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

DISCUSSION PAPER SERIES 202/26

Marriage Market Equilibrium with Matching on Latent Ability: Identification using a Compulsory Schooling Expansion

Dan Anderberg, Jesper Bagger, V. Bhaskar, Tanya Wilson

www.rfberlin.com

JULY 2026

Marriage Market Equilibrium with Matching on Latent Ability: Identification using a Compulsory Schooling Expansion

Authors

Dan Anderberg, Jesper Bagger, V. Bhaskar, Tanya Wilson

Reference

JEL Codes: D10, D13, I26, J12

Keywords: Marriage, Assortative mating, Return to education, Ability

Recommended Citation: Dan Anderberg, Jesper Bagger, V. Bhaskar, Tanya Wilson (2026): Marriage Market Equilibrium with Matching on Latent Ability: Identification using a Compulsory Schooling Expansion. RFBerlin Discussion Paper No. 202/26

Access

Disclaimer

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

Marriage Market Equilibrium with Matching on Latent Ability: Identification using a Compulsory Schooling Expansion*

Dan Anderberg[†] Jesper Bagger[‡] V. Bhaskar[§]
Tanya Wilson[¶]

July 15, 2026

Abstract

We develop an identification strategy for equilibrium matching models with latent traits and apply it to the 1972 Raising of the School-Leaving Age (RoSLA) reform in England and Wales. Our identification exploits the RoSLA-induced discontinuity in the distribution of qualifications to disentangle the contributions of qualification and ability to marital surplus. Both are valued and are complements in the marital surplus function. Ability increases the probability of ever marrying while basic qualification attainment does not. Hence, the observed gap in marriage rates between basic qualified and unqualified individuals is entirely due to selection on ability. The RoSLA worsened marital prospects of low ability individuals through general equilibrium effects.

Keywords: Marriage, Assortative mating, Return to education, Ability
JEL Classification: D10, D13, I26, J12

*We thank the editor, three anonymous referees, Dan Ackerberg, Pierre-André Chiappori, Monica Costa Dias, Arnaud Dupuy, Jeremy Fox, John Knowles, Jungyoon Lee, Cezar Santos, and attendants at numerous conferences and seminars for helpful comments.
[†]Royal Holloway; IFS; CESifo. **E-mail**: dan.anderberg@rhul.ac.uk
[‡]University of Edinburgh; Aarhus University; IZA; RFBerlin. **E-mail**: jbagger@ed.ac.uk
[§]University of Texas at Austin; CEPR. **E-mail**: v.bhaskar@austin.utexas.edu
[¶]University of Glasgow. **E-mail**: tanya.wilson@glasgow.ac.uk

1. Introduction

Structural equilibrium models of matching typically require researchers to specify the systematic traits on which agents match. Yet many economically important traits are latent and correlated with observed characteristics. This creates an identification problem. This paper exploits a major policy reform that changed the mapping from latent ability to educational qualifications to identify and estimate an equilibrium matching model of the marriage market in which educational qualifications and unobserved ability are both important.

To put this in context: estimating the returns to education remains a central issue in economics. As David Card's influential survey (Card, 2001) notes, accounting for unobserved heterogeneity is critical, since individuals with higher educational attainment also have latent traits, e.g. ability, that make them different from those with lower educational attainment. Much empirical research in the past three decades has focused on finding quasi-experiments or instrumental variables to disentangle the effects of education and ability on *individual* pecuniary and non-pecuniary outcomes (Oreopoulos and Salvanes, 2011).

One avenue for a return to education that has recently received increased attention is via the marriage market (Chiappori, Iyigun, and Weiss, 2009; Chiappori, Salanié, and Weiss, 2017). Since marriage formation is an *equilibrium* phenomenon, this literature has relied on structural models of the marriage market, as pioneered by Choo and Siow (2006), to estimate the payoffs in marriage to the two genders by their level of educational attainment. In this literature, matching is on traits that are observable to an econometrician with no role for latent variables, like ability.

This paper bridges the two methodological approaches by developing an identification strategy for equilibrium matching models with latent traits. The strategy exploits reforms that change the mapping between observable characteristics and latent types while leaving the latent trait distribution unchanged. This is our main methodological contribution. We demonstrate it using the 1972 Raising of the School-Leaving Age (RoSLA) order in England and Wales: our identification theorem establishes that such a change can disentangle the separate roles of latent traits and observed characteristics in equilibrium matching.

The same logic applies whenever a reform (i) alters selection into an observable category and thereby makes that category more informative about a latent trait, (ii) leaves the distribution of the latent trait unchanged, and (iii) operates in a market where the treated and untreated cohorts interact in equilibrium. The last condition is essential and is what distinguishes our approach: identification comes from the matches that form between cohorts facing different mappings, so they must share a common market rather than occupy separate ones. Such reforms are not uncommon. The introduction of military conscription can decouple selection into service from latent characteristics such as health or preferences; policies mandating participation in high-stakes examinations can transform educational credentials from a noisy proxy into a sharp indicator of underlying ability. In any such setting, the resulting cross-regime variation in the joint distribution of observables and latent traits can be exploited to identify the role of those traits in equilibrium behaviour.

The RoSLA raised the school leaving age in the UK from 15 to 16 for individuals born September 1957 or later. In England and Wales, unlike in the US, students are eligible to leave school at the end of the academic year in which they reach the minimum school-leaving age, rather than on their birthday. Since basic qualifications[1] are awarded based on exams taken at the end of the academic year in which a student turns 16, the RoSLA effectively required students to remain in school until they were eligible to sit these exams. The RoSLA produced a sharp increase in the share of school leavers obtaining a basic qualification, yet a substantial fraction still left school unqualified even after the reform. We interpret these facts as follows. Prior to the reform, unqualified school leavers included both individuals who lacked the academic ability to obtain a basic qualification and those who chose not to pursue one due to the opportunity cost—namely, the earnings they would forgo by staying in school. By compelling students to remain in school until exam eligibility, the RoSLA sharply reduced this opportunity cost for academically able students, who consequently became much more likely to obtain a basic qualification. In effect, while the pre-reform unqualified pool was composed of both those who lacked the ability to obtain a qualification and those who chose not to, the post-reform pool consisted almost entirely of the former.

The marriage market consequences of the RoSLA are clearly visible in data on the 1953–1960 academic cohorts in England and Wales. The never-married rate for unqualified men and women increased discretely and permanently at the RoSLA threshold. Qualification homogamy among the unqualified increased, while overall qualification homogamy remained largely unchanged. Finally, the very first RoSLA-exposed cohort was more likely to marry with a zero age gap than previous and subsequent cohorts. These patterns are jointly inconsistent with either ability

[1] A basic qualification is an Ordinary Level (O-Level) or a Certificate of Secondary Education (CSE) qualification, the first tier of accredited examinations in England and Wales.

or qualification alone driving marriage market outcomes. If only ability mattered, the reform—which left the ability distribution unchanged—would have had no marriage market consequences at all. If only qualification mattered, an estimated qualification-only model predicts higher marriage rates for the unqualified post-RoSLA—the opposite of what we observe. Together, these observations call for a model in which both ability and qualification enter the marital surplus function.

Formally, identification exploits marital mixing across the RoSLA threshold between pre- and post-RoSLA cohort members who interact in a common marriage market but face different joint distributions of ability and qualifications. The identifying variation comes from the reform-induced change in the relationship between qualifications and ability. In the post-reform cohorts, ability is much more strongly aligned with qualification attainment. Two key assumptions underpin our identification: that the RoSLA did not shift the underlying ability distribution—a standard assumption in reform-based empirical work—and that the marital surplus function was invariant to the reform. The latter assumption is non-trivial, since the reform could have altered labor market returns to qualification and thereby the surplus from marriage to a qualified spouse; we provide auxiliary empirical evidence for the plausibility of this invariance assumption.

We bring this identification strategy to bear on data from the 1953–1960 academic cohorts in England and Wales, and obtain three main findings. First, ability and qualifications both enter the marital surplus function, with the impact of ability at least as important as that of a basic qualification; moreover, own and spousal ability are complements in the surplus function as are own and spousal qualification. Second, strikingly, ability increases the probability of ever marrying while basic qualification attainment does not—hence the observed gap in marriage rates between qualified and unqualified individuals is entirely due to selection on ability: unqualified individuals marry less frequently not because they lack qualifications, but because they tend to have low ability. Moving from low to medium ability—even without gaining a qualification—increases the ever-marrying rate by up to 10 percentage points for men and 5 percentage points for women. Third, simulation of the counterfactual marriage market where the RoSLA was not implemented reveals that the reform increased marital mixing between medium and high ability individuals, worsened marriage market outcomes for low-ability individuals, and generated spillover effects on pre-reform cohorts not directly exposed to the change in school-leaving age.

1.1. Related literature

Choo and Siow (2006, henceforth CS) is the seminal empirical implementation of the Becker-Shapley-Shubik transferable utility marriage market model (Shapley and Shubik, 1971; Becker, 1973). Marriage market participants are grouped into a finite number of observable types, with a systematic marital surplus function that depends on a couple's type-profile, and with additively separable individual i.i.d. random preferences over partner type. The original CS framework has been extended in various directions and applied in many contexts—see Dupuy and Galichon (2014), Galichon and Salanié (2021), Choo (2015), Brandt, Siow, and Vogel (2016), Mourifié and Siow (2021), and Chiappori, Salanié, and Weiss (2017). Notably, Chiappori, Salanié, and Weiss (2017) use the CS approach to decompose changes in marriage patterns over time into supply and preference changes. Noting that assortative matching occurs when the marital surplus—reflecting preferences—exhibits supermodularity, they argue, using data on educational attainments and marital outcomes in the US of cohorts born over a long time-frame (1940s through to the 1970s), that assortative matching in the US has increased, and particularly so at the top of the education distribution.[2]

The literature has to date remained focused on marital sorting on directly observable traits, with only two notable exceptions: Chiappori, Costa Dias, and Meghir (2018) and Adda, Pinotti, and Tura (2025). Chiappori, Costa Dias, and Meghir (2018) introduce latent ability in the CS framework. They develop a dynamic model of investment in human capital, marriage and household labor supply, where matching takes place on the basis of a one-dimensional variable—human capital—which depends on latent ability and acquired education. Their identification strategy rests on a fully dynamic model and upon observing individual's earnings. We explore a static model where identification requires data only on aggregate marriage frequencies by cohort and qualification. A further difference is that the type space for the marriage market in Chiappori, Costa Dias, and Meghir (2018) is two-dimensional—human capital and cohort—while our type space is three-dimensional—ability, qualifications and cohort.

Adda, Pinotti, and Tura (2025) depart from the CS-framework by introducing a match-specific i.i.d. random component to marital utility ("love")—allowing for sorting on this unobserved component—and also endogenous separations. They study how legal status affects marriage choices of natives and migrants in Italy,

[2] Arcidiacono, Beauchamp, and McElroy (2016) consider utility to be imperfectly transferable. Following Dagsvik (2000), they set up and estimate a matching model, where individuals can also choose the terms of the relationship, which they use to analyze high-school relationships.

leveraging plausibly exogenous variation in individual legal status from the 2004-2007 EU enlargements. Our analysis also obtains identification through exogenous variation—here, in qualification attainment—driven by a major policy reform. We retain the key CS-assumption of additively separable i.i.d. random preferences over traits observable in the marriage market—age, ability, and qualification. However, crucially, we allow one trait, ability, to be unobserved by the econometrician but potentially correlated with qualification.

Recently, Eika, Mogstad, and Zafar (2019) used data from the US Current Population Survey to show that homogamy in the US increased gradually between 1962 and 2013 among the less educated, but decreased somewhat among the highly educated. Their findings are contrasted by Chiappori, Costa Dias, Meghir, and Zhang (2025) who argue that assortative matching has indeed increased in the US, but has done so particularly at the top of the education distribution. Our approach is methodologically distinct from Chiappori, Costa Dias, Meghir, and Zhang (2025): whereas their approach is to study how marital surplus discretely shifted over time by comparing the educational attainment distributions and the marriage outcomes of individuals born in the 1950s to those born in the 1970s, our approach is to focus on cohorts born closely around a threshold date for a reform that discontinuously shifted the supply of qualifications, with the cohort window being small enough that the marital surplus can plausibly be assumed to be stable over the window. We find that the RoSLA worsened the marital outcomes particularly of the unqualified.

Our work is also related to the substantial literature on the effects of the RoSLA on individual outcomes. Grenet (2013) finds that the pecuniary returns to the RoSLA-induced schooling were modest. With the exception of reduced criminality (Machin, Marie, and Vujić, 2011; Bell, Costa, and Machin, 2022), there is also no evidence of nonpecuniary returns, see Clark and Royer (2013) and Silles (2011) for health and fertility outcomes. Generally, European evidence on the pecuniary returns to schooling identified from compulsory schooling laws (Pischke and von Wachter, 2008; Devereux and Hart, 2010; Grenet, 2013) show lower returns than evidence from North America (Angrist and Krueger, 1991; Oreopoulos, 2006). Our substantive results are largely consistent with this finding: we find that educational qualifications do matter for marital surplus, but latent ability matters more.

2. Data Sources

Our empirical analysis focuses on England and Wales and combines data from the Labour Force Survey, Population Statistics from the Office for National Statistics (ONS), and the 2011 Census. Part of our analysis also uses the National Child Development Study. Further details of our data sources can be found in the Supplementary Appendix.

The academic year in England and Wales runs from the 1st of September to the 31st of August in the following calendar year. Members of the 1957 academic cohort—the first RoSLA affected cohort—are born between September 1957 and August 1958. We focus on a set of academic cohorts $\mathcal{C} = \{1953, ..., 1960\}$ born around the RoSLA threshold. We will work with three ordered qualification levels, $z \in \mathcal{Z} = \{z^0, z^1, z^2\}$, representing no academic qualifications, a basic academic qualification, and an advanced academic qualification, respectively. A basic qualification refers to an Ordinary Level (O-Level) or a Certificate of Secondary Education (CSE) qualification, which is the very first tier of accredited examinations in England and Wales. An advanced qualification refers to an Advanced Level (A-Level) qualification or higher.

2.1. Labour Force Survey

The Labour Force Survey (LFS) is the largest regular household survey in the United Kingdom and is intended to be representative of the UK population. Between 1983 to 1991, the LFS was annual and from 1992 onwards it has been conducted quarterly. The LFS contains information on year and month of birth for each household member and on relationships between household members. Detailed qualifications information has been included since 1984. We pool all individuals observed in the 1984 - 2014 LFS, born in the UK and resident in England and Wales, and who are from some academic cohort $c \in \mathcal{C}$. We use the LFS data to estimate the impact of the RoSLA on academic qualification rates and specifically to characterize the marital matching patterns by cohort and qualification profiles.[3] Our focus is on formal marriage as a long-term individual outcome, and as noted

[3] The LFS records current marriage partners only, so we assume that spousal qualification and ability profiles in observed marriages are representative of first marriages. If marital dissolution depends systematically on the qualification and ability profiles of couples, this assumption may be violated. As a robustness check, we have re-estimated the model on couples aged 40 or below, where selection from dissolution is reduced. The estimated surplus functions and implied complementarities are qualitatively very similar to the baseline. Details of this robustness check are provided in the Supplementary Appendix.

Table 1: Pooled LFS Sample of Individuals from Academic Cohorts 1953-1960

	Males			Females		
	All	Single	Married	All	Single	Married
Age in Years	38.86	38.05	39.25	38.99	39.44	38.80
	[9.11]	[9.56]	[8.86]	[9.12]	[9.41]	[8.98]
Academic cohort	1956.60	1956.93	1956.44	1956.60	1956.86	1956.49
	[2.29]	[2.28]	[2.28]	[2.30]	[2.30]	[2.29]
No Qualification z^0	0.35	0.38	0.34	0.32	0.35	0.31
Basic qualification z^1	0.38	0.36	0.39	0.43	0.40	0.45
Advanced qualification z^2	0.27	0.26	0.27	0.24	0.25	0.24
Observations	147,878	47,832	100,046	156,549	47,059	109,490

Notes: The sample pools all individuals observed in the 1984-2014 Labour Force Surveys from academic cohorts 1953-1960 with non-missing information on age, qualification and marital status. Standard deviations in square brackets. Basic qualification refers to CSE/O-Level qualifications, and advanced qualifications refer to A-Level or higher qualifications. See main text for details.

just below, we use census data to characterize the proportion never-married by age 45 by gender, academic cohort and qualification level.[4]

Table 1 gives descriptive statistics for the LFS data by gender and marital status. 68 (70) percent of men (women) are married at the time of the interview. The average age is 39 for both men and women and the average academic cohort is close to 1956.5.[5] For both men and women, the basic qualification ($z = z^1$) is the most common, followed by no qualification ($z = z^0$), and then by advanced qualifications ($z = z^2$).

2.2. *ONS Population Statistics and Census Data*

We use birth statistics from the ONS to calculate academic cohort size by gender for England and Wales.[6] The UK experienced a baby boom starting in the mid-1950s and peaking in 1964. Hence, as a general characterization, the cohorts that

[4]The increase in never-marriage in the affected cohorts is driven by individuals without partners rather than substitution into long-term cohabitation. Retrospective cohabitation histories from the General Household Survey (1998–2000) show that cohabitation in these cohorts appears primarily as a transitional stage preceding marriage rather than a long-term substitute for it. Haskey (2001) documents a similar pattern for England and Wales over this period.

[5]We make no further use of the age variable as we instead use direct measurements at population-level of the fraction never-married by age 45 by gender, cohort and qualification as described below.

[6]We further apply gender-cohort mortality rates to calculate academic cohort size at age 25. The gender-specific mortality rates by age were obtained from the ONS's principal projection of historic and projected mortality rates from the 2010-based UK Life Tables. Our focus on English and Welsh birth cohorts also means that we are ignoring migration when calculating the relative populations supplies.

we are studying were steadily increasing in size by on average 3 percent per year, from around a cohort size of 650,000 towards the mid-1950s to close to 800,000 by 1960.

