

Marriage market equilibrium, qualifications, and ability

Dan Anderberg, Jesper Bagger, Varsha Bhaskar, Tanya Wilson

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Marriage market equilibrium, qualifications, and ability

Abstract

We study marital sorting on academic qualifications and latent ability in an equilibrium marriage market model using the 1972 UK Raising of the School-Leaving Age (RoSLA) legislation as a natural experiment that induced a sudden, large shift in the distribution of academic qualifications in affected cohorts, but plausibly had no impact on the distribution of ability. We show that a [Choo and Siow \(2006\)](#) model with sorting on cohort, qualifications, and latent ability is identified and estimable using the RoSLA-induced population shifts. We find that the RoSLA isolated low ability individuals in the marriage market, and affected marital outcomes of individuals whose qualification attainment were unaffected. We also decompose the difference in marriage probabilities between unqualified individuals and those with basic qualifications into causal effects stemming from ability and qualification differences. Differences in marriage probabilities are almost entirely driven by ability.

JEL-Codes: D100, D130, I260, J120.

Keywords: marriage, qualifications, assortative mating, latent ability.

Dan Anderberg

*Royal Holloway, Institute for Fiscal Studies
United Kingdom - TW20 0EX Surrey
dan.anderberg@rhul.ac.uk*

Jesper Bagger

*Royal Holloway, Institute for Fiscal Studies
United Kingdom - TW20 0EX Surrey
jesper.bagger@rhul.ac.uk*

Varsha Bhaskar

*University of Texas at Austin
Austin / Texas / USA
v.bhaskar@austin.utexas.edu*

Tanya Wilson

*University of Glasgow
Glasgow / United Kingdom
tanya.wilson@glasgow.ac.uk*

March 15, 2019

We thank Dan Akerberg, Monica Costa Dias, Jeremy Fox, John Knowles, Cezar Santos, attendants at the Barcelona GSE Summer Forum, 2015 Society of Labor Economists conference, the 2016 Notre Dame Economics of the Family conference, the 2017 Economics of Gender workshop in Edinburgh, the 2018 Families and the Macroeconomy workshop in Mannheim, and seminar participants at Bangor, Birkbeck, CESifo, IFS, Iowa State, Kent, Luxembourg University, National University of Singapore, Ohio State, Ragnar Frisch Centre, Royal Holloway, and the University of Texas at Austin for helpful comments. Bhaskar thanks the National Science Foundation for its support via grant # SES 1629055.

1. INTRODUCTION

It is well-established that individuals marry assortatively on qualification attainment (Mare, 1991). Whether this reflects marital surplus complementarities in qualifications or in latent traits correlated with qualification attainment, such as ability, is less clear but nonetheless important. For example, misinterpreting assortative mating on ability as sorting on qualifications results in flawed inference regarding the impacts of education policies and rising educational attainment on intra-household resource allocation, social mobility, and intergenerational transmissions. However, disentangling the roles of qualifications and ability in the marriage market is challenging because ability is latent, i.e. observed by marriage market participants, but not the econometrician, yet correlated with educational attainment, and because the realized marriage market allocation is an equilibrium outcome. Indeed, on the one hand, Choo and Siow (2006) pioneered empirically tractable equilibrium analysis of the marriage market, but their analysis and the literature it spawned has dealt exclusively with sorting on observable traits. On the other hand, the validity of standard reduced form empirical strategies such as regression discontinuity designs, well suited to deal with confounding latent variables, rely on the absence of equilibrium effects.

This paper combines the two approaches and exploit quasi-experimental variation in qualification rates stemming from the 1972 UK Raising of the School-Leaving Age (the RoSLA) legislation to disentangle the roles of academic qualifications and ability in shaping marriage market outcomes within an Choo-Siow equilibrium model of the marriage market. Our analysis makes several contributions. We first document marriage market responses to the RoSLA. Of particular interest is the observation that the RoSLA sharply and permanently increased the never-married rates of the least qualified men and women, even though it dramatically reduced the population supplies of these individuals. A “standard” specification of the Choo-Siow model with marital surplus depending on academic cohorts and qualifications fails to reproduce this response. Second, we show that the RoSLA induced shifts in the qualification distributions permit identification of a richer model where marital surplus also depends on the latent ability of the married individuals; here, ability is correlated with qualification attainment through selection. Structural estimates of the latent ability model reveals that the surplus matrix exhibits complementarities in ability and in qualifications, and confirms that sorting on ability is central to fitting the marriage market response to the RoSLA. Third, using the estimated model, we illustrate how marriage market equilibrium adjustments to RoSLA affected the partner-choices of individuals whose qualification attainments were unaffected, and disentangle the effects of ability and qualifications on marriage probabilities.

