographic adjustment reduces reliance on marriage migration and attenuates the initial rise in intermarriage. While marriage migration remains persistently lower following the policy change, integration through intermarriage is substantially weaker in the long run than in the short run.

The two-generation framework captures demographic feedback through fertility and religious transmission but abstracts from changes in the outside marriage market itself. In the long run, tighter spouse visa policies may weaken the international connections that sustain the relevant outside marriage market. A persistent reduction in marriage migration can erode kinship networks and information channels that facilitate subsequent matches abroad, shrinking the accessible outside pool and reducing the outside market value beyond the direct effect of the policy. This network erosion would reinforce the demographic forces described above, implying that marriage migration could decline by more than suggested by the counterfactuals in Section 8.

9.3 Two-Generation Matching with Endogenous Preferences

In the benchmark model, endogamy preferences are treated as exogenous and fixed across cohorts. This assumption is standard in one-period matching frameworks with a fixed population distribution and has no implications for identification or equilibrium predictions in a static setting. However, when marriage migration and fertility alter group sizes over time, the assumption becomes more restrictive. Preferences and population composition may evolve jointly, and abstracting from this interaction can understate long-run responses to demographic change.

Theoretical work on cultural transmission and intergenerational preference formation provides a rationale for allowing preferences to respond endogenously to group size (Bisin and Verdier, 2001; Bisin et al., 2004). In these models, parents gain utility from transmitting cultural traits to their children, and the effectiveness of transmission depends on both household composition and the surrounding social environment. When the same-group community is large, cultural transmission is easier and the incentive to marry within the group is weaker. When the group is small, incentives for endogamy are stronger because within-group marriage becomes a more important channel for sustaining cultural identity.

To incorporate this mechanism, I extend the two–generation framework by allowing the strength of religious endogamy to depend on group size in a reduced-form way. Let $q_M^{(g)}$ denote the Muslim population share in generation g. I specify endogamy

preferences as

$$\gamma_M^{(g)} = \gamma_M^0 / q_M^{(g)},$$

where γ_M^0 is the baseline parameter. This inverse relationship captures the core implication of cultural-transmission models (that incentives for within-group marriage are stronger when group size is small) without committing to a specific microfoundation for the adjustment process. Under this formulation, preferences and population composition co-evolve. Higher fertility or marriage migration increase $q_M^{(g)}$, which in turn lowers $\gamma_M^{(g)}$ and weakens incentives for religious homogamy in the subsequent generation. An analogous formulation applies to non-Muslims. This symmetry ensures that differences in outcomes arise from initial group sizes and estimated baseline preferences, rather than from asymmetric modeling assumptions.

Figure 19. Sensitivity of Matching Outcomes to the Muslim Population Share

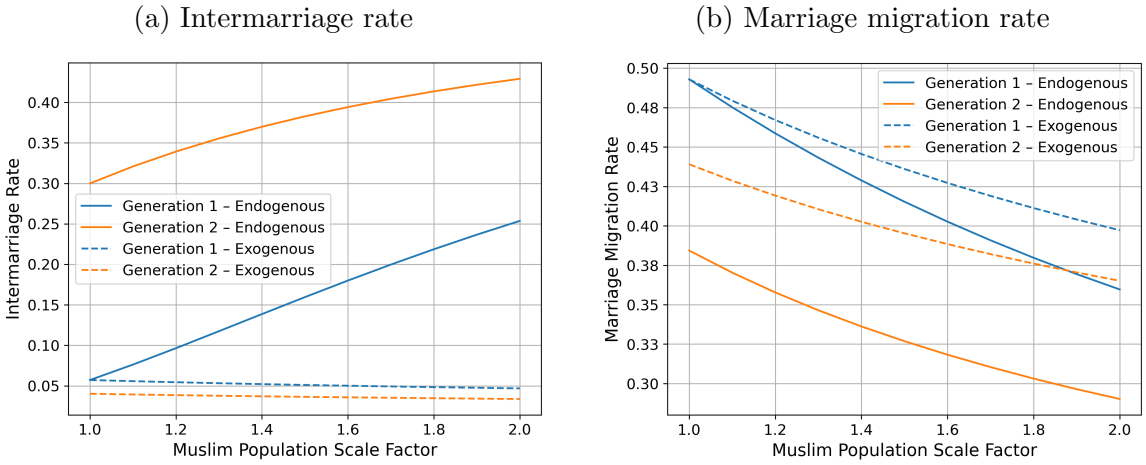

Notes. The horizontal axis reports the Muslim population scaled relative to its baseline level.

Figure 19 illustrates the quantitative implications of endogenous preferences by comparing two specifications. In the benchmark case, endogamy preferences are held fixed across generations. In the alternative specification, preferences adjust endogenously with population composition.

Panel (a) reports intermarriage rates. When preferences are exogenous, increases in the Muslim population share have only a limited effect on intermarriage. By contrast, under endogenous preferences, population growth leads to a pronounced increase in intermarriage, particularly in the second generation. As the group expands, reliance on within-group marriage for cultural transmission weakens, raising the equilibrium incidence of interreligious matches. Therefore, population growth

alone (holding preferences fixed) does not generate assimilation through intermarriage.

Panel (b) reports marriage migration rates. With fixed preferences, marriage migration declines modestly as the domestic marriage market thickens. When preferences respond to group size, the decline is substantially larger and becomes more pronounced across generations. Endogenous preference adjustment amplifies the static effect by further reducing incentives to search for partners in the outside market.

Overall, the figure shows that the exogenous-preference model provides a static benchmark, while the endogenous-preference model introduces a dynamic feedback that alters long-run predictions. An increase in the Muslim population, even in the absence of policy intervention, can raise intermarriage and reduce marriage migration once preference adjustment is taken into account.

10 Conclusion

This study provides the first economic analysis of marriage migration among ethnic minorities in a unified equilibrium framework. It explains why marriage migration rates are substantially higher among Muslims in the UK and clarifies how preferences and outside options jointly shape observed matching patterns. The results also inform the likely effects of policies that restrict spousal migration.

The analysis shows that the high incidence of marriage migration among Muslims is driven primarily by strong preferences for religious and ethnic endogamy rather than by economic migration incentives alone. Importantly, these endogamy preferences are not unique to Muslims. Estimated gains from within-group marriage are of similar magnitude across religious minorities, including Hindus and Sikhs. What distinguishes Muslims is their higher outside market value. When strong endogamy preferences interact with a more attractive external partner pool, a larger share of within-group marriages is realized through international rather than domestic matches, generating substantially higher marriage migration rates.

The elevated outside market value for Muslims is largely attributable to the greater availability of potential partners among Pakistani and Bangladeshi Muslims and to the persistence of transnational social and kinship networks linking the UK to countries of origin. Stronger connections to origin countries lower the effective cost of marrying abroad and expand the set of feasible matches. In turn, repeated cross-market marriages reinforce these transnational ties, creating a self-reinforcing mechanism

through which high marriage migration rates can persist across generations.

This paper introduces a novel approach to studying marriage migration, providing new insights despite data limitations. The estimation relies on information from a single receiving market, which constrains the analysis of alternative opportunities available to potential migrants. Future work could extend the approach to settings with data from both sending and receiving countries, or to environments with multiple internal marriage markets, such as regions within a country. More broadly, the framework could be applied to other forms of migration involving joint decisions, including migration for education or family reunification beyond marriage.

References

Adda, J., Pinotti, P., and Tura, G. (2025). There's more to marriage than love: the effect of legal status and cultural distance on intermarriages and separations. *Journal of Political Economy*, 133(4):1276–1333.

Adida, C. L., Laitin, D. D., and Valfort, M.-A. (2014). Muslims in france: identifying a discriminatory equilibrium. *Journal of Population Economics*, 27(4):1039–1086.

Adida, C. L., Laitin, D. D., and Valfort, M.-A. (2016). *Why Muslim integration fails in Christian-heritage societies*. Harvard University Press.

Ahn, S. Y. et al. (2020). Matching across markets: An economic analysis of cross-border marriage. *University of Illinois mimeo*.

Ahn, S. Y. and Lubotsky, D. (2024). Assimilation of marriage migrants and the role of language: Evidence from south korea.

Allen, C. (2016). *Islamophobia*. Routledge.

Anderberg, D. and Vickery, A. (2021). The role of own-group density and local social norms for ethnic marital sorting: Evidence from the UK. *European Economic Review*, 138:103774.

Australian Department of Home Affairs (2023). Partner visa program report, programme year 2022–23. `https://www.homeaffairs.gov.au/research-and-statistics/statistics/visa-statistics`. Accessed 8 July 2025.

Banerjee, A., Duflo, E., Ghatak, M., and Lafortune, J. (2013). Marry for what? caste and mate selection in modern india. *American Economic Journal: Microeconomics*, 5(2):33–72.

Beauchamp, A., Calvi, R., and Fulford, S. (2021). Terms of engagement: Migration, dowry, and love in indian marriages.

Bech-Sørensen, J. and Pollet, T. V. (2016). Sex differences in mate preferences: A replication study, 20 years later. *Evolutionary Psychological Science*, 2:171–176.

Beck-Gernsheim, E. (2007). Transnational lives, transnational marriages: a review of the evidence from migrant communities in europe. *Global networks*, 7(3):271–288.

Becker, G. S. (1973). A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846.

Becker, G. S. (1974). A theory of marriage: Part ii. *Journal of political Economy*, 82(2, Part 2):S11–S26.

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

Bisin, A., Patacchini, E., Verdier, T., and Zenou, Y. (2008). Are muslim immigrants different in terms of cultural integration? *Journal of the European Economic Association*, 6(2-3):445–456.

Bisin, A., Topa, G., and Verdier, T. (2004). Religious intermarriage and socialization in the united states. *Journal of political Economy*, 112(3):615–664.

Bisin, A. and Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic theory*, 97(2):298–319.

Bonjour, S. and De Hart, B. (2013). A proper wife, a proper marriage: Constructions of 'us' and 'them'in dutch family migration policy. *European Journal of Women's Studies*, 20(1):61–76.

Bonjour, S. and Kraler, A. (2015). Introduction: Family migration as an integration issue? policy perspectives and academic insights. *Journal of Family Issues*, 36(11):1407–1432.

Buss, D. M. (1989). Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures. *Behavioral and brain sciences*, 12(1):1–14.

Casey, L. (2016). The casey review: A review into opportunity and integration.

Charsley, K. (2005). Unhappy husbands: Masculinity and migration in transnational pakistani marriages. *Journal of the Royal Anthropological Institute*, 11(1):85–105.

Charsley, K. (2007). Risk, trust, gender and transnational cousin marriage among british pakistanis. *Ethnic and Racial Studies*, 30(6):1117–1131.

Charsley, K. (2013). Transnational marriage. In *Transnational marriage*, pages 12–31. Routledge.

Charsley, K. (2018). Closeness and distance in pakistani transnational cousin marriage. *Families, Relationships and Societies*, 7(3):521–526.

Charsley, K., Bolognani, M., Ersanilli, E., and Spencer, S. (2020). *Marriage migration and integration*. Springer.