We further commissioned tabulated data from the ONS based on the 2011 Census in order to characterize never-married rates, by gender, academic cohort and qualification level. As the census is fixed in time, we adjusted the tabulated data to account for first marriages occurring past the age of 45, leaving us with a measure of the proportion never-married by age 45 by gender, academic cohort and qualification level.[7]

2.3. National Child Development Study

The National Child Development Study (NCDS) initially sampled individuals born in England, Scotland and Wales in one particular week of March, 1958. The initial survey was known as the Perinatal Mortality Study and established a sample of 17,415 cohort members. The NCDS follows the cohort members over time through a set of regular "data sweeps" that combine survey information with administrative data from e.g. schools. We make use of the 1974 (at age 16), 1991 (at age 33), and 2000 (at age 42) sweeps with measurements of academic ability, and information on preferred and realized school leaving age, qualification attainment, and marital status. To be consistent with the LFS data, we restrict the NCDS sample to individuals born in England and Wales. Further details of the NCDS data are provided in the Supplementary Appendix.

We use the NCDS to empirically document how the RoSLA changed selection on ability into basic qualification attainment. This informs the specification of our extended CS model. We also compute NCDS never-married rates to further validate the specification. However, a full-fledged analysis of the marriage market impact and implications of the RoSLA using the NCDS would not be feasible. The NCDS contains only limited information on spouses, certainly not measured ability, is small compared to the LFS and Census datasets, and contains individuals only from the very first RoSLA-affected cohort.

[7]We use information from ONS Cohabitation and Cohort tables for England and Wales on the proportions of never-married individuals by cohort over single years of age. However, first marriages past the age of 45 are rare, whereby the adjustments are very small.

3. The 1972 Raising of the School Leaving Age

3.1. The Reform

In the UK, individuals become eligible to leave school at the end of the academic year in which they reach the minimum age requirement, rather than on the day they reach the specified compulsory school leaving age. In March, 1972, the UK Government introduced Statutory Act 444, known as the Raising of School Leaving Age (RoSLA) Order. The RoSLA came into operation on September 1st, 1972, raising the minimum school leaving age requirement by one year to age 16, affecting individuals born from 1st September, 1957 onwards. In England and Wales, the RoSLA prompted a dramatic 25 percentage points increase in the proportion of individuals leaving education at age 16, rather than at age 15. The effect was limited to those individuals at the lower end of the education distribution, with evaluations routinely concluding that the reform had no effect on the probability of leaving at ages 17 or above (see e.g. Chevalier, Harmon, Walker, and Zhu, 2004; Silles, 2011; Clark and Royer, 2013).

In England and Wales there are two levels of national accredited examinations sat during school. The first tier, which we label a basic qualification, led (for the cohorts of interest here) to O-Level or CSE qualifications.[8] The examinations are available at the end of the academic year in which an individual turns 16, with students typically taking examinations in a minimum of five subjects. A qualification is awarded if the student's performance exceeds the minimum passing grade. If an individual chooses to remain in school beyond that point, then at the end of the academic year in which they turn 18, a second tier of academic examinations can be taken, leading to A-Level qualifications which also serves as a prerequisite for entry to higher education.

The RoSLA-mandated increase in the school leaving age required students to remain in school for the year in which the first level of academic qualifications are conferred. The RoSLA thereby not only impacted the duration of schooling, but also raised the likelihood of leaving school with an academic qualification for those students with the academic ability to achieve a passing grade.

3.2. The Impact of the RoSLA on Qualifications

As exposure to the RoSLA was determined by a threshold date of birth, we estimate the reform-induced shifts in the qualification rates using a regression dis-

[8]Since 1987, the first tier qualification label is the General Certificate of Secondary Education (GCSE).

continuity design (RDD), see e.g Hahn, Todd, and van der Klaauw (2001) and Imbens and Lemieux (2008).

Let w be an individual's date of birth, the running variable. There is a deterministic mapping from date of birth w to academic cohort $c \in \mathcal{C}$ and we normalize w such that $\mathbb{1}(w \geq 0)$ indicates RoSLA exposure. Let $z \in \mathcal{Z}$ denote an individual's (highest) academic qualification level, and let y^z be an indicator for attaining level z. Cast in the potential outcomes framework, $y^z = y_0^z \, \mathbb{1}(w < 0) + y_1^z \, \mathbb{1}(w \geq 0)$, where y_1^z and y_0^z are the potential outcomes with and without RoSLA exposure, respectively.

The RDD regression of interest is the linear probability model

$$y^z = \alpha_0 + t_0(w) + t_1(w)\mathbb{1}(w \geq 0) + \varphi \, \mathbb{1}(w \geq 0) + \epsilon, \tag{1}$$

where $t_0(w)$ and $t_0(w) + t_1(w)$ are the pre- and post-RoSLA trends, respectively, with $t_0(0) = t_1(0) = 0$. If the potential outcomes $E[y_0^z|w]$ and $E[y_1^z|w]$ are continuous functions of w, the OLS estimator of φ in (1) consistently estimates $E\left[y_1^z - y_0^z|w = 0\right]$, the causal effect of the RoSLA on the level-z qualification rate at the reform threshold $w = 0$.

We estimate (1) by gender and qualification levels $z \in \mathcal{Z}$ using the LFS data described above (Table 1). We include linear trends in the running variable w on either side of the threshold $w = 0$, as advocated by Gelman and Imbens (2019), and employ the Lee and Card (2008) corrected standard errors to account for specification errors arising from the LFS recording date of birth by month, an inherently discrete variable.

Figure 1 is a graphical rendition of the RDD and the estimated regression functions. It plots the observed proportion holding each level of qualification for each month of birth in our observation window. There is a sharp drop in the rate of holding no qualification for each gender at the RoSLA threshold, with a corresponding discontinuous increase in the proportion holding a basic qualification. The RoSLA did not have any effect on the rate of holding an advanced qualification.[9]

[9] A modest annual cyclical pattern emerges post-RoSLA with students born in the first five months of each academic year having an elevated likelihood of having no academic qualifications. This is due to the "Easter Leaving Rule" in force at the time under which students born between September to January were allowed to leave school at the end of the Spring term, about three months earlier than the rest of their academic cohort. The variation in educational attainment generated by this rule was minor relative to that generated by the RoSLA and largely faded in the early 1960s. Combining the cohort variation in attainment generated by the RoSLA with within-cohort variation generated by the Easter Leaving Rule to decompose the labor market return to education into the effect from holding a qualification and the effect of schooling length,

Figure 1: Distribution of Academic Qualifications by Month of Birth Relative to RoSLA Threshold

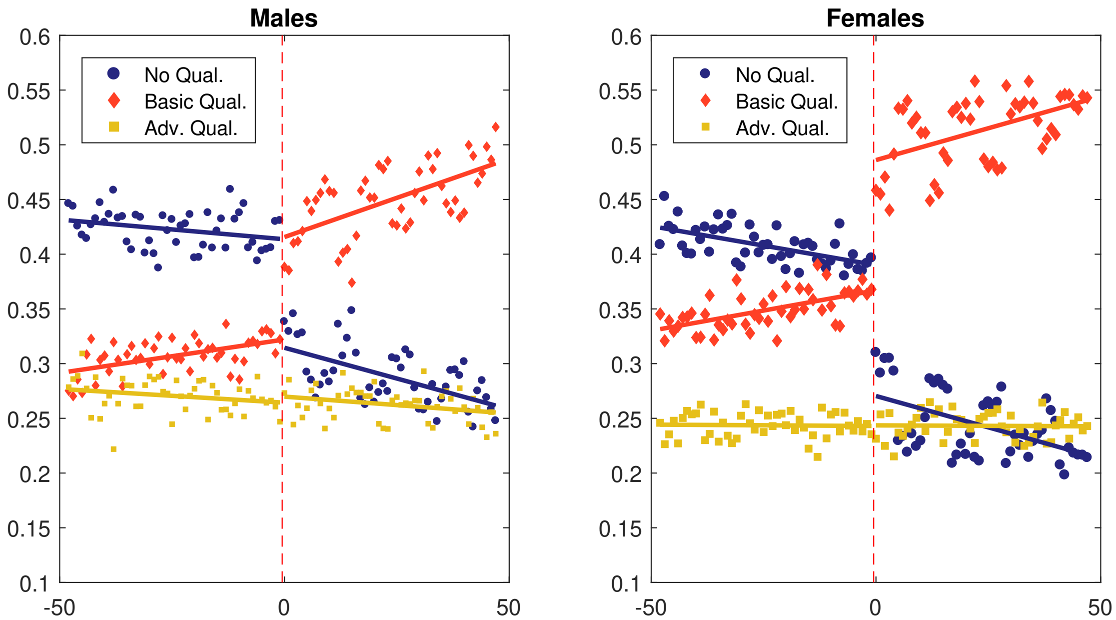

Notes: The horizontal axes show month of birth normalized to 0 at the RoSLA threshold. The vertical axes show qualification rates.

Table 2 reports estimates of φ in (1) by gender and qualification level $z \in \mathcal{Z}$, with and without the inclusion of demographic controls. The RoSLA reduced the fraction holding no qualification by about 10 percentage points for men and by about 12 percentage points for women, effect sizes that are very precisely estimated and in line with previous studies evaluating the RoSLA-effect on qualification rates.[10] The fact that the proportion of individuals leaving education at age 16 rather than at age 15 increased by 25 percentage points, but the qualification rate increased only by 10-12 percentage points leaves room for significant selection, an issue we turn to next.

3.3. RoSLA Effects by Ability: Evidence from the NCDS

All born in March, 1958, the NCDS cohort members were in the first RoSLA-treated academic cohort. The 1974 NCDS data sweep was conducted as the NCDS cohort members were turning 16, and a few months before they were due to sit

Dickson and Smith (2011) find that qualifications drive most of the returns to education. Their findings motivate our focus on qualifications as the key measure of educational attainment.

[10] See e.g. Chevalier, Harmon, Walker, and Zhu (2004), Dickson and Smith (2011), Grenet (2013), and Dickson, Gregg, and Robinson (2016). Grenet (2013) highlights a slightly larger effect of RoSLA on the qualification rate of women as compared to men, a feature that is also evident in Figure 1.

Table 2: Regression Discontinuity Estimates of the Impact of the RoSLA on Qualification Rates

	Males		Females	
	(i)	(ii)	(iii)	(iv)
No Qualification	−0.099*** (0.008)	−0.106*** (0.006)	−0.120*** (0.009)	−0.128*** (0.007)
Basic qualification	0.094*** (0.008)	0.106*** (0.005)	0.119*** (0.010)	0.131*** (0.007)
Advanced qualification	0.006 (0.005)	−0.000 (0.005)	0.001 (0.005)	−0.002 (0.004)
Controls	No	Yes	No	Yes
Observations	147,878	147,878	156,549	156,549

Notes: The sample used in each regression is described in Table 1. We report the estimated coefficient on a RoSLA dummy for being born September, 1957 or later from separate regressions with the dependent variable being a dummy for having a particular level of qualification attainment. Distance of date of birth from the September, 1957 threshold measured in months is used as the running variable and is included in linear form and interacted with the RoSLA dummy. The demographic controls include a third degree polynomial in age, month of birth dummies, and year of interview dummies. Standard errors are clustered on the running variable. *, **, and *** indicates statistical significance at the 10, 5, and 1 percent levels, respectively.

the O-level/CSE exams that would give them their first academic qualification. The 1974 NCDS questionnaire asked if they would have preferred to leave school at age 15 had that option still existed. We use that information here, along with information on actual qualification attainment (measured at age 33 in the 1991 NCDS data sweep) and, critically, measures of academic ability to study how the response to the RoSLA relate to ability.[11]

3.3.1. Preferred and Realized School-Leaving Age and Qualifications. The 1974 sweep asked the NCDS cohort two questions about preferred school-leaving age. First, "At what age do you think you are most likely to leave school?", with the lowest option being 16, consistent with this being the minimum school-leaving age for their cohort. Second, after being reminded that they were the first cohort required to stay in school until age 16, they were asked "In your case, do you wish that you could have left when you were 15?". From these two questions, we define an individual's *preferred school-leaving age* with options 15, 16 and 17+. The 1991 sweep, completed at age 33, collected information on the cohort members' *realized*

[11]The information on preferred and actual school-leaving age was recently used by Geruso and Royer (2018). However, the authors did not relate this information to measures of ability.

school-leaving age and *qualification attainment*.

Using this information, we classify the NCDS cohort members as:[12]

- Never-takers: individuals who preferred to leave school at age 15, and who, when leaving school at age 16, still left without any academic qualifications
- Compliers: individuals who preferred to leave school at age 15, but who, when leaving school at age 16, now left with a basic academic qualification
- Always-takers: individuals who preferred to stay until age 16 even if leaving at age 15 had been an option, and who left school with a basic qualification

Additionally we define as a separate class,

- Stayers: individuals who stayed past age 16 and gained an advanced qualification.

3.3.2. Ability. The NCDS includes measures of academic ability at age 16, comprising a reading comprehension test, a mathematics test, and teacher assessments in English, mathematics, and science. We interpret these as noisy indicators of an underlying latent academic ability, which we estimate using these observed measures (see Supplementary Appendix). The resulting ability measure is standardized.

Figure 2 shows the cumulative distribution of the NCDS ability measure for Never-takers, Compliers, Always-takers, and Stayers.[13] The distributions for Compliers and Always-takers are remarkably similar, indicating that individuals who were induced by the RoSLA to obtain a basic qualification (Compliers) have comparable academic ability to those who would have obtained the same qualification regardless (Always-takers). Additionally, the distribution for Stayers clearly stochastically dominates those of both Compliers and Always-takers, which in turn clearly stochastically dominate the distribution for Never-takers.

In summary, Figure 2 suggests that it is descriptively meaningful to characterize individuals whose qualification outcomes were influenced by the RoSLA (Compliers) as: (i) more academically able than the least able group (unqualified

[12]Material in the Supplementary Appendix verify that (i) preferred school-leaving ages in the NCDS align closely with the pre-RoSLA LFS qualification distribution, (ii) preferred school-leaving age aligns well with realized school-leaving age under the constraint that leaving at age 15 is not an option, and (iii) the NCDS qualifications closely resemble the LFS qualifications for the first RoSLA-affected cohort.

[13]Figure 2 pools data across genders. The Supplementary Appendix reports mean ability by gender and classification.

Never-takers), (ii) less able than the most academically able group (Stayers with advanced qualification), and (iii) similar in ability to those who would have obtained a basic qualification even without the reform (Always-takers). This pattern will inform the structure of our marriage market model, which categorizes individuals into three latent ability types—low, medium, and high—with the policy-induced qualification margin concentrated among those with medium ability.

Figure 2: Ability Distributions by Individual Classification in the NCDS

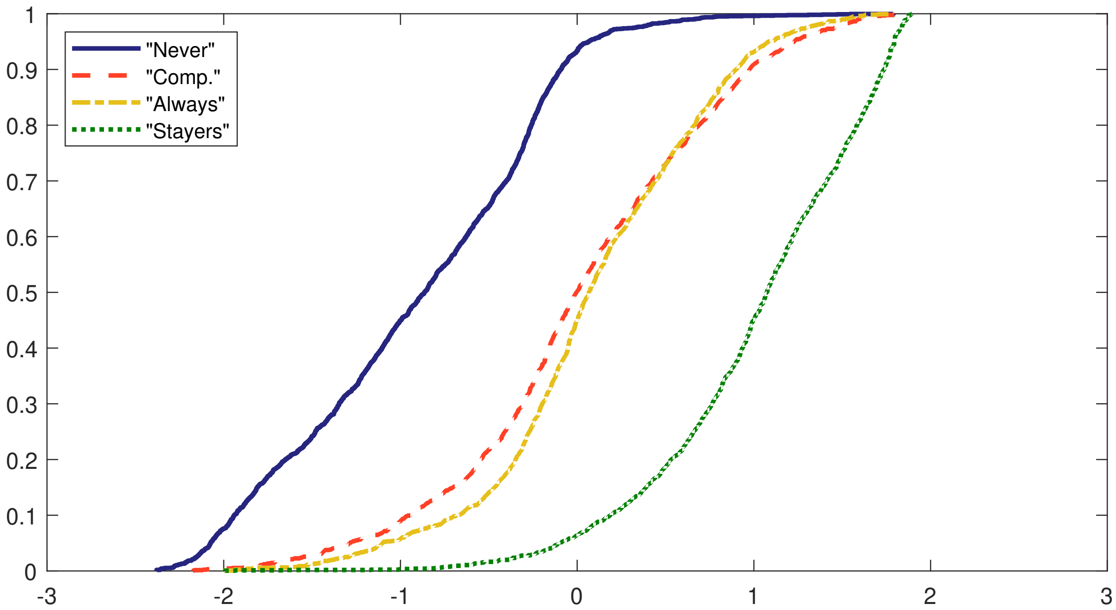

Notes: The sample size is 6,365 individuals from the NCDS, with 2,054, 1,329, 1,280, and 1,702 individuals in the Never-takers, Compliers, Always-takers, and Stayer classes. The measure of ability is constructed as a latent factor reflected in a reading comprehension test, a mathematics test, and teacher assessments in English, mathematics, and science at the age of 16. See the Supplementary Appendix for details.

4. Marriage Market Outcomes

The RoSLA induced a dramatic, permanent shift in the qualification distribution, reducing the likelihood of leaving school with no academic qualifications, while increasing the likelihood of leaving with a basic qualification. In addition, individuals who as a consequence of RoSLA left school with a basic academic qualification were positively selected in the sense that they were the most able among the individuals who, absent the RoSLA, would have left school without a qualification. We now document how these population shifts impacted the marriage market.

4.1. Never-Married Rates

Figure 3 shows the never-married rate (by age 45) by gender, cohort, both without conditioning on qualification and by level of qualification, computed from census data. The never-married rates are increasing across cohorts overall and for all qualification levels. The never-married rates of men are also above those of women. There is no discernible impact of the RoSLA on the unconditional never-married rates, nor on the never-married rates of basic and advanced qualified men and women. However, it is evident that the never-married rate of unqualified men and women increases discretely and permanently at the RoSLA threshold and that the jump is accompanied by an increase in the trend.

For women, RoSLA marks the clear reversal in the historical relationship between qualifications and marriage: prior to the reform, more educated women were less likely to marry—partly due to marriage bars in certain professions—but the first RoSLA-affected cohorts mark the point at which unqualified women become the least likely to ever marry. For men, the unqualified were already less likely to marry before the reform, but the RoSLA induces a distinct and persistent widening of this gap. While the relationship between education and marriage has been studied extensively, the pinpointed patterns we document around the RoSLA appear to be new to the literature, emerging from our focus on a narrow set of cohorts around a sharp policy threshold.