The first tier of UK academic qualifications, henceforth denoted a basic qualification, is obtained at the end of the academic year in which a student turns 16. Prior to the RoSLA, around 40 percent of a cohort left school without obtaining a formal academic qualification, and, depending on gender, 30-35 percent with a basic qualification. The rest left with what we label an advanced qualification: a formal academic qualification obtained at age 18 or higher through post-compulsory education. By raising the school leaving age from 15 to 16, RoSLA sharply reduced the likelihood of leaving school with no academic qualifications and a corresponding increase in the likelihood of leaving with a basic qualification. RoSLA had no impact on the rate of leaving with an advanced qualification.

Our first contribution is to document the marriage market response to the RoSLA. We show that the never-married (at age 45) rates of unqualified men and women exhibit a discrete and permanent upward shift in the first RoSLA-treated cohort. No corresponding shifts occurred for those with basic or advanced qualifications. We also confirm overall strong assortative mating on qualifications, and highlight a distinct and permanent increase in assortative mating among those holding no academic qualification after the RoSLA, with no effect on assortative mating among those holding basic or advanced qualifications. Finally, we document a temporary and small shift in the husband-wife age gap distribution: the very first RoSLA-treated academic cohorts were more likely to marry among themselves than earlier and later cohorts.

This latter empirical fact provides the first direct empirical evidence on how age and qualifications are traded off in the marriage market. Such trade-offs are what one would expect based on standard equilibrium marriage market theory (Shapley and Shubik, 1971; Becker, 1973) and confirms that academic qualifications matter in the marriage market. The increase in the never-married rates among those with no qualifications is harder to reconcile with that same body of theory. Consider for example Choo and Siow (2006)’s seminal empirical implementation of the Becker-Shapley-Shubik transferable utility marriage market model. The Choo-Siow framework groups marriage market participants into a finite number of so-called systematic types, with individual-level separable i.i.d. preference shocks over a partner’s systematic type.¹ A fitted Choo-Siow model with systematic types defined as the intersection between cohort and qualification, what we refer to as a standard specification, in fact predicts a sharp downward shift in the never-married rates of the unqualified following RoSLA. The intuition is compelling: RoSLA makes unqualified individuals scarce, which, in equilibrium, inflates both the transfers these individuals extract upon marrying, and the rate at which they marry.²

We take the failure of the standard Choo-Siow specification as evidence that the data features a richer systematic type space. If an additional trait is correlated with qualification attainment, then a shock to the qualification distribution will lead to a shift in the relationship between qualifications and the latent characteristic. Such a compositional effect is not accounted for in the standard specification. In the context of the RoSLA, individual ability is an obvious candidate trait to consider, but ability is a latent variable and such traits are not easily encompassed in the Choo-Siow framework.³ Indeed, the empirical tractability of the Choo-Siow model derives in large parts from an observable systematic type space and unobserved separable individual-level preference heterogeneity, which cannot be interpreted as latent systematic traits.

In our second contribution, we use the RoSLA-induced discontinuous shift in the qualification distribution to identify and estimate a Choo-Siow style marriage market model with a latent

¹Recent papers that follow Choo and Siow (2006) include Dupuy and Galichon (2014), Galichon and Salanié (2015), Choo (2015), Brandt, Siow, and Vogel (2016), Mourifié and Siow (2017), Chiappori, Salanié, and Weiss (2017), and Chiappori, Costa Dias, and Meghir (2018).

²This intuition refers to comparative statics results related to population shifts in the Choo and Siow (2006) model involving one type on one side of the market (see Decker, Lieb, McCann, and Stephens, 2012). The RoSLA-induced population shift, of course, involved more than one type on both sides of the market.

³Based on a dynamic equilibrium marriage market model, Knowles and Vandenbroucke (2018) argues that compositional changes in latent preferences for marriage and fertility among singles resulting from World War I account a large share of the rise in the female post-war marriage probabilities.

systematic type dimension. Our extended systematic type-space is the intersection of academic cohort, qualification, and ability. Ability—which in our model can be low, medium or high—is a latent trait, correlated with observed qualifications through a simple individual-level selection model of qualification attainment. The model interprets the RoSLA as reducing the opportunity cost of obtaining a basic qualification, resulting in a tightening of the selection on ability into qualification-levels: those who complied with the RoSLA by obtaining a basic qualification were medium ability individuals who, in the absence of the reform, would have remained unqualified. This interpretation is corroborated by empirical evidence of selective RoSLA-responses, whereby those who, even after the reform, did not obtain any qualification were negatively selected in terms of socio-economic background characteristics. From a marriage market perspective, the RoSLA therefore has two effects. On the one hand, it made unqualified individuals relatively scarce. On the other hand, it lowered the average ability among the unqualified. If ability carries value in the marriage market, the latter composition effect reduces the marital prospects of the unqualified, thus pushing up their never-married rates.