Charsley, K., Bolognani, M., and Spencer, S. (2017). Marriage migration and integration: Interrogating assumptions in academic and policy debates. *Ethnicities*, 17(4):469–490.

Charsley, K. and Ersanilli, E. (2019). The 'mangetar trap'? work, family and pakistani migrant husbands. *Norma*, 14(2):128–145.

Charsley, K. and Liversage, A. (2015). Silenced husbands: Muslim marriage migration and masculinity. *Men and Masculinities*, 18(4):489–508.

Charsley, K., Storer-Church, B., Benson, M., and Van Hear, N. (2012). Marriage-related migration to the UK. *International Migration Review*, 46(4):861–890.

Chiappori, P.-A. (2017). *Matching with transfers*. Princeton University Press.

Chiappori, P.-A., Fiorio, C. V., Galichon, A., and Verzillo, S. (2022). Assortative matching on income. Technical report, JRC Working Papers in Economics and Finance.

Chiappori, P.-A., Nguyen, D. L., and Salanié, B. (2019). Matching with random components: simulations. *Work. Pap., Columbia Univ., New York Google Scholar Article Location*, 5.

Chiappori, P.-A., Oreffice, S., and Quintana-Domeque, C. (2018). Bidimensional matching with heterogeneous preferences: education and smoking in the marriage market. *Journal of the European Economic Association*, 16(1):161–198.

Chiappori, P.-A., Salanié, B., and Weiss, Y. (2017). Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–67.

Choo, E. and Siow, A. (2006). Who marries whom and why. *Journal of political Economy*, 114(1):175–201.

Ciscato, E., Galichon, A., and Goussé, M. (2020). Like attract like? a structural comparison of homogamy across same-sex and different-sex households. *Journal of Political Economy*, 128(2):740–781.

Constable, N. (2010). *Cross-border marriages: Gender and mobility in transnational Asia*. University of Pennsylvania Press.

Corblet, P. (2022). Education expansion, sorting, and the decreasing education wage premium. *Job Market Paper*.

Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.

Dale, A. (2008). Migration, marriage and employment amongst indian, pakistani and bangladeshi residents in the uk. *University of Manchester, CCSR Working Paper*, 2:1–10.

Dupuy, A. (2021). Migration in china: To work or to wed? *Journal of Applied Econometrics*, 36(4):393–415.

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

Dustmann, C. and Görlach, J.-S. (2016). The economics of temporary migrations. *Journal of Economic Literature*, 54(1):98–136.

Eika, L., Mogstad, M., and Zafar, B. (2019). Educational assortative mating and household income inequality. *Journal of Political Economy*, 127(6):2795–2835.

Galichon, A. and Salanié, B. (2017). IIA in separable matching markets.

Galichon, A. and Salanié, B. (2022). Cupid's invisible hand: Social surplus and identification in matching models. *The Review of Economic Studies*, 89(5):2600–2629.

Galichon, A. and Salanié, B. (2024). Estimating separable matching models. *Journal of Applied Econometrics*, 39(6):1021–1044.

Gardner, K. (2006). The transnational work of kinship and caring: Bengali–british marriages in historical perspective. *Global Networks*, 6(4):373–387.

Georgiadis, A. and Manning, A. (2011). Change and continuity among minority communities in britain. *Journal of Population Economics*, 24(2):541–568.

Glenn, E. N. and Glenn, E. (2010). *Issei, Nisei, war bride: Three generations of Japanese American women in domestic service*. Temple University Press.

Goodhart, D. (2013). *The British dream: Successes and failures of post-war immigration*. Atlantic Books Ltd.

Gould, E. D. and Klor, E. F. (2016). The long-run effect of 9/11: Terrorism, backlash, and the assimilation of muslim immigrants in the west. *The Economic Journal*, 126(597):2064–2114.

Gualdani, C. and Sinha, S. (2023). Partial identification in matching models for the marriage market. *Journal of Political Economy*, 131(5):1109–1171.

Hall, T., Manning, A., and Sumption, M. (2023). Projecting UK net migration.

Hitsch, G. J., Hortaçsu, A., and Ariely, D. (2010). What makes you click?—mate preferences in online dating. *Quantitative marketing and Economics*, 8:393–427.

Huschek, D., De Valk, H. A., and Liefbroer, A. C. (2012). Partner choice patterns among the descendants of turkish immigrants in europe. *European Journal of Population= Revue Europeenne De Demographie*, 28(3):241.

Jacquet, A. and Montpetit, S. (2022). Veiling and economic integration of muslim women in france.

Kawaguchi, D. and Lee, S. (2017). Brides for sale: Cross-border marriages and female immigration. *Economic Inquiry*, 55(2):633–654.

Kenrick, D. T. and Keefe, R. C. (1992). Age preferences in mates reflect sex differences in human reproductive strategies. *Behavioral and Brain Sciences*, 15(1):75–91.

Kibria, N. (2012). Transnational marriage and the bangladeshi muslim diaspora in britain and the united states. *Culture and Religion*, 13(2):227–240.

Kofman, E. (2004). Family-related migration: a critial review of european studies. *Journal of Ethnic and Migration studies*, 30(2):243–262.

Labour Force Survey (2011). http://doi.org/10.5255/UKDA-Series-2000026.

Lee, H.-K. (2008). International marriage and the state in south korea: Focusing on governmental policy. *Citizenship studies*, 12(1):107–123.

Lievens, J. (1999). Family-forming migration from turkey and morocco to belgium: The demand for marriage partners from the countries of origin. *International migration review*, 33(3):717–744.

Lindquist, J., Xiang, B., and Yeoh, B. S. (2012). Opening the black box of migration: Brokers, the organization of transnational mobility and the changing political economy in asia. *Pacific affairs*, 85(1):7–19.

Manning, A. and Roy, S. (2010). Culture clash or culture club? national identity in britain.

Muslim Council of Britain (2015). British muslims in numbers: A demographic, socio-economic and health profile of muslims in britain. Muslim Council of Britain.

Muslim Council of Britain (2024). Census 2021 – striving for fairness. Muslim Council of Britain. Analysis of Census 2021 data on Muslims in England and Wales.

OECD (2017). International migration outlook 2017. http://dx.doi.org/10.1787/migr_outlook-2017-en.

Office for National Statistics (2011). 2011 census for england and wales. https://www.ons.gov.uk/census/2011census/2011censusdata/censusmicrodata/securemicrodata.

Office for National Statistics (2012). Religion in england and wales 2011. *Statistical Bulletin*, Office for National Statistics. Released 11 December 2012.

Office for National Statistics (2015). 2011 census analysis: Ethnicity and religion of the non-UK born population in england and wales. Office for National Statistics. Published 18 June 2015.

Office for National Statistics (2022). Religion, england and wales: Census 2021. *Statistical Bulletin*, Office for National Statistics. Released 29 November 2022.

Office for National Statistics (2024). Long-term international migration, table ipsz4: Main reason for migration. `https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/internationalmigration`. Accessed 8 July 2025.

Phillips, R., Chambers, C., Ali, N., Pande, R., and Hopkins, P. (2020). Mobilizing pakistani heritage, approaching marriage. *Ethnic and Racial Studies*, 43(16):1–19.

Portes, A. and Rumbaut, R. G. (2001). *Legacies: The story of the immigrant second generation*. Univ of California Press.

Rumbaut, R. G. (2008). Reaping what you sow: Immigration, youth, and reactive ethnicity. *Applied development science*, 12(2):108–111.

Shaw, A. (2001). Kinship, cultural preference and immigration: consanguineous marriage among british pakistanis. *Journal of the Royal Anthropological Institute*, 7(2):315–334.

Shaw, A. (2014). Drivers of cousin marriage among british pakistanis. *Human Heredity*, 77(1-4):26–36.

Shaw, A. and Charsley, K. (2006). Rishtas: adding emotion to strategy in understanding british pakistani transnational marriages. *Global networks*, 6(4):405–421.

Siow, A. (2015). Testing becker's theory of positive assortative matching. *Journal of Labor Economics*, 33(2):409–441.

Sumption, M. and Vargas-Silva, C. (2019). Love is not all you need: Income requirement for visa sponsorship of foreign family members. *Journal of Economics, Race, and Policy*, 2:62–76.

University of Essex, I. f. S. a. E. R. (2024). Understanding society: Waves 1-14, 2009-2023 and harmonised bhps: Waves 1-18, 1991-2009. DOI: http://doi.org/10.5255/UKDA-SN-6614-20. 19th Edition.

Wang, H.-z. and Chang, S.-m. (2002). The commodification of international marriages: Cross-border marriage business in taiwan and viet nam. *International migration*, 40(6):93–116.

Weiss, Y., Yi, J., and Zhang, J. (2018). Cross-border marriage costs and marriage behavior: Theory and evidence. *International Economic Review*, 59(2):757–784.

Wray, H., Agoston, A., and Hutton, J. (2014). A family resemblance? the regulation of marriage migration in europe. *European Journal of Migration and Law*, 16(2):209–247.

Yang, W.-S. and Lu, M. C.-W. (2010). *Asian cross-border marriage migration: Demographic patterns and social issues*, volume 2. Amsterdam University Press.

Zappi, S. (2003). French government revives assimilation policy. *Migration Policy Institute*, 1:2003.

Zhou, M. (1997). Segmented assimilation: Issues, controversies, and recent research on the new second generation. *International migration review*, 31(4):975–1008.

A Appendix Figures and Tables

A.1 Figures

Figure A1. Marriage Migration and Intra-ethnic Marriages

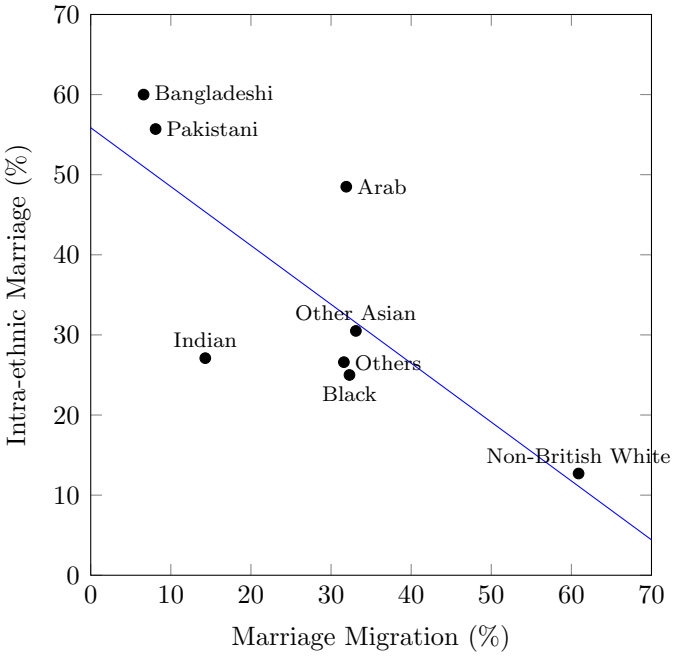

Notes. Marriage migration is defined as cases in which the spouse arrived in the UK at age 18 or older. Intra-ethnic marriage refers to marriages where both spouses report the same ethnic group. *Source.* Census for England and Wales, 2011.