A naive regression discontinuity estimate on the two sets of eight gender-specific data points in Figure 3 suggests the never-married rate jumps at the threshold by about 2.5 percentage points for unqualified men and 2 percentage points for unqualified women.[14] Month of birth in the LFS make that data better suited for an RDD. However, the restriction to individuals aged 45 or above substantially reduces the LFS sample sizes. In the Supplementary Appendix we present the results from estimating an RDD specification, by gender, where we, in order to retain precision, impose common trends across qualification groups on either side of the RoSLA threshold. The resulting LFS estimated discontinuities are fully consistent with those suggested by the census data, indicating that the never-married rates increased at the RoSLA threshold specifically for unqualified men and women.

Using the 2000 NCDS sweep (at age 42), the Supplementary Appendix verifies that there is no difference in the never-married rate of "Compliers" and

[14] These RoSLA impact estimates are obtained in a population size weighted RDD regression with differential linear trends on either side of the RoSLA threshold, ignoring clustering in the running variable.

Figure 3: Never-Married (by Age 45) Rates by Cohort, Gender, and Qualification

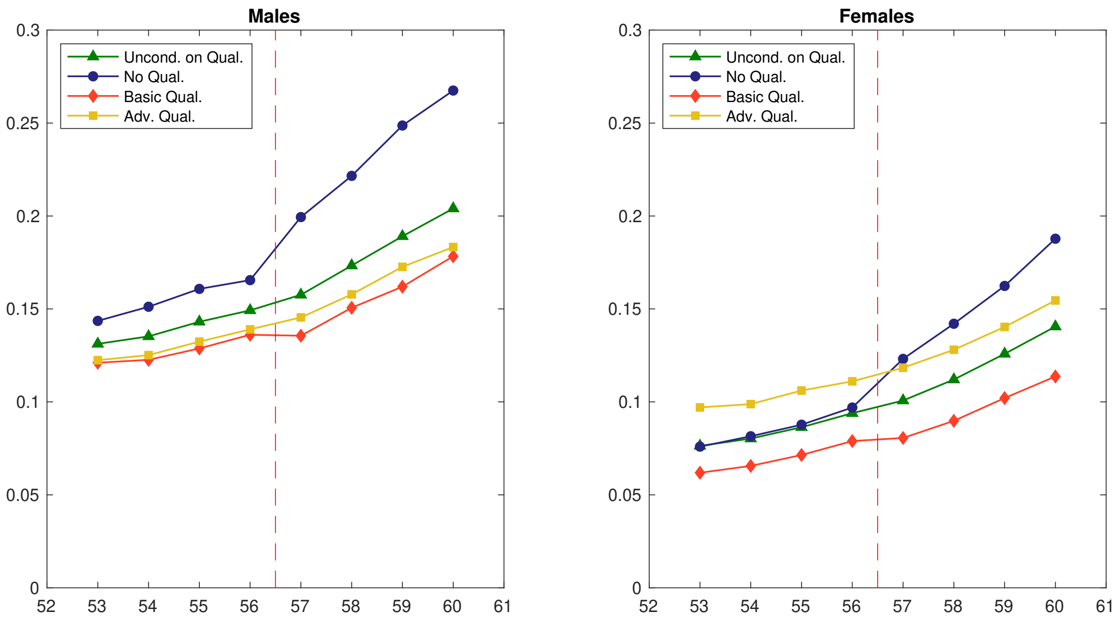

Notes: The horizontal axes show academic cohorts. The vertical axes show never-married (by age 45) rates.

"Always-takers". Both groups attained a basic qualification as the NCDS cohort was RoSLA-treated, and the finding supports the notion that "Compliers" and "Always-takers" not only had the same ability, but also the same post-reform marital outcomes.

4.2. Qualification Homogamy

Our LFS sample of married individuals features strong qualification homogamy. Figure 4 shows that, pooling across cohorts, for each gender and qualification, the qualification most frequently held by the spouse is the same as that held by the respondent.

To descriptively explore how marital sorting on qualification changed around the RoSLA we use a qualification homogamy index used also by Eika, Mogstad, and Zafar (2019). Let $g \in \{m, f\}$ index gender, where m stands for "male" and f for "female", and use prime to indicate own variables, and double-prime for the spouse. Then, define

$$S^g(z, c) \equiv \frac{\Pr^g(z' = z, z'' = z | c' = c)}{\Pr^g(z' = z | c' = c) \Pr^g(z'' = z | c' = c)}, \quad g = m, f, \tag{2}$$

Figure 4: Qualification Homogamy by Gender

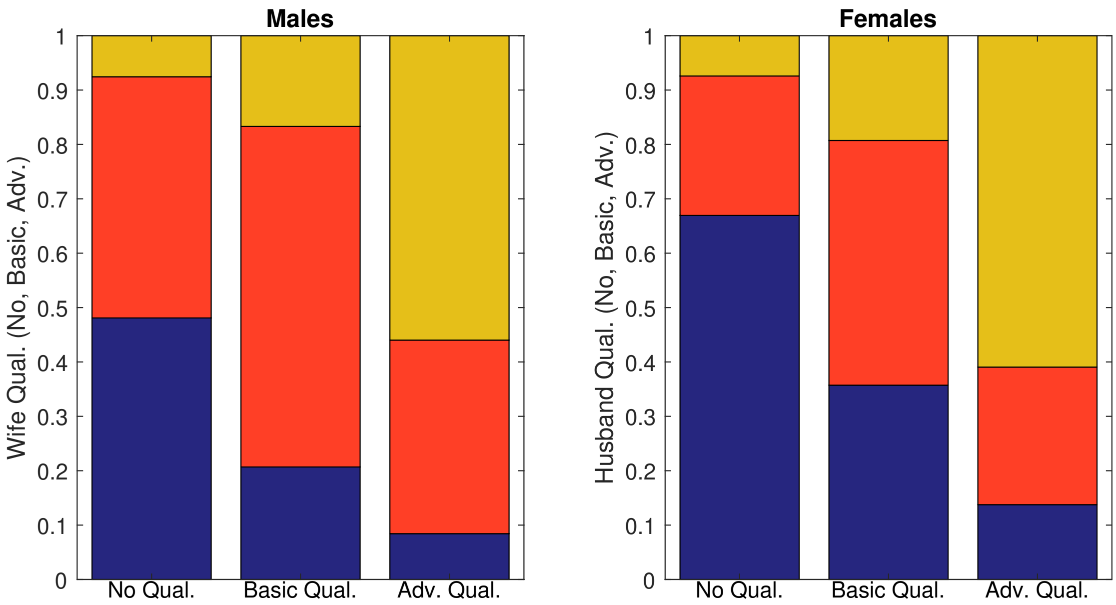

Notes: Each panel shows the spousal qualification distribution by own qualification level.

for married individuals of gender g and cohort c. The numerator is the probability that, for a randomly drawn individual in this subpopulation, both the individual and their spouse have qualification z. The denominator is the product of two marginal probabilities: that the individual has qualification z, and that the spouse has qualification z. Under random matching, $S^g(z,c)$ equals unity; it exceeds unity when it is more common for both partners to hold qualification z than random matching would imply. A high $S^g(z,c)$-value will reflect a high propensity to marry an equally-qualified partner, but will also reflect small marginals resulting from small population shares or high single rates. As the index does not disentangle these mechanisms we deploy it merely as a descriptive device to pinpoint changes in marriage market outcomes around the RoSLA that our structural model should be ale to replicate.

We compute $S^g(z,c)$ by qualification, gender and cohort using the LFS sample of married individuals. The cohort profiles of the homogamy indices are shown in Figure 5 for married men (left panel) and for married women (right panel). The aggregate index—which aggregates across qualifications—is flat for both genders, but masks considerable qualification heterogeneity.

First, the advanced qualified always has the highest index value and the basic qualified the lowest, with the unqualified in-between. Given the caveat above, we do not over-interpret this level ranking: the advanced qualified are the smallest of the three groups in the married population and the basic qualified the largest (Table 1), so the ordering is what the smaller marginals alone would tend to produce.

Second, there is a permanent increase in the homogamy index after the RoSLA for unqualified men and women, alongside a modest decline for the basic qualified. As both movements coincide with reform-induced shifts in the qualification marginals each may reflect a combination of mechanical and behavioral adjustments.[15] Our structural model later interprets them as reflecting, in part, increased marriage between low-ability individuals and increased mixing between medium- and high-ability individuals.

Figure 5: Qualification Homogamy by Cohort, Gender, and Qualification

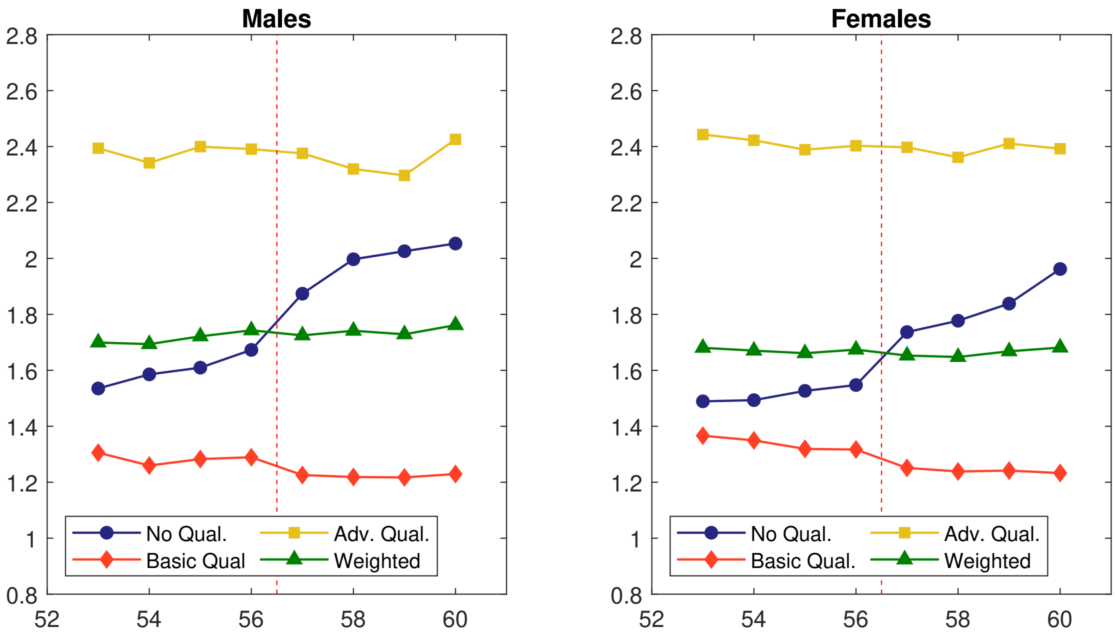

Notes: The horizontal axes show academic cohorts. The vertical axes show the homogamy index, see (2).

Following Chiappori, Costa Dias, Meghir, and Zhang (2025), we also report the odds-ratio measure of assortative matching and its change from pre- to post-RoSLA (Supplementary Appendix Table S.6). The odds ratio has the advantage of being invariant to changing margins in the supply of qualification types, which

[15]The unqualified become a smaller share of the married population, due to the higher qualification rate (Figure 1) and their reduced likelihood of marriage (Figure 3), while the basic qualified become a larger share.

makes it well suited to isolating changes in sorting from the supply shift induced by the reform. We apply it to two binary partitions: (i) the "extensive" qualification margin (no qualification vs. some qualification, either a basic or advanced), and (ii) the "intensive" qualification margin (advanced qualification vs. either a basic or no qualification). The point estimates move in opposite directions across the two margins. On the extensive margin the odds ratio rises from pre- to post-RoSLA (by about 0.5), consistent with the unqualified becoming relatively more isolated in the marriage market; on the intensive margin it falls (by about 0.56), consistent with marriages between advanced- and basic-qualified individuals becoming more common as the supply of basic-qualified individuals increased. Both changes are, however, imprecisely estimated: the 95% confidence intervals for the pre-to-post change narrowly include zero in each case. We therefore read the odds-ratio evidence as suggestive corroboration of the direction of the homogamy-index movements rather than as independently conclusive.[16]

4.3. Age Gaps

The husband-wife age gap is the difference in their academic cohorts, that is $d \equiv c' - c$, where c and c' are the cohort of the husband and the wife, respectively. The age gap distribution is skewed towards positive gaps where the husband is older than the wife. Indeed, the most common age gaps are 1 year, 2 years and 0 years (see Supplementary Appendix).

To explore whether the age gap distribution was affected by the RoSLA, we consider how the cohort-specific distributions differs from the aggregate distribution in the key cohorts around the reform threshold. For each gender, each age gap $\delta \in \{-3, -2, ..., 3\}$ and for each cohort $c \in \mathcal{C}$ we regress the indicator $\mathbb{1}\left(d = \delta\right)$ on a cohort-dummy for being from cohort c. This way we determine, for each gap in this central range, whether men and women from cohort c had a different likelihood of being married with this particular age gap compared to men and women from all other cohorts in $\mathcal{C}$. The coefficient estimates from these regressions are rendered graphically in Figure 6.

[16]Under the Choo and Siow (2006) framework types differ only in observable qualification, changes in the odds ratio can be interpreted as changes in the surplus matrix. Our analysis instead exploits the narrow reform window to motivate the assumption of a locally stable surplus matrix and to attribute the observed change in the observed assortative matching to the reform-induced change in the relationship between observed qualifications and latent ability. The reduced odds ratio on the intensive margin is also consistent with the reduced-form evidence in Geruso and Royer (2018) that indicates that some of those who gained a basic qualification due to the RoSLA responded also by marrying qualified partners at a rate higher than mechanically implied if they were to marry an unchanged set of partners.

The top row of Figure 6 shows the results for men in each cohort 1956 to 1958. In each sub-figure, a vertical line has been added to delineate when the wife is drawn from a pre- versus post-RoSLA cohort. The first RoSLA affected cohort, that is the 1957 cohort, is the only cohort with statistically significantly different age gap frequencies, being about one percentage point more likely to be married with an age gap of either 0 or +1. Conversely, they were less likely to be married with a negative age gap.

The bottom row of Figure 6 shows the corresponding results for women. The vertical lines in this case delineate when the husband is drawn from a pre- versus post-RoSLA cohort. Consistent with the findings for men, the most notable deviations are for early post-reform women—born in the 1957 and 1958 academic cohorts—who were about one percentage point more likely to marry with an age gap of 0 and +1 respectively (thus both marrying 1957 cohort men).

The evidence thus suggests that the RoSLA led to a temporary and quantitatively modest shift in the age-gap distribution, with early RoSLA-affected men and women slightly more likely to marry each other. This is inconsistent with qualifications playing no role in the marriage market, and underlines the importance of accounting for how qualifications shape matching patterns among those who marry.[17] At the same time, this evidence pertains to sorting among those who do marry and says nothing directly about whether qualifications affect the probability of marrying at all—a distinct question we address when we estimate the effects of qualifications and ability on the probability of ever marrying.

5. Model

Our theoretical analysis builds on Choo and Siow (2006)'s empirical framework for the transferable utility matching model of Shapley and Shubik (1971) and Becker (1973).

5.1. *The Choo-Siow Framework*

Assume continuum sets of men and women, each of finite Lebesgue measure. Each individual belongs to one of a finite set of systematic types $\mathcal{X}$, and their systematic type is publicly observed by all potential partners. Each individual must choose a type of partner in $\mathcal{X}_+ = \mathcal{X} \cup \{0\}$, where choosing 0 corresponds to remaining single. A match between a type-x man and a type-y woman results in a systematic match surplus Σ_{xy}. Man i is also subject to a vector of preference shocks, $(\varepsilon_i(y))_{y \in \mathcal{X}_+}$,

[17] Geruso and Royer (2018) conduct a similar exercise with largely consistent findings.

Figure 6: Deviations from the Aggregate Age Gap Distribution by Cohort and Gender

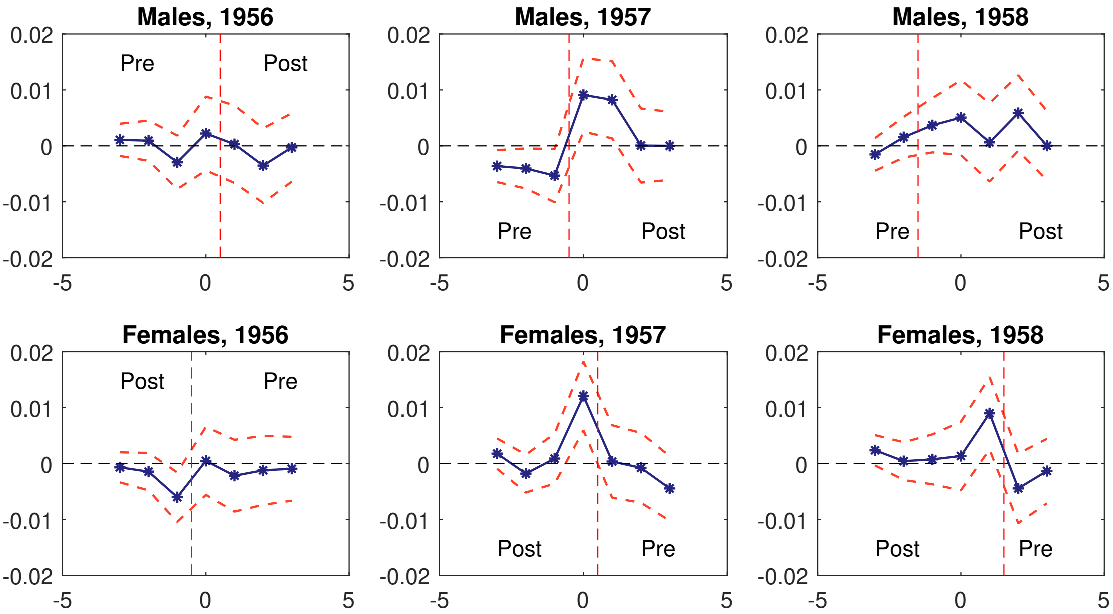

Notes: The horizontal axes show age gap $\delta \in \{-3, -2, ..., 3\}$. The vertical axes show regression coefficients from a regression of $\mathbb{1}(d = \delta)$ on a cohort-dummy for being from cohort c. The vertical dashed lines in the top panels delineate when the wife is drawn from a pre- versus post-RoSLA cohort. The vertical dashed lines in the bottom panels delineate when the husband is drawn from a pre- versus post-RoSLA cohort.

where $\varepsilon_i(y)$ is an idiosyncratic incremental payoff (to man i) from matching with *any* woman of type $y \in \mathcal{X}$ (or remaining single if $y = 0$). Similarly, woman j is subject to a vector of shocks $(\varepsilon_j(x))_{x \in \mathcal{X}_+}$ over male systematic types. The total surplus generated by a match between man i of type x and woman j of type y is

$$\Sigma_{xy} + \varepsilon_i(y) + \varepsilon_j(x). \tag{3}$$

The preference heterogeneity components $\varepsilon_i(y)$ and $\varepsilon_j(x)$ are separable, entering (3) additively. They depend only on the partner's type, not their individual identity. Transferable utility implies that the total surplus can be divided between the two parties in any way the couple chooses.

Under the assumption that the systematic type of every individual is observed by all market participants, existence and essential uniqueness of a stable matching is guaranteed (Decker, Lieb, McCann, and Stephens, 2012).[18] Chiappori, Salanié,

[18]Note that existence and essential uniqueness applies equally to our generalized model with a latent ability, since we shall assume throughout that market participants observe the types of

and Weiss (2017) show that, under separability, a stable matching is characterized by a pair of vectors $(U_{xy}, V_{xy})_{x,y\in\mathcal{X}}$ such that for each $x, y \in \mathcal{X}$, $U_{xy} + V_{xy} = \Sigma_{xy}$. U_{xy} specifies the systematic component of the payoff to a type-x man when he matches with a type-y woman while V_{xy} specifies the corresponding payoff to the woman. The systematic payoff from remaining single, denoted by U_{x0} and V_{0y}, is normalized to zero. Each man i of type x maximizes his overall payoff $U_{xy} + \varepsilon_i(y)$ by his choice of $y \in \mathcal{X}_+$, and each woman j of type y maximizes $V_{xy} + \varepsilon_j(x)$ by her choice of $x \in \mathcal{X}_+$. As a consequence of these choices, marriage markets clear: for each pair (x, y), the measure of type x men who choose to marry type y women equals the measure of type y women who choose to marry type x men.

Let $\mu^m_{y|x}$ denote the equilibrium probability that a man of type x chooses $y \in \mathcal{X}_+$, and similarly, let $\mu^f_{x|y}$ denote the equilibrium probability that a woman of type y chooses $x \in \mathcal{X}_+$. The preference shocks $\varepsilon_j(x)$ and $\varepsilon_i(y)$ are assumed to be i.i.d. Type I Extreme Value distributed,[19] such that

$$\log\left(\frac{\mu^m_{y|x}}{\mu^m_{0|x}}\right) = U_{xy}, \text{ and } \log\left(\frac{\mu^f_{x|y}}{\mu^f_{0|y}}\right) = V_{xy}. \tag{4}$$

Decker, Lieb, McCann, and Stephens (2012) show that the marriage market equilibrium exhibits natural comparative statics.

In large marriage markets, where the Law of Large Numbers hold, the population frequencies in (4) can be replaced by their sample counterparts. Hence, if types are observable to the econometrician, U_{xy} and V_{xy}, and consequently also Σ_{xy}, are identified.[20]

The identification argument in the standard CS framework relies on the systematic types of the marriage market participants being observable to the econometrician, and thus restricts attention to marital sorting on observables. This is unfortunate, not least because the empirical marriage market literature typically considers only relatively low-dimensional systematic type-spaces, leaving plenty of

all potential partners.

[19]The Extreme Value distribution assumption is standard in the literature. A notable exception is Galichon and Salanié (2021), who show that the Choo-Siow framework remains tractable with minimum distributional assumptions.

[20]We opt to characterize the equilibrium in terms of type-conditional choice probabilities as this allows a more direct discussion of how choice frequencies vary across types and across policy regimes. More conventionally, we could characterize the equilibrium directly in terms of measures: let N_{xy} denote the measure of matches between men of type x and women of type y, and let N^m_{x0} and N^f_{0y} denote the measures who remain single, of type x men and type y women respectively. In this alternative notation, (4) would take the form $\log(N_{xy}/N^m_{x0}) = U_{xy}$ and $\log\left(N_{xy}/N^f_{0y}\right) = V_{xy}$.

room for latent traits to confound the analysis.[21]

The focus on low-dimensional observable type-space also yields stark comparative static predictions for marriage market equilibrium responses to changes in the population distribution resulting from a large education reform like the RoSLA. Consider for instance a decrease in the supply of males of observable type x, say, men with no academic qualifications. The CS framework stipulates that men of this type will obtain a larger share of the marital surplus from a marriage to any type of woman. As a result, type x men are predicted to marry at a higher rate, i.e. the never-married rate $\mu^m_{0|x}$ is predicted to fall. Of course, the never-married rate of the relatively scarce post-RoSLA unqualified men markedly increased at the reform threshold, cf. Figure 3.[22] These counterfactual predictions, along with the NCDS evidence presented above, motivate extending the model to include a latent marriage market trait, namely ability.

5.2. A Model with Latent Ability

In our empirical setting, an individual's type has three dimensions. First, by their date of birth, an individual belongs to an academic cohort $c \in \mathcal{C}$. Second, they are endowed with ability $a \in \mathcal{A}$. Finally, they hold an academic qualification, $z \in \mathcal{Z}$. Hence, an individual's type is a triple $(c, a, z) \in \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$. There are three relevant qualification levels, $\mathcal{Z} = \{z^0, z^1, z^2\}$, referred to as no qualification, a basic qualification, and an advanced qualification, respectively. Informed by the NCDS evidence in Figure 2, we consider a parsimonious specification with three ability levels, $\mathcal{A} = \{a^0, a^1, a^2\}$, which we refer to as low, medium, and high ability, respectively. The academic cohorts in $\mathcal{C}$ are split into the pre- and post-RoSLA cohorts, $\mathcal{C}_0 = \{1953, ..., 1956\}$ and $\mathcal{C}_1 = \{1957, ..., 1960\}$.

An individual's systematic type is observed by prospective partners in the marriage market. Ability, however, is a latent trait, i.e. unobserved by the econometrician, and the identification arguments referenced above therefore no longer apply. Identification of the extended CS model with latent ability will instead come from the fact that the RoSLA created a discontinuity in qualifications between cohorts, but not in ability. Hence, as a first step we impose an assumption of continuity in the distribution of the latent ability, including across the RoSLA

[21] Notable exceptions include Chiappori, Oreffice, and Quintana-Domeque (2012) and Dupuy and Galichon (2014) who exploit exceptionally rich datasets to analyze marital sorting with a high-dimensional systematic type-space.

[22] The RoSLA of course induced shifts in the population distributions for both genders. Nevertheless, we confirm below that the comparative statics intuition is useful for understanding the marriage market responses to the RoSLA.

threshold. Whilst the distribution of ability may change over time, for instance through improvement in the quality of primary- and lower level secondary schooling at the national level, there is no reason to expect a discontinuity in this process at the 1957 cohort. Recall that w is date of birth, see (1), and let $\Pr^g(a|w)$ be the proportion of individuals of gender g, born at time w, who have ability level $a \in \mathcal{A}$. We then make the following assumption:

Assumption 1. $\Pr^g(a|w)$ is continuous in birth date w for ability $a \in \mathcal{A}$, and gender $g \in \{m, f\}$.

As the RoSLA sharply increased the proportion holding a basic qualification rather than no qualification, we focus on modeling a shift in the population proportions with $z = z^0$ and $z = z^1$. The NCDS ability distributions in Figure 2 suggest that the shift took place within the subpopulation of individuals with medium ability, i.e. with $a = a^1$.

We assume that having low ability $a = a^0$ precludes the attainment of even a basic qualification, thus always leaving them unqualified, $z = z^0$. For an individual with medium ability $a = a^1$, a basic qualification $z = z^1$ (but not an advanced qualification $z = z^2$) would be attainable, but whether they pursue it or not would depend on the opportunity cost of doing so. Individuals with high ability, $a = a^2$, are assumed to always find it worthwhile to attain an advanced qualification, $z = z^2$.

To obtain a basic qualification, $z = z^1$, an individual must (i) attend school up to and through the academic year in which they turn 16, and (ii) pass the exams held at the end of that year. While passing the exams generally depends upon ability, attending school entails an opportunity cost, an important component of which is foregone labor earnings. Pre-RoSLA, an individual would forego labor income if they attended school beyond age 15, but paid employment ceased to be an option for a 15-year old in the post-RoSLA regime.[23] Hence, the reform reduced the opportunity cost of acquiring a basic qualification, $z = z^1$, without altering the ability required to so. In this setup, consistent with the NCDS evidence in Figure 2, the RoSLA impacted the qualification rate of medium ability individuals, $a = a^1$, who would be weighing the costs and benefits of sitting the exams for the basic qualification level.

[23] Indeed, the compulsory school leaving age and the minimum age requirements for paid employment are statutorily linked: the 1933 Children and Young Persons Act for England and Wales sets out the conditions under which young people can work and essentially prohibits anyone working full time whilst they are at school.

Assumption 2 formalizes the above by parameterizing the proportions of individuals of medium ability $a = a^1$, gender g, and cohort $c \in \mathcal{C}_k$, by RoSLA-regime $k = 0, 1$, who attains a basic qualification:

Assumption 2. Any individual of low ability, $a = a^0$, is always unqualified, $z = z^0$, and any individual with high ability, $a = a^2$, always attains an advanced qualification, $z = z^2$. A gender g and cohort $c \in \mathcal{C}_k$ individual of medium ability, $a = a^1$ attains a basic qualification, $z = z^1$, with probability γ_k^g and remains unqualified, $z = z^0$, with complementary probability $1-\gamma_k^g$, for $k = 0, 1$, indicating pre- and post-RoSLA cohort respectively, where $\gamma_1^g > \gamma_0^g$.

Figure 7 illustrates the $\mathcal{A} \to \mathcal{Z}$ mapping from Assumption 2. Note that there are four possible ability-qualification profiles: (a^0, z^0), (a^1, z^0), (a^1, z^1) and (a^2, z^2). Within this set, there is variation in ability (low versus medium) among the unqualified, and there is variation in qualifications (unqualified versus basic) among medium ability individuals. Let $\mathcal{X} \subset \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$ be the set of full types that can arise and note that $|\mathcal{X}| = 32$.[24]

Figure 7: Assumed Relation between Ability and Qualification by Education Regime

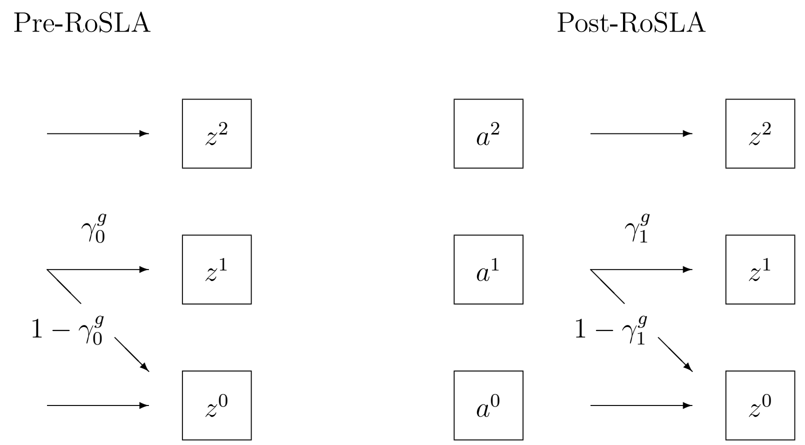

As the reform eliminated the foregone earnings cost of acquiring a basic qualification there is a strong *ex ante* reason to believe that, after the RoSLA, those

[24]Two further remarks are in order. First, Assumption 2 implicitly stipulates that whether or not a medium-ability individual gained a basic qualification was triggered by some further unobserved heterogeneity, which is unrelated to potential marital surplus. Second, restricting γ_0^g and γ_1^g to be constant across cohorts within each respective regime is essential for our use of Assumption 2 to recover population type distributions for all cohorts, and it is tenable due to the short observation window.

who had the ability to gain a qualification would do so. This would imply that γ_1^g would be unity and we will show below that this hypothesis is not rejected by the data. The significance of this finding will be clear just below where we use a simple(r) 2 cohorts $\times$ 2 ability-types $\times$ 2 qualification-types specification to argue that, specifically for the case where $\gamma_1^g = 1$, $g = m, f$, the marital surplus matrix Σ_{xy} is identified.

For identification, we impose the following assumption on the systematic component of the marital surplus matrix Σ_{xy}. Assumption 3, restricts Σ_{xy} to be additively separable in ability-qualification and cohort profile.

Assumption 3. The systematic surplus function Σ_{xy} is additively separable,

$$\Sigma_{(c,a,z),(c',a',z')} = \zeta\left(a, z, a', z'\right) + \lambda\left(c, c'\right).$$

Assumption 3 imposes that the systematic component of marital surplus, ζ, is stable across cohorts and thus invariant to the RoSLA: the reform affected marital outcomes through changes in the distribution of types rather than through shifts in the marriage market value of ability and qualifications. This assumption is not innocuous. In principle, the reform could have altered the economic returns to qualifications—for example by sharpening the negative ability signal that having no qualifications sends to the labor market—which would in turn affect marital returns. We provide empirical support for Assumption 3 in the Supplementary Appendix. There we revisit and extend the analysis of Del Bono and Galindo-Rueda (2006) using the "Easter Leaving Rule" which generated limited but exogenous *within-cohort* variation in qualification attainment both *before and after* the RoSLA, to gauge whether labor market returns to a basic qualification changed across regimes.[25] We find no evidence of any systematic change in the labor market returns to a basic qualification after the RoSLA. This suggests that the reform operated through changes in selection into qualifications rather than through a revaluation of qualifications in the labor market.[26]

[25] The Easter Leaving Rule (see footnote 9) created within-cohort differences in qualification attainment between individuals born September–January and those born February–August, both before and after the RoSLA. We exploit this variation to estimate reduced-form qualification, employment and wage discontinuities at the January–February threshold within each regime.

[26] The assumed separability between ability-and-qualifications on the one hand and cohort profile on the other is also directly testable in the post-RoSLA regime when $\gamma_1^g = 1$ as it implies a double-difference restriction on marriage frequencies. The null is rejected in only two of 36 cases (see Supplementary Appendix section S.4.5 for details). This test however does not speak to issue of stability across the regimes.

5.3. Identification in the $2 \times 2 \times 2$ Case

Consider the $2 \times 2 \times 2$ case with one pre-RoSLA cohort c^0 and one post-RoSLA cohort c^1, two ability types $\{a^0, a^1\}$, and two qualification types $\{z^0, z^1\}$. The ability-qualification mapping is as in Figure 7, but restricted to $\gamma_1^g = 1$ and ignoring the (observable) advanced-qualified high-ability type. In this case, there are five types for each gender, and thus 25 couple type-profiles plus 10 singles rates. However, since ability is latent, two types are not directly observable, or distinguishable, namely those without qualifications in the pre-RoSLA cohort, who may be of either low ability or medium ability. Due to the presence of these indistinguishable types the majority of the matching and single rates—namely, 20 out of 35—are not directly observable. Our main methodological contribution is the following identification theorem:

Theorem 1. *Suppose there is one pre-RoSLA cohort c^0 and one post-RoSLA cohort c^1, two ability levels $\{a^0, a^1\}$, and two qualification levels $\{z^0, z^1\}$. Under Assumptions 1–3, specialized to the $2 \times 2 \times 2$ environment with $\gamma_1^m = \gamma_1^f = 1$, the complete matching pattern is identified. Consequently, the systematic surplus matrix Σ, and the systematic match values U and V, are identified.*

Proof. See Appendix. □

The $2 \times 2 \times 2$ model has 5 types of each sex, three of which are distinguishable, while two are indistinguishable. The three distinguishable types are individuals with a basic qualification in either cohort, and unqualified individuals in the post-RoSLA cohort—who, by Assumption 2 with $\gamma_1^m = \gamma_1^f = 1$, must be of low ability. The two indistinguishable types are the unqualified in the pre-RoSLA cohort, who may be of low or medium ability. Thus, we observe 20 entries in the matching matrix. The proof of Theorem 1 shows that it is possible to identify the complete matching pattern, i.e. all 35 entries in the matching matrix.

The proof proceeds as follows. First, matches between distinguishable types identify the cohort component of surplus. This allows the observed post-reform low-ability group to serve as an anchor for recovering the latent pre-reform low-ability group. The assumption that the distribution of ability is the same across cohorts identifies the overall size of the latent group in the pre-reform cohort. These steps lead to a single non-linear equation in a single unknown that has a unique solution, enabling us to recover the latent matching frequencies. Thus we identify the complete matching matrix, from which identification of the surplus matrix follows exactly as in Choo and Siow (2006).

The main contribution of this theorem lies in identifying an equilibrium matching model when a policy reform changes the mapping between an observable characteristic and a latent trait, while leaving the latent trait distribution unchanged.

We note that it is essential for the identification result that individuals from the pre- and post-RoSLA regimes interact in a common marriage market. This highlights that our empirical approach of modeling a common marriage market encompassing a policy reform is distinctly different to approaches that study how marriage patterns change between earlier and later cohorts that operate in separate marriage markets.

Finally, a paradoxical observation: the fact that the mapping between ability and qualifications is not one-to-one in the pre-RoSLA cohort presents a challenge as well as an opportunity. The challenge is that posed by latent types. However, by overcoming this challenge, we are able to decompose the roles of ability and qualifications in the marriage market. If we had data only on post-RoSLA cohorts, then we would be able to infer the ability of each individual, but we would not be able to separate out the roles of the two variables in the marriage market.

5.4. Empirical Specification

Our complete empirical specification enriches the $2 \times 2 \times 2$ version considered in Theorem 1 by (i) expanding the type-space to include the high ability type, $a = a^2$, who, per Assumption 2, always gains an advanced academic qualification, $z = z^2$, (ii) extending the observation window to eight cohorts, four on either side of the reform threshold, (iii) allowing for trends in the ability distributions in accordance with Assumption 1, and (iv) allowing $\gamma_1^m \leq 1$ and $\gamma_1^f \leq 1$.