Figure A2. Share of Partnered Men in Cohabiting Unions, by Age and Religion

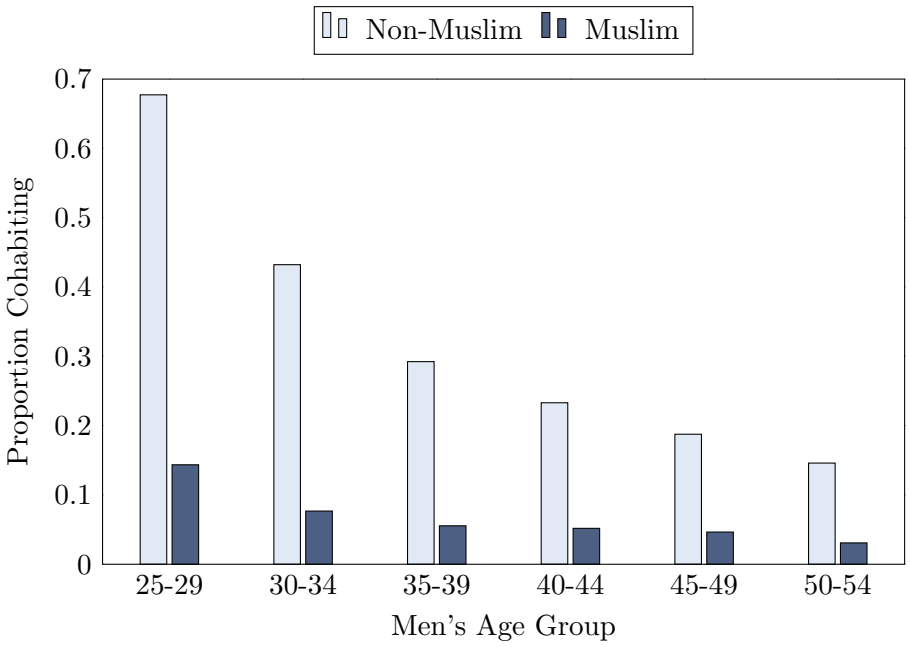

Notes. The figure reports the share of partnered men in each age group who are cohabiting rather than legally married. Religious affiliation is defined by the man. *Source*: Census for England and Wales, 2011

Figure A3. Share Actively Practicing Individuals, by Religion and Nativity

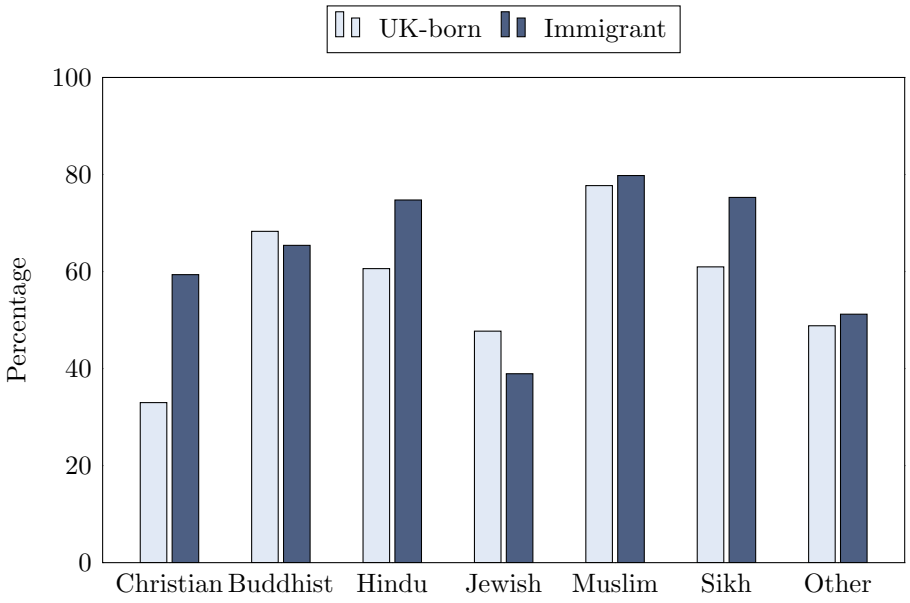

Notes. The figure reports the percentage of individuals within each religious group who answer 'yes' to the question 'Do you consider that you are actively practising your religion?'. *Source.* Citizenship Survey, 2010-2011

Figure A4. Inter-religious and Inter-ethnic Marriage Rates, by Ethnicity and Religion

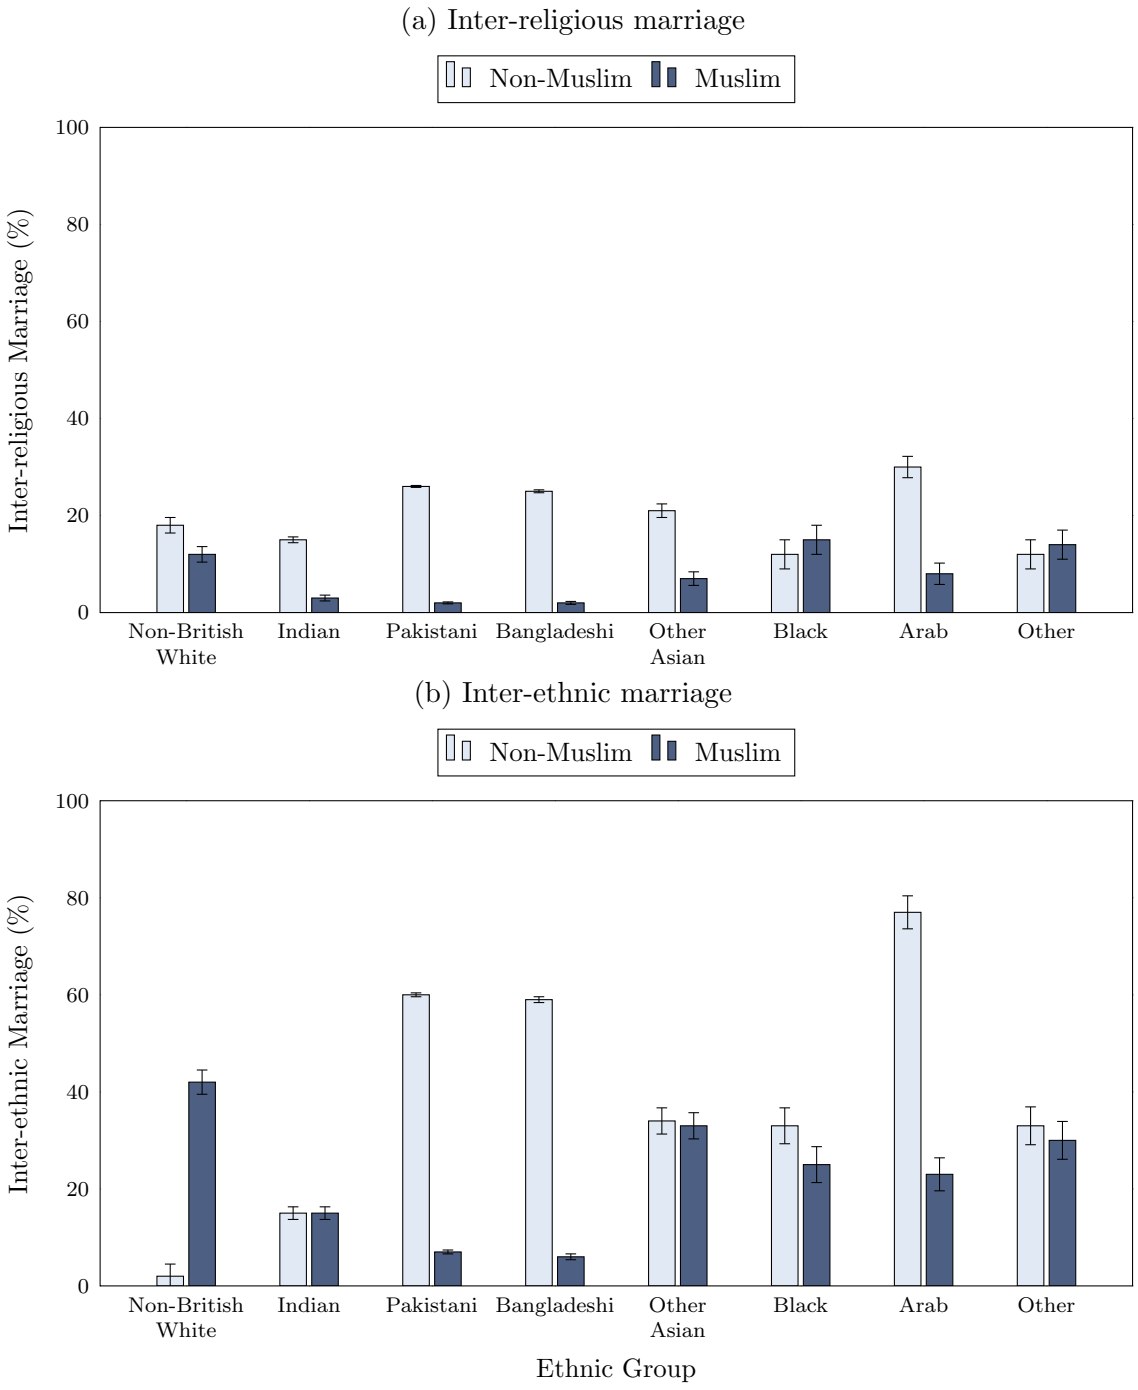

Notes. Error bars denote 95 percent confidence intervals. *Source.* Census for England and Wales, 2011

Figure A5. Baseline Marginal Effects of Marriage Migration on Spouse's Education

(a) Muslim men

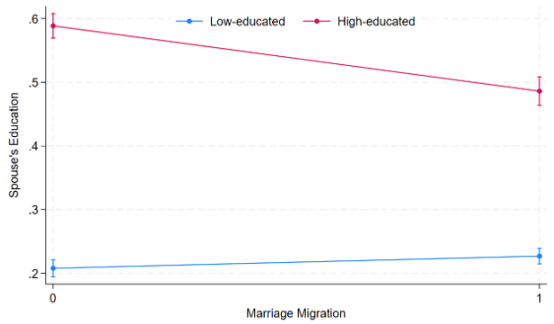

(b) Muslim women

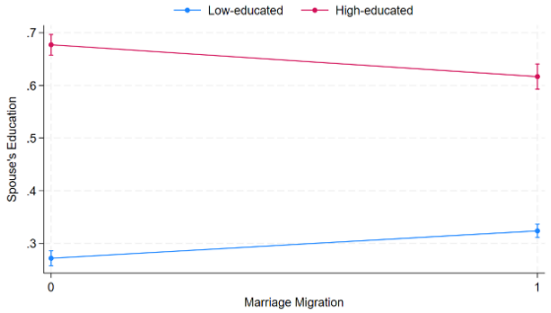

(c) Non-Muslim men

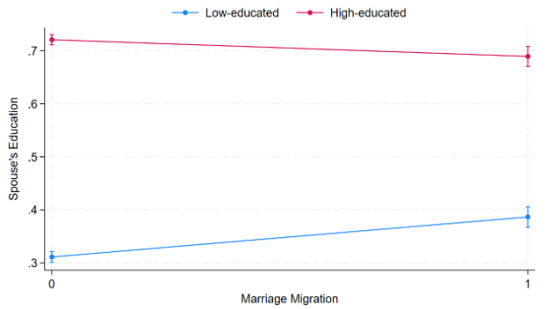

(d) Non-Muslim women

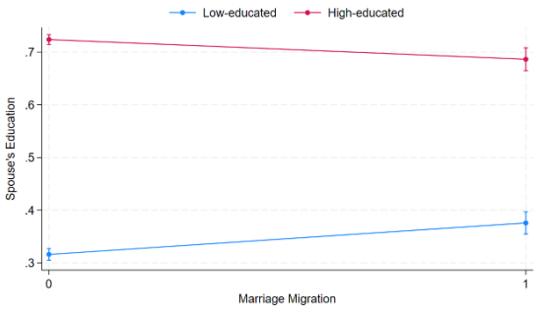

Notes: The figure reports average marginal effects from logit regressions of an indicator for whether the spouse holds a college degree on marriage migration, the individual's own college education, and their interaction.