Introducing the (a^2, z^2)-type is empirically important as the advanced qualified is a sizeable subpopulation who occasionally marry partners with basic or no qualification. However, their inclusion raises no identification issues as they are, per Assumption 2, an observable type.

Including additional cohorts raises two issues. First, our RDD analysis in Section 3.2 identifies the increase in the basic qualification rate only at the reform threshold. Cohort variation in γ_k^g, within regime $\mathcal{C}_k$, $k = 0, 1$, is not distinguishable from variation in the frequency of ability levels a^1 and a^0. Assumption 2 makes clear that we treat γ_k^g, $g = m, f$ as regime-specific cohort-invariant parameters (note, however, that the ability distribution trends continuously by cohort to exactly match the observed qualification distributions by gender and cohort). Second, there are evident trends in never-married rates (see Figure 3). Ignoring such trends risk mis-characterizing discontinuities at the RoSLA threshold. Furthermore, with $|\mathcal{C}|^2 = 64$ possible cohort-profiles, it is natural to restrict how marital surplus depends on the husband-wife cohort profile (c, c'). For these reasons, we postulate the following surplus structure,

$$\Sigma_{(c,a,z),(c',a',z')} = \zeta\left(a, z, a', z'\right) + \lambda\left(c' - c\right) + \tau^m\left(c, a\right) + \tau^f\left(c', a'\right). \tag{5}$$

where, with a slight abuse of notation, (5) redefines $\lambda(\cdot)$ as a function of the husband-wife age gap $c' - c$ only, and includes trends, $\tau^g(c, a)$ for $g = m, f$, that are gender- and ability-specific. We specify $\lambda(\cdot)$ as fully non-parametric at age-gaps around zero, impose linear effects of age gap outside the range $\{-3, -2, ..., 3\}$, and normalize $\lambda(0) = 0$. Marital surplus trends are continuous piecewise linear functions of academic cohort with a single knot at the RoSLA threshold.[27]

Empirically, a proportion of individuals in our sample marry spouses from out-of-sample cohorts (pre -1953 or post-1960). To accommodate this, we extend $\mathcal{X}_+$ slightly to $\mathcal{X}_+ \equiv \mathcal{X} \cup \{0, \text{pre}, \text{post}\}$, where "pre" and "post" indicate marriages to partners born prior to 1953 and after 1960, respectively. We model the probability of marrying an out-of-sample-cohort spouse by own gender and cohort to approximate the age gap function $\lambda(\cdot)$, thereby accounting for such marriages in a model-consistent way.[28]

Finally, allowing for the possibility that $\gamma_1^m \leq 1$ or $\gamma_1^f \leq 1$ seems prudent from an empirical point of view, but as we shall see, the case $\gamma_1^m = \gamma_1^f = 1$ covered in Theorem 1 in fact offers the best fit to the data.

6. Estimation

With four (a, z) combinations, $\zeta(a, z, a', z')$ is a 4×4 matrix of 16 parameters. The parameterized age-gap preferences and trends add 22 parameters. We stack these 38 unknown systematic surplus parameters in $\boldsymbol{\theta}$. We must also handle $\boldsymbol{\gamma} \equiv (\gamma_0^m, \gamma_1^m, \gamma_0^f, \gamma_1^f)$, the proportions of medium ability men and women who gain basic qualification pre- and post-RoSLA. We do so by estimating, in a first step, the likelihood function concentrated with respect to $\boldsymbol{\theta}$ to obtain the model's fit to the data as a function of these proportions.

We aggregate the equilibrium in full types $x, y \in \mathcal{X}$ to characterize the implied

[27] The inclusion of ability-specific trends formally violates the additive structure in Assumption 3, but they are identified from observable trends by qualification and gender, and crucially, they facilitate the identification of the RoSLA threshold discontinuities. Details of the specification of $\lambda(\cdot)$ and $\tau^g(\cdot, a)$ are provided in the Supplementary Appendix.

[28] Details of the approach can be found in Section S.2 of the Supplementary Appendix.

choice frequencies in terms of observable types. The aggregation is done according to the distribution of ability a conditional on gender g, cohort c and qualification z, $\Pr^g(a|c,z;\boldsymbol{\gamma})$, which depends on $\boldsymbol{\gamma}$. Let $\widetilde{x},\widetilde{y}\in\widetilde{\mathcal{X}}$ denote an observable type, where $\widetilde{\mathcal{X}}\equiv\mathcal{C}\times\mathcal{Z}$. Observable choices are in the set $\widetilde{\mathcal{X}}_+\equiv\widetilde{\mathcal{X}}\cup\{0\}$. Let $\widetilde{\mu}^g_{\widetilde{y}|\widetilde{x}}(\boldsymbol{\theta},\boldsymbol{\gamma})$ denote the equilibrium probability that a gender-g individual of observable type $\widetilde{x}\in\widetilde{\mathcal{X}}$ makes observable partner-choice $\widetilde{y}\in\widetilde{\mathcal{X}}_+$, including singlehood.

6.1. The Likelihood Function

The data is a pair of gender-specific matrices of marriage frequencies. A generic entry $M^g_{\widetilde{x}\widetilde{y}}$ in the gender-g data matrix contains the number of gender-g persons of observable type $\widetilde{x}\in\widetilde{\mathcal{X}}$ who made observable partner choice $\widetilde{y}\in\widetilde{\mathcal{X}}_+$. The model is estimated by Maximum Likelihood. The log-likelihood function is

$$\log\mathcal{L}(\boldsymbol{\theta},\boldsymbol{\gamma})=\sum_{g\in\{m,f\}}\sum_{\widetilde{x}\in\widetilde{\mathcal{X}}}\sum_{\widetilde{y}\in\widetilde{\mathcal{X}}_+}M^g_{\widetilde{x}\widetilde{y}}\log\widetilde{\mu}^g_{\widetilde{y}|\widetilde{x}}(\boldsymbol{\theta},\boldsymbol{\gamma}).$$

A detailed derivation is provided in the Supplementary Appendix.

6.2. Estimation Procedure

Let y^{z^1} be an indicator for holding a basic qualification level $z=z^1$ and w the date of birth relative to the RoSLA threshold. Assumptions 1 and 2 imply

$$\frac{\gamma^g_1}{\gamma^g_0}=\frac{\lim_{w\downarrow 0}E^g[y^{z^1}|w]}{\lim_{w\uparrow 0}E^g[y^{z^1}|w]},\ \text{for } g=m,f. \tag{6}$$

The RDD estimator of (1) for qualification z^1 provides precise estimates of (6) as $\gamma^m_0/\gamma^m_1=0.755$ and $\gamma^f_0/\gamma^f_1=0.739$ (with bootstrapped standard errors 0.013 and 0.014, clustered on the running variable w) and we fix the ratios at these values. This leaves γ^m_1 and γ^f_1 as the two free parameters in $\boldsymbol{\gamma}$.[29]

While γ^m_1 and γ^f_1 are trivially constrained to lie within the unit interval, the qualification distribution imposes a tighter lower bound. Specifically, the post-RoSLA basic qualification rate for gender-g individuals is the product of γ^g_1 and the proportion of gender-g individuals with ability level a^1. However, this proportion is bounded above by the complement of the proportion of individuals with

[29]Note that $\gamma^g_0/\gamma^g_1=1-\varphi^g_{z^1}/\lim_{w\downarrow 0}E^g[y^{z^1}|w]$ where the estimated values of $\varphi^g_{z^1}$ from (1) are 0.106 (0.005) for men and 0.131 (0.007) for women (see Table 2), and estimates of $\lim_{w\downarrow 0}E^g[y^{z^1}|w]$ are obtained as the empirical rate of holding qualification z^1 among individuals in the first RoSLA-affected cohort, 0.433 (0.004) for men and 0.499 (0.004) for women.

qualification z^2 and ability a^2. As a result, consistency with observed qualification rates implies a tighter lower bound on γ_1^g of $\underline{\gamma}_1^g = 0.69$, $g = m, f$.[30]

We note here that, at the lower bound $\gamma_1^g = 0.69$ for $g = m, f$, all unqualified individuals have medium ability and no one has low ability; hence, there is no latent ability and our latent ability model collapses to the standard CS model.

Figure 8 displays the log-likelihood function concentrated over $\boldsymbol{\theta}$, evaluated initially under the additional restriction $\gamma_1^m = \gamma_1^f \equiv \gamma_1$. The resulting profile $\log \mathcal{L}^c(\gamma_1) \equiv \log \mathcal{L}(\hat{\boldsymbol{\theta}}(\gamma_1), \gamma_1)$ depends on the common value of γ_1, where $\hat{\boldsymbol{\theta}}(\gamma_1)$ is the maximum likelihood estimator of $\boldsymbol{\theta}$ for given γ_1, subject to $\gamma_0^m/\gamma_1^m = 0.755$ and $\gamma_0^f/\gamma_1^f = 0.739$. It measures the model's fit to the data as a function of γ_1.

The point labeled "Pref. Spec." marks the maximum of this function, identifying the maximum likelihood estimate of γ_1 as $\hat{\gamma}_1 = 1$, and thus of γ_0^m and γ_0^f as $\hat{\gamma}_0^m = 0.755$ and $\hat{\gamma}_0^f = 0.739$.[31] Notably, this estimate of γ_1 implies that all individuals with medium ability obtain a basic qualification after RoSLA, consistent with a substantial reduction in the opportunity cost of doing so. The corresponding maximum likelihood estimate of the parameters in the marital surplus function, $\boldsymbol{\theta}$, is $\hat{\boldsymbol{\theta}}(\gamma_1^m = 1, \gamma_1^f = 1)$.[32]

For comparison, the point labeled "CS" shows the likelihood value from a separately estimated CS model without latent ability (see Supplementary Appendix). This model corresponds to our latent ability specification with $\gamma_1 = 0.69$ and yields the lowest likelihood value, indicating a particularly poor fit. The poor fit stems from distinctly erroneous predictions regarding the impact of RoSLA on unqualified single rates (see Supplementary Appendix S.4.4, Figure S.6).

The estimates of γ_1^m and γ_1^f remain $\hat{\gamma}_1^m = \hat{\gamma}_1^f = 1$ even in a specification that allows them to differ from each other (Supplementary Appendix, Figure S.3). Hence, even when accounting for potential gender differences in the propensity to gain a basic qualification, the model specification that best fits the data continues to imply that all individuals with medium ability obtain a basic qualification after the RoSLA. Accordingly, our estimate of $\boldsymbol{\theta}$ remains $\hat{\boldsymbol{\theta}}(\gamma_1^m = 1, \gamma_1^f = 1)$.

[30] Specifically, $\underline{\gamma}_1^g \equiv \max_{c \in \mathcal{C}_1} \left\{ \Pr^g(z = z^1 | c \in \mathcal{C}_1)/[1 - \Pr^g(z = z^2 | c \in \mathcal{C}_1)] \right\}$.

[31] As robustness we have conducted the analysis for $\hat{\gamma}_0^g$ between 0.715 and 0.775. This affects the point estimate of the marital surplus matrix, but does not qualitatively affect our conclusions.

[32] The concentrated likelihood is a continuous function of γ_1, but with a kink at $\gamma_1 = 0.81$. Numerical examination reveals that, at $\gamma_1 = 0.81$, two different values of $\boldsymbol{\theta}$ fit the data equally well. Specifically, the kink is the intersection of the paths of two local maximizers, $\hat{\boldsymbol{\theta}}^1(\gamma_1)$ and $\hat{\boldsymbol{\theta}}^2(\gamma_1)$. For $\gamma_1 < 0.81$, the global maximizer is $\hat{\boldsymbol{\theta}}^1(\gamma_1)$, while for $\gamma_1 > 0.81$, it is $\hat{\boldsymbol{\theta}}^2(\gamma_1)$. Hence, the restriction $\gamma_1^m = \gamma_1^f = 1$ in Theorem 1 remains important for identification of $\boldsymbol{\theta}$ in the richer specification with eight cohorts, three ability-types, three qualification levels, age gap preferences and trends. Fortunately, Figure 8 shows that the best fit is obtained with $\hat{\gamma}_1^m = \hat{\gamma}_1^f = 1$ where the maximum likelihood estimate of $\boldsymbol{\theta}$ is unique.

The remainder of the paper focuses on the estimated marital surplus function and its implications for marriage rates and sorting.

Figure 8: The Log-Likelihood Function Concentrated over $\boldsymbol{\theta}$, subject to $\gamma_1^m = \gamma_1^f$

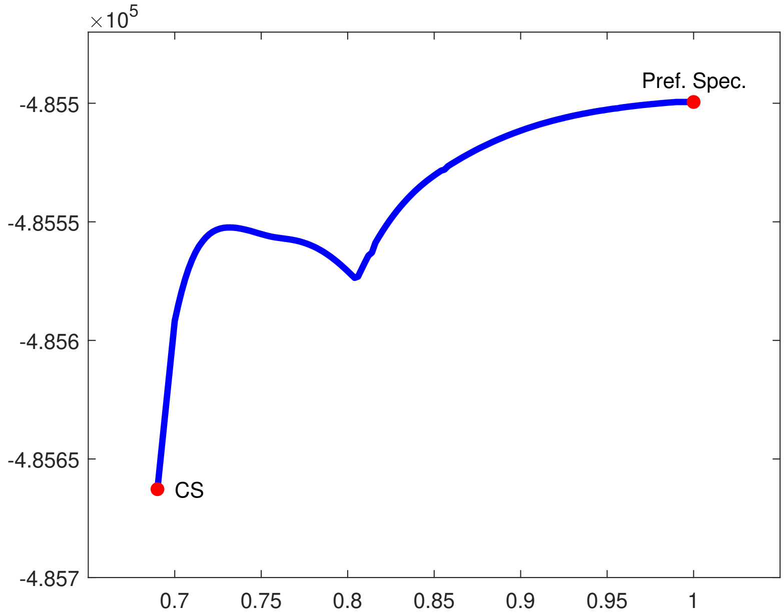

Notes: "Pref. Spec" indicates the maximum likelihood value. The "CS" indicates the likelihood value associated with the CS model without latent ability.

7. Results, Model Fit, and Specification Tests

7.1. Ability-Qualification Profile Contributions to Marital Surplus

Table 3 gives the estimates of the systematic surplus function $\zeta(a, z, a', z')$ defined over ability-qualification profiles.[33] As the surplus levels are not immediately interpretable we focus our discussion of Table 3 on comparing surplus values across type profiles with particular emphasis on partner trait complementarities.

The upper left 2×2 sub-matrix in Table 3 gives the estimated surpluses from marriages where both spouses are unqualified, but with low or medium ability. This sub-matrix also trivially exhibits increasing differences. Among the unqualified there is strong complementarity with respect to ability.

Similarly, the center 2×2 sub-matrix gives the estimated surpluses from marriages where both spouses have medium ability, but with or without a basic qualification. This sub-matrix too exhibits increasing differences, indicating strong complementarity with respect to holding a qualification.

[33]The estimated parameters of the age gap and trend functions are presented in the Supplementary Appendix.

If we then disregard the medium ability, but unqualified (second column and second row), the remaining 3×3 matrix gives the estimated marital surpluses associated type-profiles where each spouse has an ability and a qualification level that are perfectly aligned. It too exhibits increasing differences, thus indicating complementarity with respect to increases in aligned ability-and-qualification.

Table 3: The Systematic Marital Surplus Function $\zeta(a, z, a', z')$

	Wife type			
	(a^0, z^0)	(a^1, z^0)	(a^1, z^1)	(a^2, z^2)
Husband type (a^0, z^0)	-0.580^{***} (0.222)	-3.061^{***} (0.544)	-2.030^{***} (0.139)	-5.042^{***} (0.159)
Husband type (a^1, z^0)	-2.697^{***} (0.559)	-1.472^{**} (0.669)	-2.234^{***} (0.407)	-4.921^{***} (0.443)
Husband type (a^1, z^1)	-2.500^{***} (0.197)	-2.501^{***} (0.338)	-1.133^{***} (0.083)	-3.064^{***} (0.090)
Husband type (a^2, z^2)	-4.656^{***} (0.208)	-4.271^{***} (0.355)	-2.310^{***} (0.089)	-0.498^{***} (0.090)

Notes: Asymptotic standard errors in parentheses. *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively.

7.2. *Model Fit*

7.2.1. Never-Married Rates Figure 9 plots empirical and model predicted never-married rates by gender, cohort, unconditional on qualification level and by level of qualification.[34] Importantly, the estimated model correctly predicts increasing never-married rates for both unqualified men and women at the RoSLA reform threshold. Figure 8 indicated that allowing for latent ability is essential to fit the observed marriage market responses to the RoSLA. The reasons is that, with latent ability, the RoSLA triggers two composition effects.

First, among the pre-reform cohorts, individuals with medium ability but no qualification are acceptable partners for those with some ability and some qualification; hence, among the unqualified in the pre-reform cohorts, those with medium

[34] Since the model abstracts from divorce and remarriage, it cannot account for gender differences in lifetime marriage incidence arising from differential remarriage rates. While no official statistic permits an ex ante adjustment for this feature, remarriage rates following divorce are approximately twice as high for men as for women (ONS, 2016), causing the model to mechanically underpredict male never-married rates and overpredict female rates in levels. In Figure 9 a scaling factor—aligning the predicted aggregate never-married rate in the baseline 1953 cohort with the observed aggregate rate—has therefore been applied to all predicted rates for each gender. This adjustment affects only the gender-specific levels and does not influence the model's predicted cohort patterns or the changes around the RoSLA threshold.