Figure A6. Adjusted Marginal Effects of Marriage Migration on Spouse's Education

(a) Muslim men

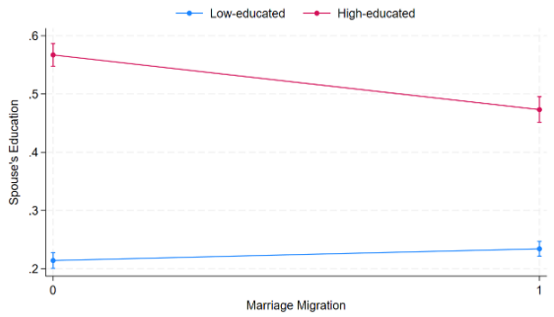

(b) Muslim women

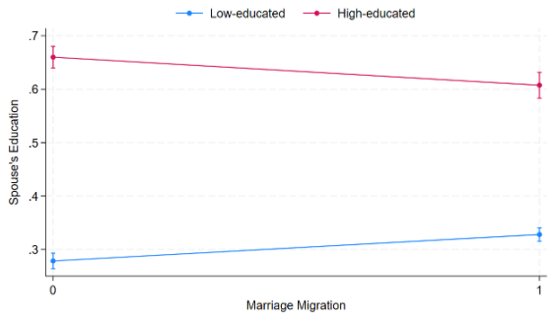

(c) Non-Muslim men

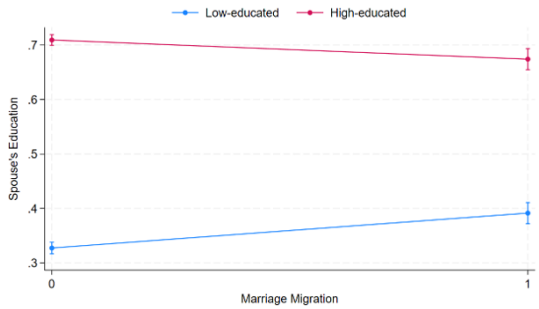

(d) Non-Muslim women

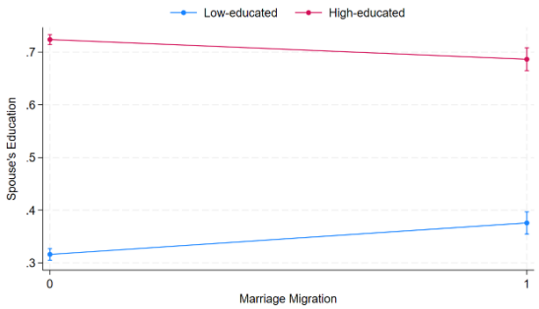

Notes: The figure reports average marginal effects from logit regressions of an indicator for whether the spouse holds a college degree on marriage migration, the individual's own college education, and their interaction. This specification includes individual-level controls, region fixed effects, and ten-year birth-cohort fixed effects.

Figure A7. Spatial Distribution of the Muslim Population, England and Wales

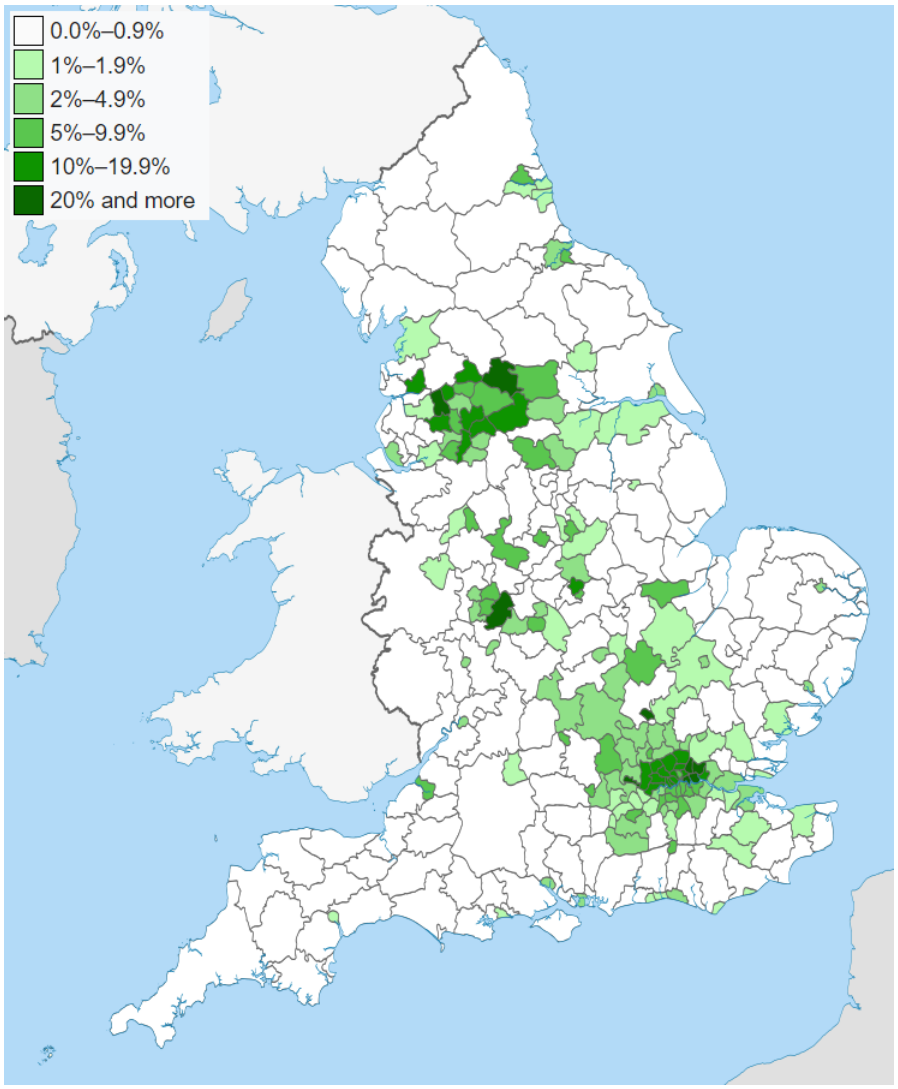

Source. Census for England and Wales, 2011

A.2 Tables

Table A1. Educational Sorting in Marriage Migration (Logit Model)

	Dependent variable: College education			
	Muslim		Non-Muslim	
	Male	Female	Male	Female
Baseline specification				
Marriage migration	-0.110***	-0.139***	-0.038***	-0.098***
	(0.008)	(0.008)	(0.008)	(0.009)
Adjusted specification				
Marriage migration	-0.096***	-0.137***	-0.052***	-0.114***
	(0.008)	(0.008)	(0.008)	(0.009)
Region & cohort FE	Y	Y	Y	Y
N	12,458	12,857	21,060	19,687

Notes. The table reports average marginal effects from logit models where the dependent variable is an indicator equal to one if the individual has a college degree. Baseline specifications include ethnic group, region, and 10-year birth-cohort fixed effects. Adjusted specifications additionally control for age and spouse's religion and ethnicity. Sample is limited to ethnic minorities. *t* statistics in parentheses. *p<0.10, **p<0.05, ***p<0.001. *Source.* Census for England and Wales, 2011.

Table A2. Age Trade-offs in Marriage Migration

	Dependant variable: Spouse's age			
	Muslim		Non-Muslim	
	Male	Female	Male	Female
Baseline specification				
Marriage migration	-0.481***	0.551***	-1.684***	0.690***
	(-6.8)	(7.6)	(-26.2)	(9.8)
Age	0.825***	0.888***	0.839***	0.851***
	(177.7)	(166.6)	(232.2)	(229.8)
Region & cohort FE	Y	Y	Y	Y
R^2	0.717	0.684	0.723	0.728
N	12,458	12,857	21,060	19,687
Adjusted specification				
Marriage migration	-0.461***	0.586***	-1.571***	0.682***
	(-6.4)	(7.9)	(-24.0)	(9.5)
Age	0.788***	0.803***	0.789***	0.783***
	(61.5)	(59.9)	(80.3)	(77.0)
Region & cohort FE	Y	Y	Y	Y
R^2	0.720	0.688	0.725	0.730
N	12,458	12,857	21,060	19,687

Notes. The table reports linear regressions where the dependent variable is the spouse's age. Baseline specifications include controls for own education, ethnic group, region, and 10-year birth-cohort fixed effects. Adjusted specifications additionally control for spouse's religion, spouse's ethnicity, and the education gap between spouses. t statistics in parentheses. *p<0.10, **p<0.05, ***p<0.001. *Source.* Census for England and Wales, 2011.

Table A3. Model Estimates with and without Geographic Restriction

	Limited Geography		England & Wales	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Endogamy preferences				
Ethnicity	10.5	7.4	10.4	7.4
Religion	23.1	23.1	22.7	22.7
Marriage migration (%)				
Overall	22.7	52.0	21.7	53.0
Outside market value	10.7	11.4	10.0	11.2
Endogamy preferences	12.0	40.6	11.8	41.8

Notes. The table compares estimated endogamy preferences and the decomposition of marriage migration under the baseline specification and a specification that restricts the sample to regions with relatively high Muslim population shares. The limited-geography sample includes the East Midlands, West Midlands, East of England, Greater London, and South East of England.

Table A4. Model Fit: Observed and Simulated Matching Outcomes

	Observed Matching		Simulated Matching	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Inter-education (%)				
Low-educated	23.8	26.6	22.7	26.9
High-educated	37.9	40.3	31.0	39.7
Inter-ethnic (%)				
White British	4.1	72.5	2.5	68.1
Other White	63.3	25.4	73.0	42.1
Indian	14.3	14.8	16.8	13.0
Pakistani/Bangladeshi	58.7	5.6	75.1	4.6
Other	30.8	22.4	37.2	39.3
Inter-religious (%)				
All	0.3	3.3	0.1	3.6
Marriage Migration (%)				
Other White	10.5	45.9	10.0	34.9
Indian	25.9	33.3	20.9	27.5
Pakistani/Bangladeshi	27.1	57.5	12.7	52.3
Other	24.9	44.6	23.3	35.3
Single (%)				
All	48.9	34.4	51.8	42.6

Notes. The table compares observed matching patterns in the data with matching outcomes simulated from the model using the estimated parameters.

Table A5. Estimated Endogamy Preferences with and without Marriage Migration

	With Migration		Without Migration	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Age				
Young & Middle-aged	6.54	7.59	6.50	7.93
	(0.47)	(0.47)	(0.47)	(0.47)
Young & Old	9.90	9.66	9.89	9.73
	(0.47)	(0.47)	(0.47)	(0.47)
Middle-aged & Old	5.81	7.28	5.59	7.49
	(0.47)	(0.47)	(0.47)	(0.47)
Education				
Low & High	3.78	3.41	4.04	3.40
	(0.21)	(0.21)	(0.21)	(0.21)
Ethnicity				
Other White & White British	6.59	7.06	5.85	4.18
	(1.31)	(1.31)	(1.31)	(1.31)
Indian & White British	16.10	14.88	14.93	13.41
	(1.31)	(1.31)	(1.31)	(1.31)
Pak/Bng & White British	16.30	9.97	14.05	7.67
	(1.31)	(1.31)	(1.31)	(1.31)
Other & White British	11.35	5.20	10.11	2.82
	(1.31)	(1.31)	(1.31)	(1.31)
Religion				
Muslim & Non-Muslim	22.59	22.59	22.67	22.67
	(1.10)	(1.10)	(1.15)	(1.15)

Notes. The table reports estimated endogamy preference parameters under a model that allows for marriage migration and a model that rules out marriage migration. The model without migration corresponds to the classic framework of Choo and Siow (2006).

Table A6. Comparison of Results by Nativity Status

	UK-born		UK-born or bred	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Endogamy preferences				
Ethnicity	10.2	6.1	10.4	7.4
Religion	22.4	22.4	22.7	22.7
Marriage migration (%)				
Overall	23.7	61.1	21.7	53.0
Outside market value	8.9	13.6	10.0	11.2
Endogamy preferences	14.8	47.4	11.8	41.8

Notes. The table compares estimated endogamy preferences and the decomposition of marriage migration under alternative definitions of the resident population. "UK-born" restricts the resident sample to individuals born in the UK. "UK-born or bred" corresponds to the baseline sample that also includes individuals who migrated to the UK before age 18.

Table A7. Simulated Outcomes under Counterfactual Preference Restrictions

		Counterfactual		
	Original	No Ethnic Preferences	No Religious Preferences	No Ethnic or Religious Preferences
Non-Muslims				
Price Gap				
Other White	1.9	2.7	1.9	2.7
Indian	1.4	3.3	1.5	3.3
Pakistani/Bangladeshi	1.7	3.4	0.7	2.8
Other	1.0	2.3	0.9	2.1
Mariage Migration	18.3	7.6	19.4	8.4
Inter-religious Marriage	0.3	0.2	15.7	6.3
Inter-ethnic Marriage				
White British	2.5	11.3	2.6	14.6
Other White	73.0	90.3	72.4	90.4
Indian	16.8	92.1	15.1	91.7
Pakistani/Bangladeshi	75.1	94.0	10.1	89.3
Other	37.2	84.7	35.4	82.8
Muslims				
Utility Gap				
Other White	0.3	1.7	2.2	2.9
Indian	1.1	2.4	2.0	3.8
Pakistani/Bangladeshi	-0.1	0.6	0.0	2.2
Other	0.2	0.9	0.7	1.8
Mariage Migration	47.4	29.2	41.0	10.0
Inter-religious Marriage	3.5	0.5	26.6	86.0
Inter-ethnic Marriage				
White British	68.1	88.7	2.8	18.5
Other White	42.1	80.1	71.9	90.6
Indian	13.0	77.5	15.9	92.9
Pakistani/Bangladeshi	4.6	31.1	8.2	86.5
Other	39.3	59.7	32.2	78.3

Notes. The table reports equilibrium outcomes simulated from the model under counterfactual restrictions on preference parameters. In the column "No Ethnic Preferences," all ethnic endogamy preference parameters are set to zero. In the column "No Religious Preferences," all religious endogamy preference parameters are set to zero. In the column "No Ethnic or Religious Preferences," both sets of parameters are set to zero. "Original" reports outcomes implied by the estimated baseline model.

Table A8. Estimated Endogamy Preferences by Religious Group

	$\mathcal{D}_{zz'} = \Phi_{zz} + \Phi_{z'z'} - \Phi_{zz'} - \Phi_{z'z}$				
Groups (z & z')	Muslim	Hindu	Sikh	Christian	No Religion
Age					
Young & Middle-aged	7.93	3.20	3.80	5.71	5.54
	(0.47)	(0.34)	(0.38)	(0.47)	(0.47)
Young & Old	9.73	5.28		7.97	7.74
	(0.47)	(0.34)		(0.47)	(0.47)
Middle-aged & Old	7.49	2.78	5.31	4.79	5.03
	(0.47)	(0.34)	(0.38)	(0.47)	(0.47)
Education					
Less than college & College or more	3.40	2.07	1.98	3.13	3.49
	(0.21)	(0.18)	(0.22)	(0.21)	(0.21)
Ethnicity					
Other White & White British	4.18			6.70	3.18
	(1.31)			(1.32)	(1.3)
Indian & White British	13.41	10.71	11.93	13.71	11.37
	(1.31)	(0.94)	(0.91)	(1.32)	(1.3)
Pak/Bng & White British	7.67			18.31	14.3
	(1.31)			(1.32)	(1.3)
Other & White British	2.82	5.30	10.97	11.15	10.01
	(1.31)	(0.94)	(0.91)	(1.32)	(1.3)
Religion					
Different religions	22.67	22.27	26.87	7.66	7.43
	(1.15)	(0.83)	(0.84)	(1.16)	(1.14)

Notes. The table reports estimated endogamy preference parameters for each religious group. $\mathcal{D}_{zz'}$ measures the relative surplus of endogamous matches compared to exogamous matches between groups z and z', holding all other characteristics constant. Parameters are obtained from separate structural estimations by religious group. Empty cells reflect combinations with insufficient observations. Standard errors are reported in parentheses.

B Derivations

This appendix derives the choice probability in equation (18) under the assumption that the idiosyncratic taste shocks are i.i.d. extreme value type I. Consider a woman of type x residing in location ℓ. For each potential partner type $y \in \mathcal{Y}$ and partner location $k \in \{1, \ldots, L\}$, her utility from matching with a man of type y in location k is

$$U(x, y, \ell, k) + \varepsilon_{iy,\ell k}$$

where $\varepsilon_{iy,\ell k}$ are i.i.d. across (y, k) with cdf $F(\varepsilon) = \exp(-\exp(-\varepsilon))$. The utility of remaining single is normalized to zero. The probability that the individual chooses option (y, k) is

$$P_{x,\ell}(y, k) = \Pr\Big\{U(x, y, \ell, k) + \varepsilon_{iy,\ell k} \geq U(x, z, \ell, m) + \varepsilon_{iz,\ell m},\ \forall (z, m)\Big\}$$

A standard property of the extreme-value assumption implies that this multinomial choice probability takes the logit form,

$$P_{x,\ell}(y, k) = \frac{\exp\big(U(x, y, \ell, k)\big)}{1 + \displaystyle\sum_{z \in \mathcal{Y}} \sum_{m=1}^{L} \exp\big(U(x, z, \ell, m)\big)} \tag{18}$$

An analogous expression holds for men.

C IIA in Two-sided Matching Models

This appendix clarifies the sense in which an Independence of Irrelevant Alternatives (IIA)–type property applies in the two-sided matching model used in this paper.

In contrast to one-sided discrete choice models, IIA cannot be interpreted as invariance of individual choice probabilities in a matching market, because matches require mutual agreement. Instead, under separable utilities and i.i.d. extreme value type I taste shocks, two-sided matching models satisfy a generalized IIA property: certain double odds ratios of matching frequencies are invariant to changes in population sizes (Galichon and Salanié, 2017).

Formally, consider any two female types $x, z \in \mathcal{X}$ and any two male types $y, t \in \mathcal{Y}$. Define the double odds ratio

$$\frac{\mu_{xy}\mu_{zt}}{\mu_{xt}\mu_{zy}}$$

Under the logit assumption, the joint surplus satisfies

$$\Phi(x, y) = 2 \ln \left(\frac{\mu_{xy}}{\sqrt{\mu_{x0}\mu_{0y}}} \right).$$

Taking logs of the double odds ratio yields

$$\ln \left(\frac{\mu_{xy}\mu_{zt}}{\mu_{xt}\mu_{zy}} \right) = \frac{1}{2} \Big[\Phi(x, y) + \Phi(z, t) - \Phi(x, t) - \Phi(z, y) \Big].$$

This expression depends only on differences in deterministic joint surplus and is independent of all subpopulation sizes. Consequently, population scale affects absolute matching frequencies but does not alter relative sorting patterns, as captured by double odds ratios. This invariance property is the appropriate analogue of IIA in two-sided matching models and underlies the identification of surplus parameters in the empirical analysis.

D IPFP Method

This appendix describes how the equilibrium matching distribution is computed given the estimated preference parameters, using an Iterative Proportional Fitting Procedure (IPFP).

As discussed in Section 4, the data contain information only on individuals residing in the destination country and on spouses who migrate for marriage. The distribution of singles in origin countries is not observed. Accordingly, when computing equilibrium matchings and counterfactuals, I treat outside market objects as fixed. In particular, the masses of single men and women in origin countries and migration utility terms are taken as exogenous to changes in domestic matching behavior.

Given joint surplus parameters Φ_{xy}, equilibrium matching flows implied by the transferable utility logit model satisfy the following conditions.

Local matches within the destination country are given by

$$\mu_{xy}^{\text{in}} = \exp\left(\frac{\Phi_{xy}}{2}\right) \sqrt{\mu_{x0}^{\text{in}}\, \mu_{0y}^{\text{in}}},$$

where μ_{x0}^{in} and μ_{0y}^{in} denote the masses of single women and single men of each type in the destination country.