Figure 9: Predicted and Empirical Never-Married Rates by Cohort, Gender, and Qualification

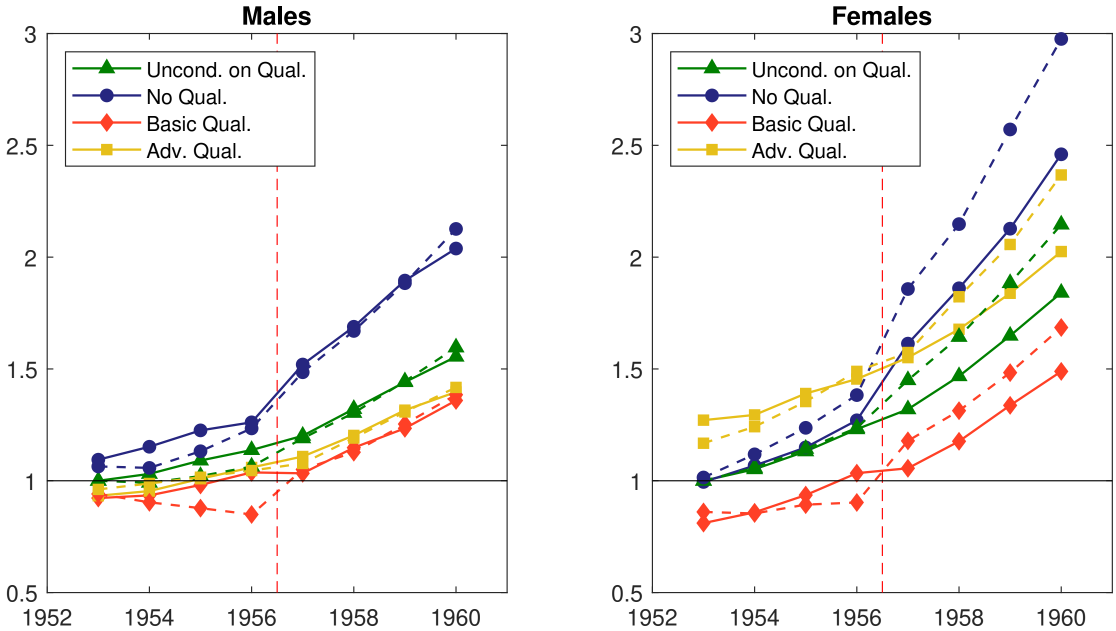

Notes: The horizontal axes show academic cohorts. The vertical axes show never-married rates in index form relative to the baseline 1953 cohort. Solid lines render the data (see Figure 3) whilst the dashed lines render model predictions. A scaling factor has been applied to all predicted rates for each gender, aligning the predicted aggregate never-married rate for the baseline 1953 cohort with the observed aggregate rate. This adjusts for remarriages being more common among men than among women leading the estimated model to mechanically underpredict (overpredict) the never-married rates among men (women) in levels. See footnote 34 for further details.

ability have lower never-married rates. As the qualification choices of the post-RoSLA cohorts eliminated the medium ability type from the unqualified group, the never-married rate among the unqualified increases.

Second, the RoSLA-induced change in qualification choices renders the post-reform type distribution polarized between individuals without ability or qualifications and individuals possessing both ability and a basic qualification. Under the estimated complementarity in qualifications in marital surplus, this impact the partner choices of individuals with medium ability. Pre-RoSLA, medium-ability individuals without qualifications would frequently marry unqualified partners, since they themselves also lacked qualifications. Post-RoSLA, however, medium-ability individuals hold a basic qualification, which under the estimated complementarity structure increases the marital surplus associated with a qualified relative to an unqualified spouse. As a result, the proportion of individuals willing to marry an unqualified individual declines in the post-RoSLA cohorts, contributing to the increase in the never-married rate among the unqualified at the reform threshold.

The estimated model also fits to the never-married rates for individuals with basic and advanced qualifications, possibly with the exception of men with a basic qualification before the reform. Here, the model predicts an increase in the never-married rate at the RoSLA reform threshold of individuals holding a basic qualification due to their increased supply. However, there is little evidence of such an effect in the data.

7.2.2. Qualification Homogamy The estimated model precisely replicates the distributions of spousal qualifications by gender and own qualification when pooling cohorts, i.e. Figure 4. These model predictions are presented in Figure S.4 in the Supplementary Appendix.

Figure 10 shows the predicted homogamy indices $S^g(z, c)$ by gender and cohort as dashed lines, with solid lines representing the data. The estimated model replicates the stability of the aggregate ("weighted") homogamy index across cohorts. It further replicates the markedly higher homogamy indices for advanced qualified individuals, and finally, it replicates well the observed increases in the homogamy indices for individuals with no qualifications at the RoSLA threshold.

7.2.3. Age Gaps The estimated model reproduces the observed age distribution very accurately (Supplementary Appendix, Figure S.5). Our descriptive analysis highlighted a small temporary shift in the age gap distribution among the cohorts born just after the RoSLA threshold (Figure 6). The estimated model captures only a fraction of this temporary shift. Hence, if anything, the model may be overstating the strength of preferences for sorting on age.

7.3. Specification Checks

Our Supplementary Appendix S.4 contains two model specification checks. First, we fit a "qualification-only" version of our model without latent ability and confirms that marital sorting on ability is essential for capturing the observed rise in never-married rates among unqualified men and women at the RoSLA threshold. Second, we develop and implement a test of the assumption of additive cohort (or age gap) preferences. The tests does not reject the additive age-gap-based

Figure 10: Predicted and Empirical Qualification Homogamy by Cohort, Gender, and Qualification

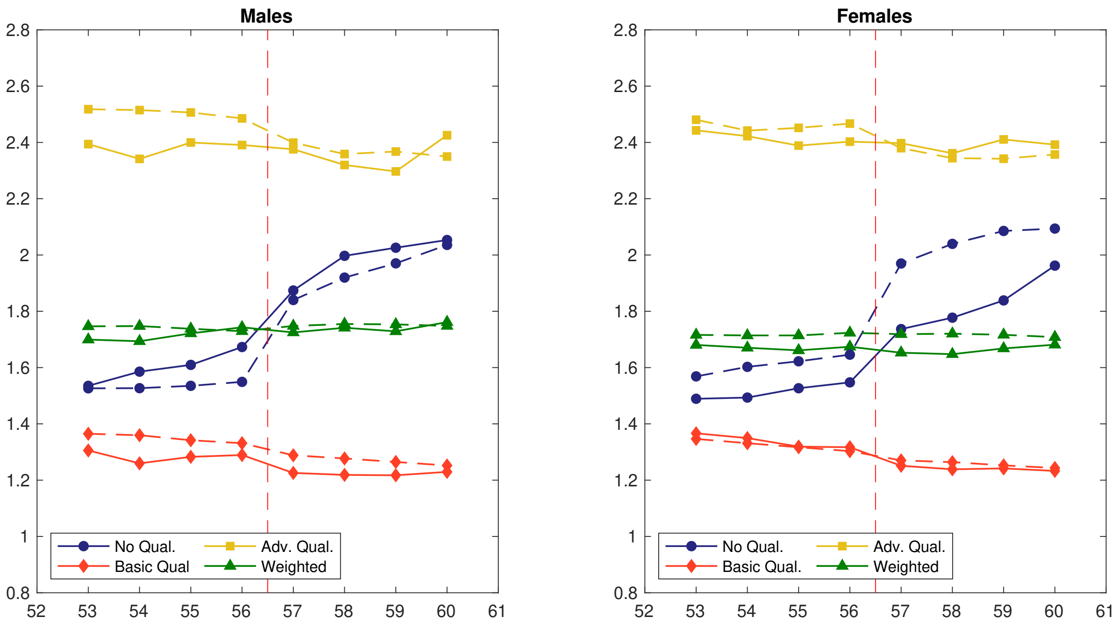

Notes: The horizontal axes show academic cohorts. The vertical axes show the homogamy index, see (2). Solid lines renders the data. Dashed lines renders model predictions.

preference structure with linear trends.[35]

8. Qualifications, Ability and Marriage Market Outcomes

The estimated structural model recovers equilibrium marriage probabilities for every existing combination of gender, ability, and qualification. This permits comparisons of otherwise identical individuals who differ only in a single characteristic, while holding the equilibrium environment fixed. Throughout this section, we refer to these comparisons as model-implied causal effects reflecting that these are structural counterfactuals implied by the estimated model rather than directly observed empirical contrasts. Our primary focus is on the causal effects of own qualification and own ability on the probability of ever marrying, $\Pr(\text{Ever married}|c, a, z) = 1 - \mu^g_{0|c,a,z}$. We also consider whether own qualification and ability predict spousal qualification. As the latter object is defined only conditional on marriage, it does not have a strict causal interpretation.

[35]Supplementary Appendix S.4 further shows that the estimated surplus structure is robust to restricting the LFS sample to couples with average age 40 or below in order to reduce any bias from selection into divorce by couple profile.

8.1. *The Causal Effects of Qualifications and Ability on Ever-Marrying*

The model-implied causal effects of ability and qualification on the ever-married probability for gender g are given by:

$$\Delta_a^g(c) \equiv \mu^g_{0|c,a^0,z^0} - \mu^g_{0|c,a^1,z^0}, \quad \text{and} \quad \Delta_z^g(c) \equiv \mu^g_{0|c,a^1,z^0} - \mu^g_{0|c,a^1,z^1}.$$

Hence, $\Delta_a^g(c)$ is the causal effect of possessing medium ability on the ever-married rate of individuals with no qualification. $\Delta_z^g(c)$ is the causal effect of basic qualification attainment on the ever-married rate of medium ability individuals. Let $\Delta^g(c)$ be the effect on the probability of ever-marrying of having both medium ability *and* holding a basic qualification compared to having low ability *and* being unqualified; clearly,

$$\Delta^g(c) \equiv \Delta_a^g(c) + \Delta_z^g(c).$$

We can identify $\Delta_a^g(c)$ and $\Delta_z^g(c)$ only for the pre-RoSLA cohorts $c \in \mathcal{C}_0$. In contrast, $\Delta^g(c)$ is identified for all sample cohorts.

Figure 11: The Causal Effects of Ability and Qualification on Marital Outcomes

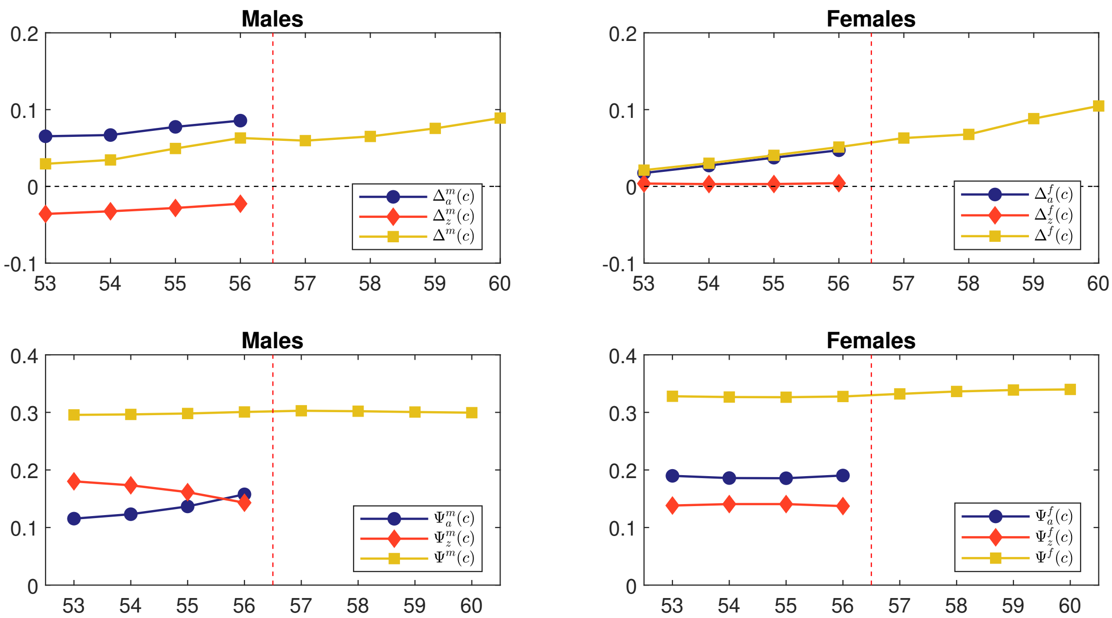

Notes: The horizontal axes show academic cohorts. The vertical axes show the effect of own ability and basic qualification attainment on the rate of (i) ever marrying (top row) and (ii) the spouse holding at least a basic academic qualification (bottom row).

The top panels of Figure 11 shows $\Delta_a^g(c)$, $\Delta_z^g(c)$, and $\Delta^g(c)$ for men and women, respectively. The gap in the ever-married probability between medium ability, basic qualified individuals and low ability, unqualified individuals (yellow line) grew substantially over the sample cohorts for both men and women, from only a couple of percentage points for those born in the early 1950s to about ten percentage points for those born in 1960. The top-right panel shows that, for medium ability women, the model-implied effect of holding a basic qualification on the probability of ever-marrying is effectively zero. The top-left panel shows that the corresponding basic qualification effect for medium-ability men is modestly negative. This should be interpreted as an equilibrium matching effect: under the estimated complementarity in qualifications, acquiring a basic qualification shifts these men toward a thinner and more competitive qualified segment of the marriage market, while the effect attenuates near the reform threshold as qualified women become more abundant. In contrast, the estimated effects of medium ability are positive for both unqualified men and unqualified women. As a result, the total gap in marriage probability $\Delta^g(c)$ between medium ability, basic qualified individuals and low ability, unqualified individuals is driven by the strongly positive causal effect of ability.

Equilibrium effects make the causal impact of qualifications on marriage outcomes considerably harder to identify than effects on individual outcomes such as earnings or health. The small model-implied effects estimated here are consistent with the instrumental variable evidence in Lefgren and McIntyre (2006) and Anderberg and Zhu (2014), but are estimated much more precisely because the structural model explicitly accounts for equilibrium matching.

Our analysis of the ever-married probability connects to the literature on "marital premia". As noted by Chiappori, Salanié, and Weiss (2017), in the CS-setting, the negative of the log never-married rate for a given type measures their average marital utility, whereby log differences in never-married rates across types serve as a measure of a "marriage premium". Our model thus shows positive ability marriage premia for both men and women. In contrast, the equilibrium marriage market return to a basic qualification marriage is zero for women even marginally negative for men.[36]

[36]To see this, note that appropriately scaled $\Delta_a^g(c)$- and $\Delta_z^g(c)$-effects can be interpreted as the marriage premium associated with (medium) ability and (basic) qualification, respectively. Specifically $-\log\left(\frac{\mu^g_{0|c,a^1,z^0}}{\mu^g_{0|c,a^0,z^0}}\right) \approx \frac{\Delta_a^g(c)}{\mu^g_{0|c,a^0,z^0}}$, and $-\log\left(\frac{\mu^g_{0|c,a^1,z^1}}{\mu^g_{0|c,a^1,z^0}}\right) \approx \frac{\Delta_z^g(c)}{\mu^g_{0|c,a^1,z^0}}$.

8.2. Ability, Qualifications and the Spousal Qualification Gap

Define $\Psi_a^g(c)$ as the difference in the probability of the spouse holding at least a basic qualification between those with medium ability a^1 and those with low ability a^0, conditional on being unqualified z^0 and married, for gender g and cohort c. Similarly, define $\Psi_z^g(c)$ as the difference in spouse qualification rate between those with qualification z^1 and those unqualified z^0, conditional on medium ability a^1 and being married, for gender g and cohort c. Finally, define $\Psi^g(c) = \Psi_a^g(c) + \Psi_z^g(c)$ as the total gap in spouse qualification rate between those of type (a^1, z^1) and those of type (a^0, z^0). $\Psi_a^g(c)$ and $\Psi_z^g(c)$ are identified only for pre-RoSLA cohorts, whereas $\Psi^g(c)$ is identified for all sample cohorts. Because the spousal qualification gap is defined conditional on marriage, it does not have strict causal interpretation. Nonetheless, the decomposition of $\Psi^g(c)$ into the quantities $\Psi_a^g(c)$ and $\Psi_z^g(c)$ is novel and informative about the relative importance of ability and qualifications in the marriage market.

The lower panels of Figure 11 shows that the probability of being married to an academically qualified spouse increases with own ability and with own (basic) qualification. For both men and women, ability and an own qualification have similar positive effects on the qualification rate of the spouse. The total effect of both ability and a qualification is about 30 percent for both men and women—which corresponds to the differences seen in Figure 4—and highly stable across cohorts. That is, conditional on marrying, own qualification attainment and own ability are both important determinants of spousal qualification type.

9. General Equilibrium Marriage Market Effects of the RoSLA

We end our analysis by demonstrating that the RoSLA affected the marital outcomes of a wide set of people, including individuals whose educational choices and outcomes were not directly affected by the reform. To do so we contrast the simulated marriage outcomes from our estimated model to those obtained from a counterfactual scenario where the RoSLA reform was never implemented.

In the counterfactual scenario we assume that the pre-reform qualification attainment process in Figure 7 continued to apply across all the sample cohorts $c \in \mathcal{C}$, providing us with counterfactual qualification distributions based on which we simulate a counterfactual marriage market equilibrium. We focus our comparison across equilibria on the outcomes by ability type as ability was, per assumption, unaffected by the RoSLA.

9.1. Never-Married Rates and the RoSLA

Figure 12 plots the difference in never-married rates with and without the RoSLA. A positive number for a particular ability-type, gender and cohort in Figure 12 implies that the reform increased the never-married rate for that particular type.

Figure 12: The Effect of the RoSLA on the Never-Married Rates by Ability, Cohort and Gender

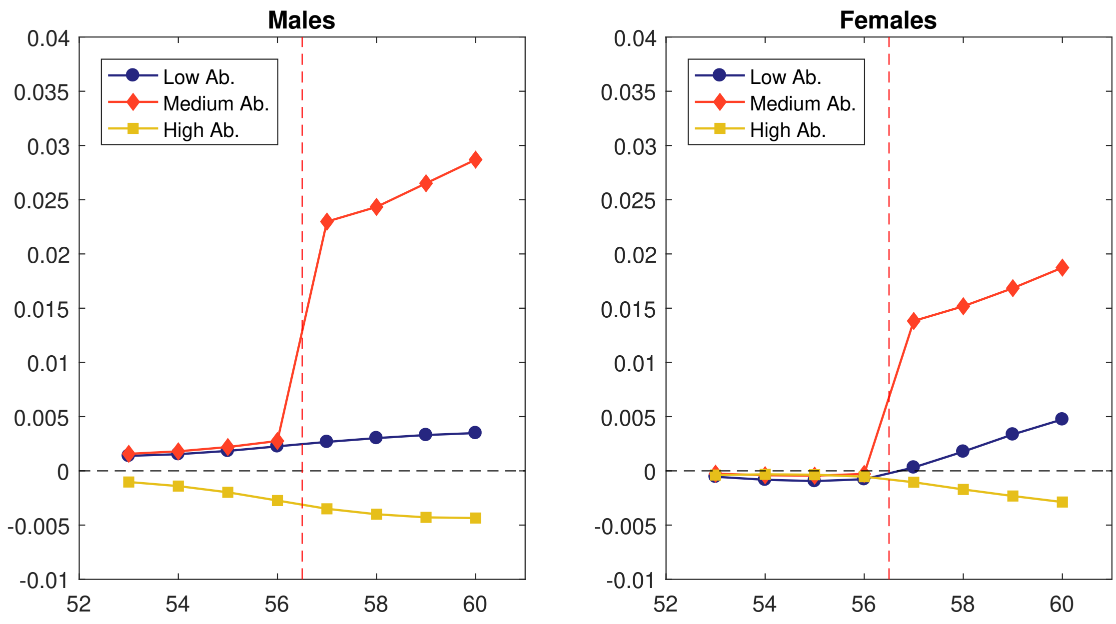

Notes: The horizontal axes show academic cohorts. The vertical axes show the difference in never-married rates with and without the RoSLA.