Marriage migration flows are given by

$$\mu_{y\to x}^{\text{m}} = \exp\left(\frac{\Phi_{xy}}{2}\right) B_{y\to x} \sqrt{\mu_{x0}^{\text{in}}},$$
$$\mu_{x\to y}^{\text{m}} = \exp\left(\frac{\Phi_{xy}}{2}\right) B_{x\to y} \sqrt{\mu_{0y}^{\text{in}}},$$

Population balance in the destination country requires that, for each type,

$$n_x^{\text{in}} = \mu_{x0}^{\text{in}} + \sum_{y\in\mathcal{Y}} \left(\mu_{xy}^{\text{in}} + \mu_{y\to x}^{\text{m}}\right),$$
$$n_y^{\text{in}} = \mu_{0y}^{\text{in}} + \sum_{x\in\mathcal{X}} \left(\mu_{xy}^{\text{in}} + \mu_{x\to y}^{\text{m}}\right).$$

Substituting the expressions for local and migration matches yields a system of equations in the unknown single masses $\{\mu_{x0}^{\text{in}}, \mu_{0y}^{\text{in}}\}$.

The system is solved using an Iterative Proportional Fitting Procedure (IPFP), which is known to converge in logit matching models with transferable utility (Galichon and Salanié, 2022).

Starting from arbitrary positive initial values $\mu_{x0}^{\text{in},(0)}$ and $\mu_{0y}^{\text{in},(0)}$, the algorithm alternates between updating single women and single men masses according to

$$\mu_{x0}^{\text{in},(2k+1)} = \left(\sqrt{ n_x^{\text{in}} + \frac{1}{4} \left[\sum_{y \in \mathcal{Y}} \exp\left(\frac{\Phi_{xy}}{2} \right) \left(B_{y \to x} + \sqrt{\mu_{0y}^{\text{in},(2k)}} \right) \right]^2 } - \frac{1}{2} \sum_{y \in \mathcal{Y}} \exp\left(\frac{\Phi_{xy}}{2} \right) \left(B_{y \to x} + \sqrt{\mu_{0y}^{\text{in},(2k)}} \right) \right)^2$$

$$\mu_{0y}^{\text{in},(2k+2)} = \left(\sqrt{ n_y^{\text{in}} + \frac{1}{4} \left[\sum_{x \in \mathcal{X}} \exp\left(\frac{\Phi_{xy}}{2} \right) \left(B_{x \to y} + \sqrt{\mu_{x0}^{\text{in},(2k+1)}} \right) \right]^2 } - \frac{1}{2} \sum_{x \in \mathcal{X}} \exp\left(\frac{\Phi_{xy}}{2} \right) \left(B_{x \to y} + \sqrt{\mu_{x0}^{\text{in},(2k+1)}} \right) \right)^2$$

Iterations continue until convergence. Once the equilibrium single masses are obtained, all local and migration matching flows are recovered directly from the expressions above.

E Toy Model

This appendix presents a toy model to illustrate the mechanisms linking endogamy preferences and outside market values to marriage migration. The figures in this section are generated from a simulated data-generating process and are not based on observed data.

I consider a symmetric environment with a single binary characteristic, group membership. Let $\mathcal{X} = \mathcal{Y} = \{N, M\}$ denote the set of types for women and men, where N denotes the majority group and M denotes the minority group. In the destination country, the population shares are $1 - \gamma$ for the majority and γ for the minority. The destination population is normalized to one. The outside marriage market is populated exclusively by minority types and has total mass α, which captures the relative size of the outside pool of potential spouses available for marriage migration (Figure E1).

Figure E1. Toy Model

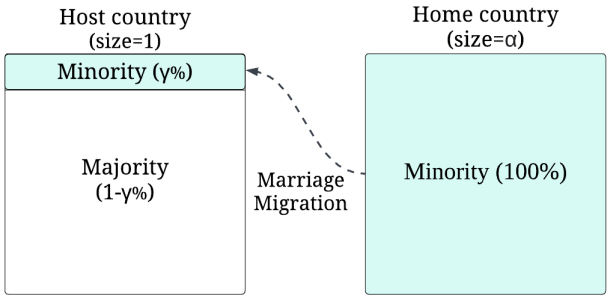

Notes. The model abstracts from all heterogeneity except group membership, with two groups labeled majority and minority.

Joint deterministic surplus from a match between types $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ is

$$\Phi(x, y, m) = a + b\, m - c \cdot \mathbf{1}\{x \neq y\},$$

where a captures the baseline gain from marriage relative to singlehood, b captures the incremental surplus associated with cross-market matching ($m = 1$), and c governs the strength of endogamy preferences.

Under the transferable-utility logit structure, local matching in the destination country satisfies

$$\mu^{\text{in}}_{xy} = \exp\left(\frac{\Phi(x, y, 0)}{2}\right) \sqrt{\mu^{\text{in}}_{x0}\mu^{\text{in}}_{0y}},$$

and cross-market matches that result in a minority spouse being brought from the outside market satisfy

$$\mu^{\mathrm{m}}_{y\to x} = \exp\left(\frac{\Phi(x,y,1)}{2}\right)\sqrt{\mu^{\mathrm{in}}_{x0}\mu^{\mathrm{out}}_{0y}},$$

with feasibility conditions defined analogously to the main model.

Two implications are immediate. First, the generalized IIA restriction implies that sorting depends on the endogamy parameter through a double-odds ratio:

$$\frac{\mu^{\mathrm{in}}_{NM}\mu^{\mathrm{in}}_{MN}}{\mu^{\mathrm{in}}_{NN}\mu^{\mathrm{in}}_{MM}} = \exp(-c).$$

Second, the relative prevalence of minority marriage migration depends on the cross-market surplus component and the availability of outside singles:

$$\frac{\mu^{\mathrm{m}}_{MM}}{\mu^{\mathrm{in}}_{MM}} = \exp\left(\frac{b}{2}\right)\sqrt{\frac{\mu^{\mathrm{out}}_{0M}}{\mu^{\mathrm{in}}_{M0}}}.$$

By simulating the model and solving for equilibrium across different parameter values, I illustrate how the key mechanisms generate variation in marriage migration and intermarriage outcomes. In Figure E2 shows the effect of migration utility b on equilibrium outcomes under $a = c = 0$, that is, in the absence of endogamy preferences. When migration entails a large disutility, all matches occur domestically. As migration utility increases, a growing share of minorities choose to marry abroad, which mechanically reduces the intermarriage rate between minorities and majorities. Higher migration utility also raises the attractiveness of marriage relative to singlehood for minorities, leading to a decline in minority singlehood. When migration utility is sufficiently large, almost all minority marriages take place abroad. These simulations highlight that migration utility alone can substantially depress intermarriage rates, even when individuals have no intrinsic preference for endogamy.

Figure E3 examines the effect of endogamy utility c holding migration utility fixed at zero ($a = b = 0$). Negative values of c correspond to a preference for intergroup marriage, while positive values reflect a preference for intragroup matching. When endogamy preferences are weak, intermarriage is prevalent and marriage migration remains limited. As endogamy utility becomes positive, matching increasingly occurs within groups. Because minorities have access to a larger pool of same-group partners abroad, stronger endogamy preferences translate into higher marriage migration

Figure E2. Equilibrium Effects of Migration Utility

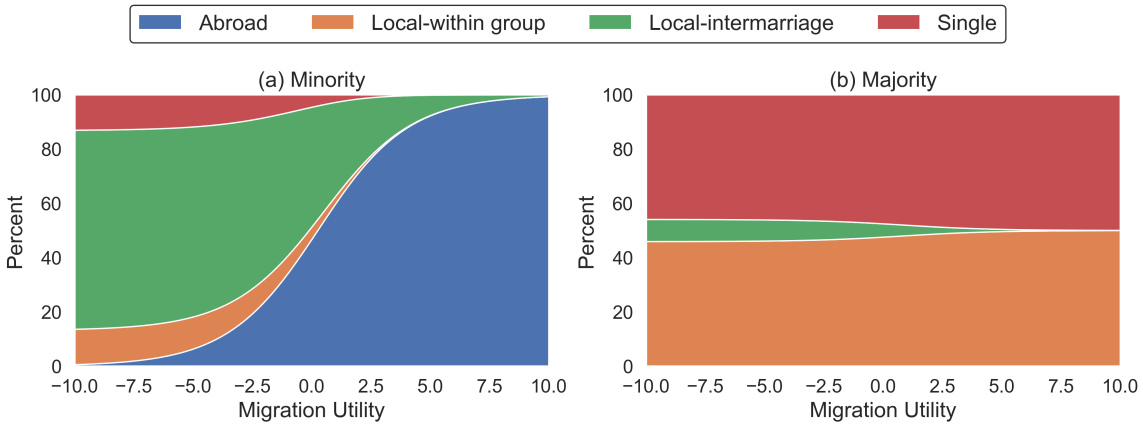

Notes. The figure reports simulated equilibrium outcomes as the migration utility parameter b varies. Preferences are given by $\Phi(x, y, m) = a + b\,m - c \cdot \mathbf{1}\{x \neq y\}$. The simulations set $a = 0$ and $c = 0$, so there are no endogamy preferences and no baseline gains from marriage relative to singlehood. Changes in equilibrium marriage migration, intermarriage, and singlehood therefore reflect variation in migration utility alone.

among minorities. As a result, increases in endogamy utility simultaneously reduce intermarriage and raise marriage migration.

Figure E3. Equilibrium Effects of Endogamy Utility

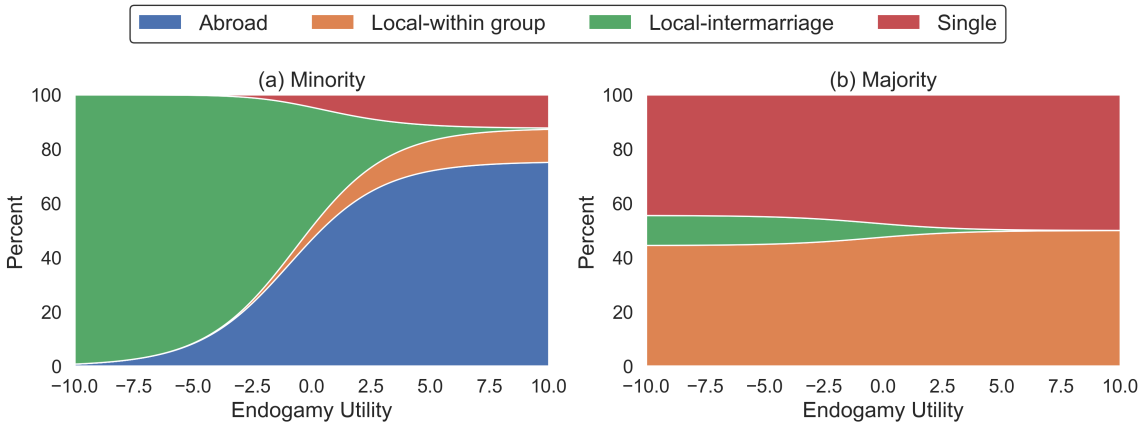

Notes. The figure reports simulated equilibrium outcomes as the endogamy utility parameter c varies. Preferences are given by $\Phi(x, y, m) = a + b\,m - c \cdot \mathbf{1}\{x \neq y\}$. The simulations set $a = 0$ and $b = 0$, so there are no baseline gains from marriage relative to singlehood and no migration utility. Changes in equilibrium marriage migration and intermarriage therefore reflect variation in endogamy preferences alone.