We begin by considering the low- and high-ability groups. Since the reform did not alter their own qualification attainment, any changes in their marriage outcomes must arise through equilibrium effects operating through the marriage market. These simulated effects therefore provide a clean illustration of the indirect consequences of the reform. We then turn to medium-ability individuals, for whom the reform both altered qualification attainment directly and changed the surrounding equilibrium. Consequently, the simulated effects for this group combine direct and equilibrium channels.

Low-ability individuals are unqualified in both regimes and, under the estimated complementarity in qualifications, are most likely to match with other unqualified individuals in the pre-RoSLA equilibrium. By reducing the supply of such partners, the RoSLA reduced the availability of suitable matches in equilibrium and increased their never-married rates, despite these individuals not being directly affected by the reform themselves. Because husbands are typically older than their wives, this general equilibrium effect emerged before the reform threshold for men, while for women it was concentrated among the post-reform cohorts.

High-ability individuals are academically qualified in both regimes and, under the estimated complementarity in qualifications, are most likely to match with other qualified individuals. By increasing the supply of qualified potential spouses, the RoSLA expanded the availability of suitable matches in equilibrium and modestly increased their probability of ever marrying, despite these individuals not being directly affected by the reform themselves. As for the low-ability type, the typical husband-wife age gap implies that these general equilibrium effects appeared earlier for men than for women. Within a few cohorts of the reform threshold, the estimated effects stabilized at around half a percentage point, increasing ever-marriage among the high-ability type by approximately the same magnitude as it reduced ever-marriage among the low-ability type.

For the medium-ability type, the simulated effects combine the direct consequences of acquiring a basic qualification with the indirect effects arising through equilibrium adjustment in the marriage market. Nevertheless, the increase in never-marriage following the reform is consistent with the estimated complementarity in qualifications: as medium-ability individuals acquire basic qualifications, the relative surplus associated with qualified matches increases, while the simultaneous expansion in the supply of qualified medium-ability individuals intensifies competition within that segment of the marriage market. Together, these forces contribute to the higher never-married rates estimated for medium-ability men and women in the post-RoSLA equilibrium.

9.2. Marital Sorting and the RoSLA

Figure 13 illustrates the effect of the reform on marital sorting on ability. It displays the impact of the reform on the distribution of spouse ability type by own ability type, cohort and gender, conditional on marriage. The top left panel shows that, following the reform, low-ability men became less likely to marry medium-ability women and correspondingly more likely to marry low-ability women. A similar pattern is observed for low-ability women in the bottom left panel. These findings are the natural counterpart to the effects on ever-marriage. By reducing the supply of suitable medium-ability unqualified partners, the reform both lowered the marriage prospects of low-ability individuals and made them more likely to marry within their own ability group rather than to "marry up" in ability terms.

The impact of the reform on high-ability individuals was the mirror image. Since this ability type acquired advanced qualifications irrespective of the reform, the simulated effects again reflect general equilibrium adjustment alone. As the reform increased the supply of academically qualified medium-ability individuals, high-ability men and women became more likely to marry medium-ability partners and correspondingly less likely to marry within their own ability group. Together with the increase in their probability of ever marrying, these findings indicate that the reform improved the marriage prospects of high-ability individuals by expanding the pool of suitable qualified partners.

Figure 13: The Effect of the RoSLA on the Distribution of Spouse Ability

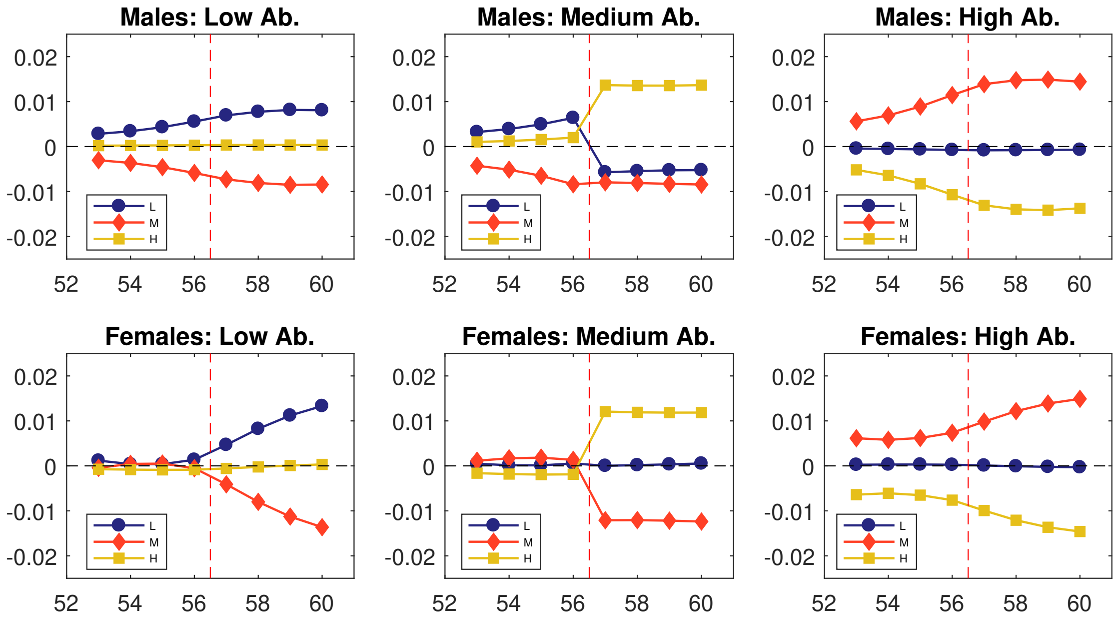

Notes: The horizontal axes show academic cohorts. The vertical axes show the difference in spousal ability rates with and without the RoSLA.

For the medium-ability type, the simulated effects again combine the direct impact of acquiring a basic qualification with the accompanying equilibrium adjustment of the marriage market. The most pronounced effect is that, following the reform, medium-ability men and women became more likely to marry high-ability partners. This is consistent with the estimated complementarity in qualifications: as medium-ability individuals increasingly acquire a basic qualification, marriages to academically qualified high-ability individuals become relatively more attractive. Since the probabilities of marrying each ability type sum to one, the increase in marriages to high-ability partners is accompanied by a corresponding reduction in marriages to low- and medium-ability partners.

As a broad characterization, based on Figures 12 and 13, the RoSLA (i) increased the never-married rates of low and medium ability individuals, while decreasing it for high ability individuals, and, (ii) increased assortative matching among low ability individuals while increased the marital mixing between high and medium ability types. Effectively, the equilibrium adjustments in response to the RoSLA left low ability individuals increasingly isolated in the marriage market.

These equilibrium responses also account for the descriptive homogamy patterns documented in Section 4. There we noted a permanent rise in the homogamy index among the unqualified after the RoSLA (Figure 5), alongside a modest decline among the basic qualified, but could not say, from the descriptive index alone, whether these movements were purely mechanical consequences of the shifting qualification marginals or reflecting more substantive forces. Our analysis of the effect on marital outcomes by ability indicate that they were substantive. The rise in unqualified homogamy reflects the increased tendency of low-ability individuals to marry one another (leftmost panels of Figure 13). The decline in basic-qualified homogamy reflects the greater mixing between medium- and high-ability individuals (middle and rightmost panels of Figure 13).

The same two movements are also registered, as margin-invariant descriptive patterns, by the rise in the extensive-margin and the fall in the intensive-margin odds ratio discussed in Section 4. Under the standard Choo–Siow reading, such odds-ratio changes would signal shifts in the marital surplus; here, with cohorts spaced too closely for the surplus to have moved, that interpretation is neither available nor needed. Both patterns arise not from any change in the marital surplus across these closely spaced cohorts, but from the reform-induced change in the ability composition underlying observed qualifications.

10. Discussion

We have documented that the 1972 Raising of the School Leaving Age (RoSLA) reform significantly reshaped marriage markets in Britain, and exploited a change in selection into basic qualification attainment to estimate implied causal effects of ability and qualification on equilibrium marriage outcomes. The RoSLA increased never-married rates among the least educated and strengthened assortative matching among less academically able and unqualified individuals. Both qualifications and ability are valued and are complements in the marriage market. Ability increases the probability of ever marrying while holding a basic qualification does not, implying that the observed gap in marriage rates between basic qualified and unqualified individuals is entirely due to selection on ability.

Our findings suggest that educational expansions may have substantial implications for inequality, operating through household formation and marital sorting, even in settings where the direct labor market returns to qualifications are only modest. In the case of the RoSLA, the reform increased marital isolation among the least educated, leaving individuals without academic qualifications more likely to remain without a partner and, among those who did partner, increasingly concentrated in low-qualified households. Additional descriptive evidence using the UK Labour Force Survey show that, by mid-life, post-RoSLA unqualified individuals were indeed more likely to live in single-adult and workless households and less likely to live in dual-earner households (see Figure S.2, Supplementary Appendix S.1.9). These patterns illustrate how an educational expansion may amplify inequality through changes in household structure and the concentration of economic disadvantage across households—complementing recent work that emphasizes the role of assortative matching and marriage market returns in shaping household inequality (Eika, Mogstad, and Zafar, 2019; Chiappori, Costa Dias, Meghir, and Zhang, 2025).

Our conclusions are, of course, drawn from a particular historical episode. The estimated effects are local to cohorts at the lower end of the education distribution whose schooling decisions were directly affected by the 1972 reform. In contemporary settings, where secondary school completion is nearly universal, the relevant educational margins for marriage market outcomes have likely shifted upward towards post-secondary and tertiary education. However, the broader mechanism identified by our analysis remains relevant: educational expansion changes not only the quantity of schooling but also the composition of educational groups, with implications for household formation and inequality. Our results show that educational policies can shape long-run inequality not only through traditional channels of wages and employment, but also through their effects on partnership formation and the distribution of economically vulnerable households.

Appendix

A. Proof of Theorem 1 in the main text

The proof is constructive. Rather than proving identification indirectly, we show how each unknown object can be recovered sequentially from observables. We first isolate the cohort effects, then recover the latent ability composition, and finally reconstruct the complete matching matrix.

We prove Theorem 1 for the simplified $2 \times 2 \times 2$ environment introduced in Section 5.3. There is one pre-RoSLA cohort, c_0, and one post-RoSLA cohort, c_1; two ability levels, $\{a_0, a_1\}$; and two qualification levels, $\{z_0, z_1\}$. Furthermore, Assumption 2 is imposed with $\gamma_1^m = \gamma_1^f = 1$, so that every medium-ability individual obtains a basic qualification after the reform.

The only difficulty is that pre-reform unqualified individuals may be either low or medium ability, whereas post-reform unqualified individuals are known to be low ability. The proof shows how information from the latter identifies the former. It proceeds in four steps.

1. We first identify the cohort component of marital surplus using only matches between *distinguishable* types.

2. We then exploit the fact that post-reform unqualified individuals are necessarily of low ability. Their observed matching behaviour provides an anchor for recovering the matching behaviour of low-ability individuals in the pre-reform cohort, where ability is unobserved by the econometrician.

3. Next, we invoke the assumption that the ability distribution is unchanged across the reform. This yields a single nonlinear equation whose unique solution identifies the number of low-ability individuals that are unmarried in the pre-reform cohort. This completes identification of the latent pre-reform ability composition.[37]

4. Finally, once every type-specific matching frequency has been recovered, the surplus matrix, together with the systematic payoffs (U, V), follows exactly as in Choo and Siow (2006).

The remainder of the appendix formalizes these four steps.

[37] Note that we do not identify the latent type of any particular individual, but only the overall composition of the population.

Given $\gamma_1^m = \gamma_1^f = 1$, there are five possible types for each gender:

$$X \equiv \left\{1 := (c_0, a_0, z_0),\ 2 := (c_0, a_1, z_0),\ 3 := (c_0, a_1, z_1),\ \tilde{1} := (c_1, a_0, z_0),\ \tilde{3} := (c_1, a_1, z_1)\right\}. \tag{A.1}$$

Types 1 and $\tilde{1}$ correspond to low-ability individuals without qualifications in the pre- and post-reform cohorts respectively. Types 3 and $\tilde{3}$ correspond to medium-ability individuals holding a basic qualification in the two cohorts. Type 2 consists of medium-ability individuals who remain unqualified in the pre-reform cohort; by construction, this type has no post-reform counterpart.

Gender is observable, and we refer to male types by $x \in X$ and female types by $y \in X$. The econometrician observes cohort and qualification, but not ability. Consequently, the two pre-reform unqualified types, $X^I = \{1, 2\}$, are observationally indistinguishable. The remaining three types,$X^D = \{3, \tilde{1}, \tilde{3}\}$, are distinguishable.

This distinction between distinguishable and indistinguishable types drives the identification argument. Matches involving only distinguishable types are directly observed, whereas matches involving the pre-reform unqualified group must be inferred. The proof shows that information from the distinguishable types, together with Assumptions 1–3, is sufficient to recover the matching frequencies involving the indistinguishable types, and hence the complete matching matrix.

Table A.1 summarizes which matching frequencies are directly observed and which must be inferred. Rows indicate male types and the columns indicate female types. Each row and column is augmented with a 0-type to indicate the unmatched singles. The cell entries denote the number of matches between the corresponding types. Question marks indicate those matches whose numbers cannot be observed.

Besides the entries shown explicitly in Table A.1, we also observe the total number of matches between any distinguishable type and the set of indistinguishable types, $\mathcal{X}^I = \{1, 2\}$, which we index by I. Thus, we observe N_{Iy} and N_{xI} for any $x, y \in \mathcal{X}^D$. We also observe the number of matches where both types are indistinguishable, N_{II}, as well as the number of singles who are indistinguishable, N_{I0}^m and N_{0I}^f.

Step 1: Identifying the cohort component of surplus

The first step is to recover the cohort component of surplus from matches involving only distinguishable types. This is possible because distinguishable types are observed in both cohorts. The only assumption needed is the additive separability

Table A.1: Observed and unobserved marriage and singlehood frequencies

Men \ **Women**	1 (a^0, z^0, c^0)	2 (a^1, z^0, c^0)	3 (a^1, z^1, c^0)	$\tilde{1}$ (a^0, z^0, c^1)	$\tilde{3}$ (a^1, z^1, c^1)	0 $\cdot$
1	?	?	?	?	?	?
2	?	?	?	?	?	?
3	?	?	N_{33}	$N_{3\tilde{1}}$	$N_{3\tilde{3}}$	N_{30}
$\tilde{1}$	?	?	$N_{\tilde{1}3}$	$N_{\tilde{1}\tilde{1}}$	$N_{\tilde{1},\tilde{3}}$	$N_{\tilde{1}0}$
$\tilde{3}$	?	?	$N_{\tilde{3}3}$	$N_{\tilde{3}\tilde{1}}$	$N_{\tilde{3}\tilde{3}}$	$N_{\tilde{3}0}$
0	?	?	N_{03}	$N_{0\tilde{1}}$	$N_{0\tilde{3}}$	

of cohort effects from ability-qualification profiles, which we restate here for the 2×2×2 specification.

Assumption A.3. Let $x = (c, a, z)$ denote the type of the man and let $y = (c', a', z')$ denote the type of a woman. The systematic surplus function Σ_{xy} is additively separable:

$$\Sigma_{xy} = \zeta\left(a, z, a', z'\right) + \lambda_{cc'}.$$

Our first step is to use the matches between observed types in the two cohorts and their singles rate to infer the cohort dimension of preferences. Since there are two cohorts, we have four parameters capturing the cohort-profile aspect of preferences: λ_{00}, λ_{01}, λ_{10}, λ_{11}. Since the equilibrium matching pattern depends only on the surplus matrix and not on its decomposition into ζ and λ, we may normalize λ_{00} to zero without loss of generality.[38]

Lemma A.1. *The cohort preference vector* $\Lambda := (\lambda_{01}, \lambda_{10}, \lambda_{11})$ *is overidentified under Assumption A.3.*

The key observation is that distinguishable types are unaffected by the latent composition problem, allowing the cohort component of surplus to be identified first.

[38] Since the equilibrium matching pattern depends only on the surplus matrix, and not on how it is decomposed, alternative normalizations have no implications for behavior.

Proof. We begin by noting that if $x, y \in \mathcal{X}^D$, then Σ_{xy} is identified. This follows from Choo and Siow (2006), who show that

$$\exp(\Sigma_{xy}) = \frac{(N_{xy})^2}{N^m_{x0} N^f_{0y}}. \tag{A.1}$$

Let i and j denote generic elements of the set of types $\{1, 2, 3\} \subset \mathcal{X}$, all belonging to cohort 0 (index i will refer to males, index j to females). We will use $\tilde{i}$ (respectively $\tilde{j}$) to denote the type from cohort 1 who has the same ability and qualifications as i (resp. j). Under Assumption A.3,

$$\Sigma_{\tilde{i}j} - \Sigma_{ij} = \lambda_{10} - \lambda_{00} = \lambda_{10},$$

independent of the values of i and j. Similarly, $\Sigma_{i\tilde{j}} - \Sigma_{ij} = \lambda_{01}$ and $\Sigma_{\tilde{i}\tilde{j}} - \Sigma_{ij} = \lambda_{11}$, independent of the ability-qualification profiles of i or j.