The final set of simulations compares the relative importance of endogamy preferences and outside market value in shaping marriage migration and intermarriage

outcomes for minorities.

Figure E4. Effect of Endogamy Preferences on Marriage Migration and Intermarriage

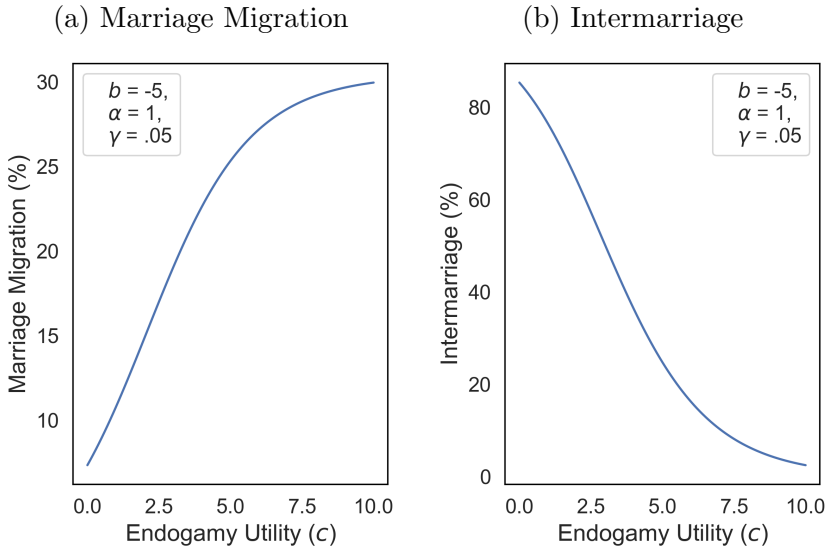

An increase in endogamy preferences directly lowers intermarriage rates by increasing the surplus required to sustain cross-group matches. At the same time, stronger endogamy preferences raise the relative attractiveness of accessing a larger pool of same-group partners abroad, thereby increasing marriage migration (Figure E4).

Outside market value affects marriage migration through two distinct channels (Figure E5). First, higher migration utility b directly raises the surplus associated with cross-market matches, increasing marriage migration. As minorities increasingly marry partners from abroad, same-group matching becomes more prevalent, which indirectly reduces intermarriage even in the absence of changes in endogamy preferences. Second, an increase in the size of the outside marriage market α lowers the effective cost of finding a spouse abroad by expanding the set of potential partners. This mechanism also increases marriage migration, though its quantitative effect is smaller than that of migration utility. In both cases, higher outside market value leads to greater marriage migration and lower intermarriage rates.

Overall, the simulations show that both endogamy preferences and outside market value can generate high marriage migration and low intermarriage. However, the underlying mechanisms differ: endogamy preferences operate primarily through sorting incentives, while outside market value operates through migration costs and the relative availability of same-group partners.

Figure E5. Effect of Outside Market Value on Marriage Migration and Intermarriage

(a) Marriage Migration

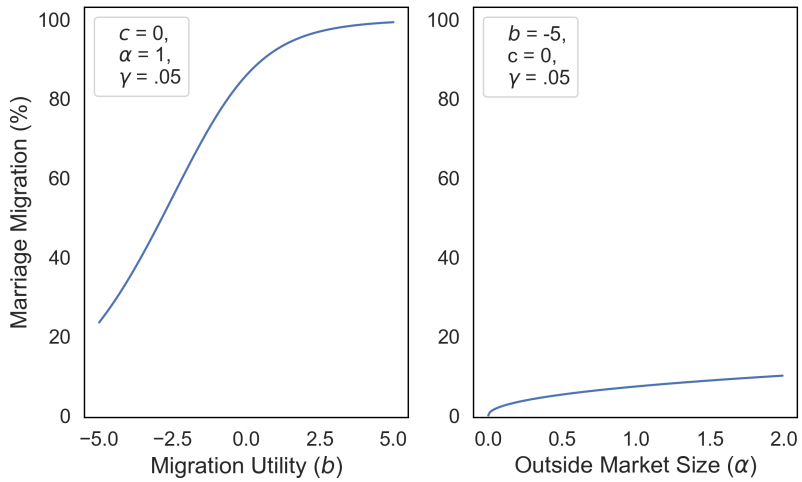

(b) Intermarriage

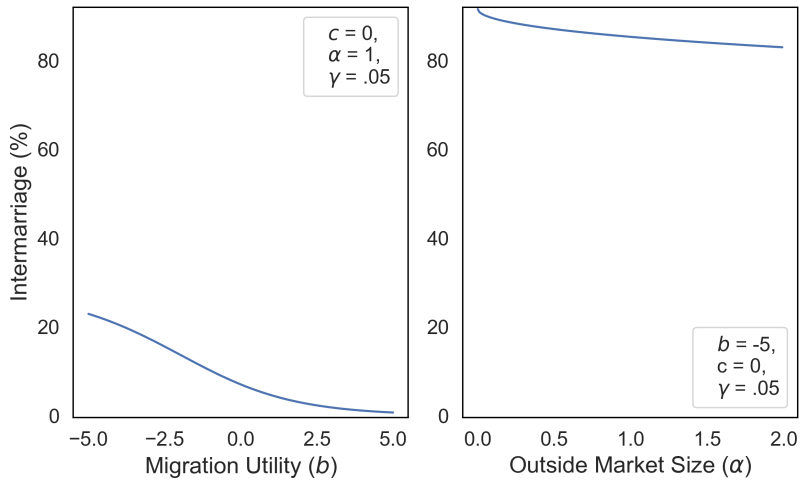

F Monte-Carlo Simulations

This appendix reports Monte Carlo simulations used to assess the finite-sample performance of the proposed estimation procedure. The simulations evaluate whether the estimator recovers the true preference parameters on average and whether its precision improves with sample size.

Individuals are divided into two groups, a majority and a minority, with minorities constituting 5 percent of the population. The deterministic joint surplus from a match between a woman i of type x_i and a man j of type y_j is specified as

$$\Phi_{ij} = \alpha_0 + \alpha_1 r_i + \alpha_2 r_j + \alpha_3 r_i r_j + \beta_1 s_i + \beta_2 s_j + \beta_3 s_i s_j + \varepsilon_{x_i y_j},$$

where r denotes education, s denotes income, and $\varepsilon_{x_i y_j}$ is a type-specific random shock that depends on the group types of the matched individuals. Minorities additionally have access to an outside marriage market. The outside market is assumed to be of the same size as the destination market, and cross-market matches generate an additional migration utility b.

For each Monte Carlo replication, I generate a synthetic matching market by drawing random utilities and equilibrium matches according to the model-implied multinomial choice probabilities. I then estimate the preference parameters using the minimum-distance procedure described in the main text. This procedure relies on a linear regression mapping between equilibrium matching frequencies and the underlying surplus parameters.

Figure F1 reports the sampling distributions of the minimum-distance estimates across Monte Carlo replications for different sample sizes ($N = 50, 500, 5000$). The true parameter values are set to $\alpha_1 = 0$, $\alpha_2 = 0$, $\alpha_3 = 5$, $\beta_1 = 1$, $\beta_2 = 1$, $\beta_3 = 3$, and $b = 2$. Dashed vertical lines indicate the true values. As sample size increases, the estimated distributions concentrate around the true parameters, indicating that the estimator is approximately unbiased and becomes increasingly precise in larger samples.

Figure F1. Monte-Carlo Simulation of Estimation Methodology

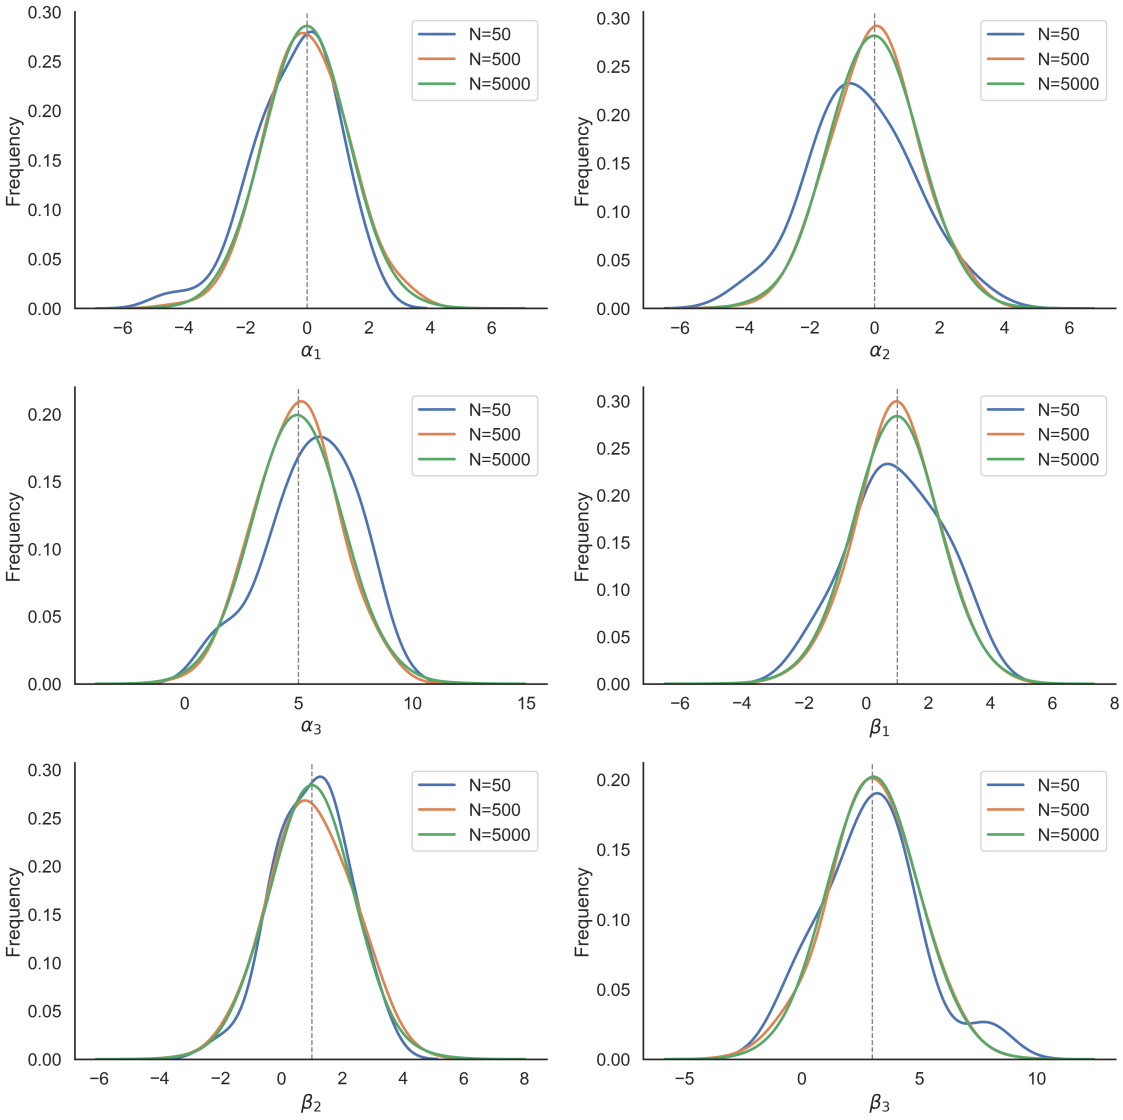

Notes. The figure reports the sampling distributions of the minimum-distance estimates across Monte Carlo replications for different sample sizes ($N = 50, 500, 5000$). Dashed vertical lines indicate the true parameter values. As the sample size increases, the distributions concentrate around the true parameters, illustrating consistency and lack of systematic bias.