We now exploit the fact that the types 3 and $\tilde{3}$ are distinguishable in both cohorts. i.e. men and women who are qualified and who are therefore of medium-ability. We observe the number of matches, N_{xy}, for any x, y belonging to $\{3, \tilde{3}\}$. Furthermore, we observe the number of single men N^m_{x0} and single women N^f_{0x}, for each type x in this set. Consequently, using (A.1) we can infer Σ_{33}, $\Sigma_{\tilde{3}3}$, $\Sigma_{\tilde{3}\tilde{3}}$, and $\Sigma_{3\tilde{3}}$. Thus, taking differences $\lambda_{10} = \Sigma_{\tilde{3}3} - \Sigma_{33}$, $\lambda_{01} = \Sigma_{3\tilde{3}} - \Sigma_{33}$, $\lambda_{11} = \Sigma_{\tilde{3}\tilde{3}} - \Sigma_{33}$, so Λ is identified.

Furthermore, since $\tilde{1}$ is also in $\mathcal{X}^D$ we can infer $\Sigma_{\tilde{1}3}$ and $\Sigma_{\tilde{1}\tilde{3}}$, and their difference is an independent measure of $\lambda_{10} - \lambda_{11}$. Similarly, as we can infer $\Sigma_{3\tilde{1}}$ and $\Sigma_{\tilde{3}\tilde{1}}$, we have an independent measure of $\lambda_{01} - \lambda_{11}$. This yields two over-identification restrictions. □

The cohort component of marital surplus, Λ, is now identified. The remaining task is to recover the latent matching frequencies involving the indistinguishable types.

Step 2. Recovering the latent matching frequencies

We now proceed to the key step in the argument. The key identification insight is that post-reform unqualified individuals are known to be low ability, whereas pre-reform unqualified individuals are a mixture of low- and medium-ability types. Having identified the cohort effects in Lemma 1, we can translate the observed matching behavior of post-reform low-ability individuals into the corresponding behavior of low-ability individuals in the pre-reform cohort. Thus, although ability is latent before the reform, the post-reform cohort provides an observable anchor for recovering the matching behavior of the latent type. The subsequent equations simply implement this mapping.

The mathematical heart of the proof is to recover the latent partition of the pre-reform unqualified population into low- and medium-ability individuals. We accomplish this by first identifying the number of low-ability men who are single. Once this quantity is known, the remaining latent matching frequencies follow by construction. Recall that the latent types are the unqualified in cohort 0, with type 1 of low ability and type 2 of medium ability. Since we observe the total number of unqualified single males in this cohort, namely $N^m_{I0} = N^m_{10} + N^m_{20}$, it suffices to identify N^m_{10}, the number of single unqualified low-ability males in cohort 0; the count N^m_{20} then follows by subtraction.

We now proceed to express every latent matching frequency in terms of a single unknown—the number of low-ability singles in the pre-reform cohort. Once this is done, we invoke an "adding-up" equation, since the total number of low-ability individuals is determined by the total cohort size and our assumption that the distribution of ability is the same across cohorts. This "adding-up" equation, expressed terms of the number of low-ability singles, yields a single non-linear equation that has a unique solution.

To proceed, we restate (A.1) for the pairs (i, j) and $(\tilde{i}, j)$:

$$\exp(\Sigma_{ij}) = \frac{(N_{ij})^2}{N^m_{i0} N^f_{0j}}, \tag{A.2}$$

$$\exp(\Sigma_{\tilde{i}j}) = \frac{\left(N_{\tilde{i}j}\right)^2}{N^m_{\tilde{i}0} N^f_{0j}}. \tag{A.3}$$

If we divide (A.2) by (A.3), the left-hand side equals $\exp(-\lambda_{10})$, and on rearranging, we get

$$N_{ij} = N_{\tilde{i}j}\sqrt{N^m_{i0}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{i}0}}}. \tag{A.4}$$

We may now replace i with 1 and $\tilde{i}$ with $\tilde{1}$ in the above to get:

$$N_{1j} = N_{\tilde{1}j}\sqrt{N^m_{10}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.5}$$

This expresses the unobserved N_{1j} in terms of the observable $N_{\tilde{1}j}$ and the single unknown N^m_{10}. Critically, the expression is linear in $N_{\tilde{1}j}$, with a coefficient that is independent of j, a property that is important for the aggregation that is to follow. We rewrite this equation for each possible value of j: first for $j = 1, 2$,

$$N_{11} = N_{\tilde{1}1}\sqrt{N^m_{10}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.6}$$

$$N_{12} = N_{\tilde{1}2}\sqrt{N^m_{10}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.7}$$

Consequently,

$$N_{11} + N_{12} = (N_{\tilde{1}1} + N_{\tilde{1}2})\sqrt{N^m_{10}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.8}$$

Although we observe neither $N_{\tilde{1}1}$ nor $N_{\tilde{1}2}$, we do observe their sum, which we denote by $N_{\tilde{1}I}$, since this is the total number of matches between unqualified males of cohort 1 and unqualified females of cohort 0. For $j = 3$, which is distinguishable, $N_{\tilde{1}3}$ is observed individually, so we keep this term separate.

$$N_{13} = N_{\tilde{1}3}\sqrt{N^m_{10}}\frac{\exp(-\lambda_{10}/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.9}$$

Similarly, one can write down the expressions for $\exp(\Sigma_{1\tilde{j}})$ and $\exp(\Sigma_{\tilde{1}\tilde{j}})$. Since the ratio of the first to the second equals $\exp(\lambda_{01} - \lambda_{11})$, we get

$$N_{1\tilde{j}} = N_{\tilde{1}\tilde{j}}\sqrt{N^m_{10}}\frac{\exp((\lambda_{01} - \lambda_{11})/2)}{\sqrt{N^m_{\tilde{1}0}}}. \tag{A.10}$$

Setting $\tilde{j} = \tilde{1}$ and $\tilde{3}$ in turn, we get the expressions for $N_{1\tilde{1}}$ and $N_{1\tilde{3}}$.

We have therefore completed step 2: every latent matching frequency has now been expressed as a function of a single unknown: the number of low-ability singles in the pre-reform cohort.

Step 3: Solving for the key unknown

We now solve for the number of low-ability singles in the pre-reform cohort.

Let T^m_1 denote the total population of type 1 males, i.e. the number of males in cohort 0 who are of low ability. Note that we observe the total number of low ability males in cohort 1, since they correspond to the number who do not get a qualification. At this point, we invoke the assumption that the low-ability share is constant across the two cohorts. This implies that we observe T^m_1, since this equals the size of cohort 0 multiplied by the observed proportion of low-ability males in cohort 1. Now,

$$T_1^m = \Big(N_{11} + N_{12} + N_{1\tilde{1}} + N_{13} + N_{1\tilde{3}}\Big) + N_{10}^m. \qquad \text{(A.11)}$$

While none of the terms on the right-hand side above are directly observed, our derivations imply that they can be written in terms of observables and a single unobservable, as follows:

$$T_1^m = \big(N_{\tilde{1}I} + N_{\tilde{1}3}\big) \exp\left(-\frac{\lambda_{10}}{2}\right) \frac{\sqrt{N_{10}^m}}{\sqrt{N_{\tilde{1}0}^m}} + \big(N_{\tilde{1}\tilde{1}} + N_{\tilde{1}\tilde{3}}\big) \exp\left(\frac{\lambda_{01} - \lambda_{11}}{2}\right) \frac{\sqrt{N_{10}^m}}{\sqrt{N_{\tilde{1}0}^m}} + N_{10}^m. \qquad \text{(A.12)}$$

This is a quadratic in the single unknown $\sqrt{N_{10}^m}$: the coefficients λ_{01}, λ_{10}, λ_{11} have been identified in Lemma A.1, and the counts $N_{\tilde{1}I}$, $N_{\tilde{1}3}$, $N_{\tilde{1}\tilde{1}}$, $N_{\tilde{1}\tilde{3}}$ and the singles $N_{\tilde{1}0}^m$ are all observed.

The right-hand side is increasing in $\sqrt{N_{10}^m}$ when $\sqrt{N_{10}^m} > 0$. It is strictly less than the left-hand side when $\sqrt{N_{10}^m} = 0$, and strictly greater when $\sqrt{N_{10}^m} = \sqrt{T_1^m}$. Consequently, there is a unique solution in the interval $(0, \sqrt{T_1^m})$, that is given by the familiar quadratic formula. Furthermore, since we observe the sum $N_{10}^m + N_{20}^m$, i.e. the total number of unqualified men in cohort 0, we also infer N_{20}^m. We are therefore able to observe directly, or recover, the number of singles for every type of man. Since the argument for women is symmetric, we also observe or infer the number of singles for every type of woman.

To summarize: we have now completed step 3 and identified a key latent variable, the number of low-ability singles in the pre-reform cohort.

Step 4: Identifying the matching matrix

We now show that we can identify all the elements of the matching matrix. We directly observe N_{ij} and infer Σ_{ij} when both i and j belong to the distinguishable set $\mathcal{X}^D = \{3, \tilde{1}, \tilde{3}\}$.

Consider then $i = 1$ and $j \in \mathcal{X}^D$. Since we infer $\Sigma_{\tilde{1}j}$ (as $\tilde{1}, j \in \mathcal{X}^D$), and we infer the elements of Λ we can infer Σ_{1j}. For example, $\Sigma_{1\tilde{1}} = \Sigma_{\tilde{1}\tilde{1}} + \lambda_{01} - \lambda_{11}$. We can therefore use equation (A.2) to infer N_{1j}, for $j \in \mathcal{X}^D$. But since we observe the sum $N_{1j} + N_{2j}$, we can also infer N_{2j}, for $j \in \mathcal{X}^D$.

The final step is to show that we can identify N_{ij} when both $i \in \mathcal{X}^I$ and $j \in \mathcal{X}^I$. Since we infer $\Sigma_{1\tilde{1}}$ and also λ_{01}, we can infer Σ_{11}. Equation (A.2) then allows us to infer N_{11}, since we also know the number of singles. Since we have already inferred the total number of type 1 men, T_1^m, this allows us to infer N_{12} as a residual. Hence, we have inferred the number of matches N_{1j} for all j. Similarly, we can infer N_{21}, and finally, N_{22}.

Having identified all elements of the matching matrix and the number of singles of each type, we can identify all entries Σ_{ij} in the surplus matrix from (A.2). Furthermore, for each possible type profile we also identify U_{xy} and V_{xy}, the systematic components of match payoffs for men and women respectively, since these are the dual variables of the surplus-maximization program (Shapley and Shubik, 1971) and can be recovered from the observed matching frequencies as in Choo and Siow (2006).

This completes the proof of the Theorem.

Remark 1. Although the cohort preference vector Λ is overidentified, we have exact identification (rather than overidentification) of the surplus matrix conditional on Λ.

Remark 2. It is critical for identification that we have inter-marriage between pre-reform and post-reform cohorts. If cohorts were separate markets—an assumption that is sometimes made in the literature—identification, and disentangling the roles of ability and qualifications, would be impossible. To see this, let us suppose that there are no matches between cohorts 0 and 1. This would correspond to $\lambda_{01} = \lambda_{10} = -\infty$. Thus, a critical step in the proof, the derivation of equation (A.4), would not be possible. Indeed, even if we had access to independent information on Λ, we would be unable to use it. The proof uses matches between type $\tilde{1}$—the unqualified individuals in cohort 1, who are necessarily low ability—to infer the matches involving type 1, the low ability individuals in cohort 0.[39]

[39] In the industrial-organization context, separate markets often aid identification, because supply is typically uniform across markets while demand varies. In our marriage-market context, both sides of the market are affected by the reform and vary across cohorts, so separate markets remove rather than provide identifying variation.

References

Adda, J., P. Pinotti, and G. Tura (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.

Anderberg, D., and Y. Zhu (2014): "What a Difference a Term Makes: The Effect of Educational Attainment on Marital Outcomes in the UK," *Journal of Population Economics*, 27(2), 387–419.

Angrist, J. D., and A. B. Krueger (1991): "Does Compulsory Schooling Attendance Affect Schooling and Earnings," *Quarterly Journal of Economics*, 106(4), 979–1014.

Arcidiacono, P., A. Beauchamp, and M. McElroy (2016): "Terms of Endearment: An Equilibrium Model of Sex and Matching," *Quantitative Economics*, 7(1), 117–156.

Becker, G. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813–846.

Bell, B., R. Costa, and S. Machin (2022): "Why Does Education Reduce Crime?," *Journal of Political Economy*, 130(3), 723–765.

Brandt, L., A. Siow, and C. Vogel (2016): "Large Demographic Shocks and Small Changes in the Marriage Market," *Journal of the European Economic Association*, 14(6), 1437–1468.

Card, D. (2001): "Estimating the Return to Schooling: Progress on some Persistent Econometric Problems," *Econometrica*, 69(5), 1127–1160.

Chevalier, A., C. Harmon, I. Walker, and Y. Zhu (2004): "Does Education Raise Productivity or Just Reflect It?," *Economic Journal*, 114, F499–F517.

Chiappori, P.-A., M. Costa Dias, and C. Meghir (2018): "The Marriage Market, Labor Supply and Education Choice," *Journal of Political Economy*, 126(S1), S26–S72.

Chiappori, P.-A., M. Costa Dias, C. Meghir, and H. Zhang (2025): "Changes in Marital Sorting: Theory and Evidence from the United States," *Journal of Political Economy*, 133(10), 3045–3077.

Chiappori, P.-A., M. Iyigun, and Y. Weiss (2009): "Investment in Schooling and the Marriage Market," *American Economic Review*, 99(5), 1689–1713.

Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012): "Fatter Attraction: Anthropometric and Socieconomic Matching on the Marriage Market," *Journal of Political Economy*, 120(4), 659–695.

Chiappori, P.-A., B. Salanié, and Y. Weiss (2017): "Partner Choice, Investment in Children, and the Marital College Premium," *American Economic Review*, 107(8), 2109–21067.

Choo, E. (2015): "Dynamic Marriage Matching: An Empirical Framework," *Econometrica*, 83(4), 1373–1423.

Choo, E., and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114(1), 175–201.

Clark, D., and H. Royer (2013): "The Effect of Education on Adult Mortality and Health: Evidence from Britain," *American Economic Review*, 103(6), 2087–2120.

Dagsvik, J. K. (2000): "Aggregation in Matching Markets," *International Economic Review*, 41(1), 27–58.

Decker, C., E. H. Lieb, R. J. McCann, and B. K. Stephens (2012): "Unique Equilibria and Substitution Effects in a Stochastic Model of the Marriage Market," *Journal of Economic Theory*, 148(2), 778–792.

Del Bono, E., and F. Galindo-Rueda (2006): "The Long Term Impacts Of Compulsory Schooling: Evidence from a Natural Experiment in School Leaving Dates," Centre for the Economics of Education Discussion paper 0074 16228, London School of Economics.

Devereux, P. J., and R. A. Hart (2010): "Forced to be Rich? Returns to Compulsory Schooling in Britain," *Economic Journal*, 120(549), 1345–1364.

Dickson, M., P. Gregg, and H. Robinson (2016): "Early, Late or Never? When does Parental Education Impact Child Outcomes?," *Economic Journal*, 126, F184–F231.

Dickson, M., and S. Smith (2011): "What Determines the Return to Education: An Extra Year or a Hurdle Cleared?," *Economics of Education Review*, 30(6), 1167–1176.

Dupuy, A., and A. Galichon (2014): "Personality Traits and the Marriage Market," *Journal of Political Economy*, 122(6), 1271–1319.

Eika, L., M. Mogstad, and B. Zafar (2019): "Educational Assortative Mating and Household Income Inequality," *Journal of Political Economy*, 127(6), 2795–2835.

Galichon, A., and B. Salanié (2021): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," CEPR Discussion paper 16228, CEPR.

Gelman, A., and G. Imbens (2019): "Why High-Order Polynomials Should not be Used in Regression Discontinuity Designs," *Journal of Business and Economic Statistics*, forthcoming.

Geruso, M., and H. Royer (2018): "The Impact of Education on Family Formation: Quasi-Experimental Evidence from the UK," Discussion paper, National Bureau of Economic Research.

Grenet, J. (2013): "Is Extending Compulsory Schooling Alone Enough to Raise Earnings? Evidence from French and British Compulsory Schooling Laws," *Scandinavian Journal of Economics*, 115(1), 176–210.

Hahn, J., P. Todd, and W. van der Klaauw (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69(1), 201–209.

Haskey, J. (2001): "Cohabitation in Great Britain: Past, Present and Future Trends–and Attitudes," *Population Trends*, 103(1), 4–25.

Imbens, G., and T. Lemieux (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615–635.

Lee, D. S., and D. Card (2008): "Regression Discontinuity Inference with Specification Error," *Journal of Econometrics*, 142, 655–674.

Lefgren, L., and F. McIntyre (2006): "The Relationship between Women's Education and Marriage Outcomes," *Journal of Labor Economics*, 24(4), 787–830.

Machin, S., O. Marie, and S. Vujić (2011): "The Crime Reducing Effect of Education," *Economic Journal*, 121(552), 463–484.

Mourifié, I., and A. Siow (2021): "The Cobb-Douglas Marriage Matching Function: Marriage Matching with Peer and Scale Effects," *Journal of Labor Economics*, 39(S1), S239–S274.

ONS (2016): "Population Estimates by Marital Status and Living Arrangements, England and Wales: 2002 to 2015," Statistical Bulletin, Office for National Statistics.

Oreopoulos, P. (2006): "Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter," *American Economic Review*, 96(1), 152–175.

Oreopoulos, P., and K. G. Salvanes (2011): "Priceless: The Nonpecuniary Benefits of Schooling," *Journal of Economic Perspectives*, 25(1), 159–184.

Pischke, J.-S., and T. von Wachter (2008): "Zero Returns to Compulsory Schooling in Germany: Evidence and Interpretation," *Review of Economics and Statistics*, 90, 592–598.

Shapley, L., and M. Shubik (1971): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

Silles, M. (2011): "The Effect of Schooling on Teenage Childbearing: Evidence using Changes in Compulsory Education Laws," *Journal of Population Economics*, 24(2), 761–777.