G The Sample

The analysis uses the 10% sample of the 2011 Census for England and Wales, which contains 5,693,850 individual observations. In constructing the matched-couple dataset, I exclude the following individuals:

- individuals for whom a spouse cannot be identified (3,626 observations),
- individuals in same-sex couples (10,242 observations),
- individuals who are divorced, widowed, or separated (815,349 observations).

These exclusions reduce the sample by approximately 14.6%. I then restrict attention to men aged 25–55 and women aged 23–53. For married couples, if either partner falls outside the relevant age range, the couple is excluded. After applying these age restrictions, the sample consists of 1,850,766 observations.

To address missing information, I exclude individuals with missing education data and their spouses (0.15% of observations), as well as individuals who do not report their religion (8.5% of observations). Finally, to ensure that marriage decisions are observed within the UK marriage market, I restrict the Muslim sample to individuals who were either born in the UK or arrived before age 18. This restriction is necessary because a substantial share of Muslims are first-generation immigrants who may have married prior to arrival. Applying this criterion excludes 54% of Muslim observations.

The final estimation sample contains 1,423,555 individuals, of whom 38,938 (2.7%) are identified as Muslim.

H Local Mobility and the National Marriage Market Among British Muslims

One potential explanation for the persistence of marriage migration among British Muslims is the thinness of local marriage markets, particularly in areas with relatively small co-ethnic populations. If individuals place strong weight on religious or ethnic endogamy, local demographic constraints may limit the availability of suitable partners, especially outside large urban centres.

To examine whether such constraints are alleviated by internal mobility, I use data from Muslim marriage biographies collected across the UK. These biographies contain self-reported information on individuals' willingness to relocate within the UK for marriage. Using this information, I construct a relocation score ranging from 0 (not willing to move) to 1 (fully flexible with respect to location), which captures the extent to which individuals perceive the national marriage market as accessible rather than relying on local search.

The data reveal a pronounced gender asymmetry in mobility preferences. Muslim women report an average relocation score of 0.63, indicating substantial willingness to relocate across regions for marriage. By contrast, Muslim men report a much lower average score of 0.17, suggesting considerably more limited internal mobility.

These patterns imply that, at least for women, marriage search is not primarily constrained to local markets. Instead, many Muslim women appear to participate in a national marriage market, with a high degree of geographic flexibility. This internal mobility weakens the link between local co-ethnic population size and marriage outcomes and is consistent with the finding that marriage migration rates do not vary systematically across areas with different concentrations of Muslims. From the perspective of female marriage migrants, internal relocation may partially substitute for marrying abroad by expanding the effective pool of potential partners, thereby mitigating local demographic constraints.

I Assessing Measurement Error Due to Misclassification in Marriage Migration

As discussed in Section 3, the baseline measure of marriage migration may overstate spousal immigration by classifying as marriage migrants individuals who entered the UK for other reasons and subsequently married a UK resident. To assess the magnitude of this potential misclassification, I construct a conservative lower-bound measure of marriage migration using information from the first wave of Understanding Society.

Under this alternative definition, marriage migration is restricted to cases in which the foreign-born spouse reports arriving in the UK strictly after the calendar year of marriage. This criterion isolates unambiguous instances of post-marital migration. While it eliminates false positives, it also excludes some genuine marriage migrants, including couples who marry and migrate within the same calendar year and cases in which a UK resident marries a recent arrival shortly after migration even if the match was arranged prior to entry. The resulting figures therefore provide a strict lower bound on marriage migration.

Table I1 reports marriage migration rates under the baseline and lower-bound definitions by religion and gender. The lower-bound rates are uniformly below the baseline, with a larger gap for women than for men, particularly among Muslims. This pattern reflects the fact that a non-negligible share of Muslim women marry men who migrated independently to the UK as adults and only later entered the marriage market. These marriages are captured by the baseline definition but excluded by the lower-bound measure, implying greater scope for misclassification for women.

Table I1. Marriage Migration Rates: Baseline and Lower-Bound Definitions

		Baseline	Lower bound
Non-Muslim	Men	0.228	0.090
	Women	0.196	0.157
Muslim	Men	0.511	0.395
	Women	0.537	0.436

Notes. The baseline measure classifies marriages between UK residents and foreign-born spouses as marriage migration regardless of the timing of arrival. The lower-bound measure restricts marriage migration to cases in which the foreign-born spouse reports arriving in the UK strictly after the year of marriage.

Table I2 uses this alternative definition to reassess the contribution of outside

market value to overall marriage migration. Moving from the baseline to the lower-bound measure reduces the estimated contribution of outside market value from 0.126 to 0.099 for women and from 0.118 to 0.111 for men. Although the adjustment is quantitatively meaningful for women, the overall magnitudes remain similar. This indicates that the main conclusions of the decomposition are not driven by misclassification of marriage migration. The baseline estimates can therefore be interpreted as an upper bound on the contribution of outside market value, with the lower-bound results providing a conservative robustness check.

Table I2. Baseline and Lower-Bound Estimates of Outside Market Value

	Women	Men
Lower bound	0.099	0.111
Baseline	0.126	0.118

Notes. Each entry reports the estimated contribution of outside market value to overall marriage migration under the corresponding definition of marriage migration. The lower-bound definition restricts marriage migration to cases in which the foreign-born spouse arrives strictly after the year of marriage.

J Robustness Checks

This section examines the robustness of the estimated preferences and decomposition results to alternative sample constructions and definitions of marriage migration.

J.1 Cohabitation

In the baseline analysis, cohabiting individuals are classified as single. However, cohabitation accounts for a sizable fraction of observed partnerships, particularly among younger non-Muslims (Figure A2). As a robustness check, I reconstruct the estimation sample by treating cohabiting opposite-sex couples as domestic matches rather than as singles. I continue to classify marriage migration using formal marriages only, since sponsoring a spouse visa requires a marriage certificate and cohabitation does not provide an independent channel for importing a partner.

Table J1. Sensitivity to the Treatment of Cohabitation

	Cohabitation Treated as Marriage		Cohabitation Treated as Single	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Endogamy preferences				
Ethnicity	10.2	7.3	10.4	7.4
Religion	21.9	21.9	22.7	22.7
Marriage migration (%)				
Overall	19.6	51.3	21.7	53.0
Outside market value	9.8	11.3	10.0	11.2
Endogamy preferences	9.8	40.1	11.8	41.8

Notes. This table compares estimates obtained under alternative treatments of cohabitation.

Comparing these results with the baseline estimates (Table J1), endogamy preferences remain qualitatively and quantitatively robust. Estimated religious endogamy preferences are approximately 4% lower, reflecting the higher prevalence of interreligious mixing among cohabiting couples, but the differences are not significant. The overall marriage migration rate declines by roughly 2 percentage points, as cohabiting couples are now classified as local matches, increasing the size of the domestic market. Importantly, the decomposition of marriage migration into endogamy preferences and outside market value remains largely unchanged, indicating that the main conclusions are not driven by the treatment of cohabitation.

J.2 Marriage Migration Definition: Sensitivity Analysis

This subsection investigates the sensitivity of the main results to alternative definitions of marriage migration. A first concern is that the baseline definition may include individuals who initially migrated to the UK for education and subsequently married, rather than individuals who migrated specifically for marriage. If such individuals differ systematically in their outside market value, this could affect the decomposition results.

Table J2. Sensitivity to the Definition of Marriage Migration

	Excluding Migrants with UK Degree^{+}		Baseline Definition^{++}	
	Non-Muslim	Muslim	Non-Muslim	Muslim
Endogamy preferences				
Ethnicity	10.6	7.9	10.4	7.4
Religion	22.6	22.6	22.7	22.7
Marriage migration (%)				
Overall	18.0	49.8	21.7	53.0
Outside market value	8.0	9.2	10.0	11.2
Endogamy preferences	10.0	40.6	11.8	41.8

Notes. This table reports estimates under alternative definitions of marriage migration. $^{+}$ Immigrants who obtained a university degree in the UK are excluded from marriage migration and treated as resident spouses. $^{++}$ Baseline definition used in the main analysis.

As a robustness check, I reclassify marriages in which the migrant spouse holds a UK university degree as local marriages rather than marriage migration. The results, reported in Table J2, show that this adjustment reduces the marriage migration rate by approximately 3 percentage points. Endogamy preferences remain stable, while the estimated contribution of outside market value declines slightly. This finding reinforces the conclusion that preferences for endogamy account for the bulk of observed marriage migration, even under more restrictive definitions.

In the main analysis, marriage migration is defined as a match between a UK-born individual or someone who arrived in the UK before age 18 and a spouse who migrated to the UK after age 18. This threshold is chosen to ensure that marriage decisions are made while residing in the UK. However, some individuals may have marriage arrangements prior to reaching adulthood. To assess sensitivity to this assumption, Table J3 reports results using alternative age thresholds of 16, 18, and 20.

Table J3. Sensitivity to the Age Threshold Used to Define Marriage Migration

	Threshold = 16		Threshold = 18		Threshold = 20	
	Non-Muslim	Muslim	Non-Muslim	Muslim	Non-Muslim	Muslim
Endogamy preferences						
Ethnicity	10.4	6.7	10.4	7.4	10.4	8.3
Religion	22.4	22.4	22.7	22.7	22.2	22.2
Marriage migration (%)						
Overall	22.1	58.1	21.7	53.0	20.6	39.5
Outside market value	10.0	12.0	10.0	11.2	9.7	7.9
Endogamy preferences	12.1	46.1	11.8	41.8	10.9	31.7

Notes. Marriage migration is defined as a union between a UK resident (UK-born or arrived before the specified age threshold) and a spouse who migrated to the UK at or after that threshold age. The table reports results for alternative age cutoffs used to classify early-arrival migrants as residents rather than marriage migrants.

As the threshold age increases, the measured marriage migration rate declines, as older cohorts who marry at younger ages are reclassified as local marriages. Despite this variation, the proportional contribution of outside market value to overall marriage migration remains remarkably stable across thresholds, providing further reassurance regarding the robustness of the decomposition. Religious endogamy preferences are unaffected. In contrast, estimated ethnic endogamy preferences increase with the threshold age, primarily because some interethnic marriages previously classified as marriage migration are reclassified as local matches at higher thresholds.