Our model retains the standard Choo-Siow restrictions: a large marriage market, separability, and i.i.d. Extreme Value Type I distributed preference shocks (Galichon and Salanié, 2015). With a latent systematic type, identification hinges on two additional restrictions: the husband-wife cohort-profile enters additively in the marital surplus, and that the RoSLA shifted the qualification distribution, but not the ability distribution. In addition, we use the easily verifiable data requirement that marriages occur both within and across policy regimes. The first restriction is testable and is not rejected. The second restriction is an untestable exclusion restriction. This identification configuration is novel, yet has some antecedents in the literature. Galichon and Salanié (2015) advocate restrictions on the surplus matrix to facilitate identification and hypothesis testing. Lefgren and McIntyre (2006) and Anderberg and Zhu (2014) use exclusion restrictions arising from particular features of the US and UK school systems to estimate the marital return to education by reduced form instrumental variable regressions. Finally, variation in population supplies in matching markets have been shown to aid identification (Brandt, Siow, and Vogel, 2016, Chiappori, Salanié, and Weiss, 2017, and Fox, Yang, and Hsu, 2018).⁴

We estimate the structural parameters of our proposed marriage market model and find that the estimated model is consistent with observed marriage market behavior around the RoSLA. Specifically, we reproduce the RoSLA-induced upward jump in the never-married rate of unqualified men and women and also the observed increased marital sorting among the unqualified. The estimated marital surplus matrix exhibits complementarities both with respect to ability (when both spouses are unqualified), and with respect to qualification (when both spouses are of medium ability). We formally reject the standard Choo-Siow specification where ability carries no value in the marriage market.

In our third contribution, we use the estimated model to highlight marriage market equilibrium repercussions of the RoSLA through counterfactual simulation, and to quantify the roles

⁴For example, Chiappori, Salanié, and Weiss (2017) assign successive cohorts to different marriage markets. Fox, Yang, and Hsu (2018) derive identification results pertaining to the distribution of latent characteristics in two-sided matching games with transferrable utility, requiring the econometrician to observe many markets.

of qualification attainment and ability in shaping marriage probabilities. We demonstrate that the equilibrium effects of the reform were by no means confined only to individuals whose qualification attainments were directly affected. Indeed, the RoSLA reduced the chances of low and medium ability individuals of ever marrying, and increased it for high ability individuals. Due to the typically positive husband-wife age gap, the effects of the RoSLA were felt by men belonging to pre-RoSLA cohorts. Overall, the reform increased the marital mixing between medium and high ability individuals, but left low ability individuals increasingly isolated in the marriage market.

We decompose the gap in marriage probabilities between unqualified individuals and those with basic qualifications, what we term the marriage gap, into components stemming from ability differences and components stemming from qualification differences. These components represent the causal effects of ability and qualification on marriage probabilities. The estimated effects are similar for men and women. We find that the marriage gap is driven almost entirely by ability. That is, unqualified individuals do not marry less because they are unqualified, but because they have low ability. We also consider the role of abilities and qualifications in explaining the gap in the likelihood of being married to a qualified spouse between the unqualified and individuals basic qualifications, what we term the spousal qualification gap. Here, we find that an individual’s own ability and qualification contribute equally to the observed gap. That is, holding ability fixed, the spousal qualification gap between individuals with a basic qualification and those with no qualifications is half the observed total spousal qualification gap.

We decompose the gap in ever-married rates between individuals with no and basic qualifications—the qualification marriage gap—into components stemming from ability differences and components stemming from qualification differences. These components represent the causal effects of ability and qualification on ever-married rates. The estimated effects are similar for men and women. We find that the marriage gap is driven almost entirely by own ability. That is, unqualified individuals do not marry less because they are unqualified, but because they have low ability. We also consider the role of ability and qualification in explaining the gap in the likelihood of being married to a qualified spouse between individuals with no and basic qualifications—the spousal qualification gap. Here, we find that own ability and own qualification contribute equally to the observed gap. That is, holding ability fixed, the spousal qualification gap between individuals with basic qualification and no qualification is half the observed spousal qualification gap.

The rest of the paper proceeds as follows. In Section 2 we describe the data sources used. In Section 3 we describe the RoSLA reform and outline its impact on academic qualifications. In Section 4 we describe the marriage market outcomes for a set of cohorts born around the RoSLA threshold. In Section 5 we outline the empirical model and the estimation approach. In Section 6 we present the model estimates and fit with the data. In Section 7 we use a counterfactual simulation to highlight how various cohorts and ability types were affected by the reform in terms of their marital outcomes while in Section 8 we consider ability and qualification marital premia. Section 9 concludes.

2. DATA SOURCES

Our empirical analysis combines data from the UK Labour Force Survey (LFS), Population Statistics from the Office for National Statistics (ONS), and the Censuses. Some supplementary analysis, presented in Appendix A, makes use of the Health Survey for England (HSE). Our focus will be on individuals born in different academic cohorts. The academic year runs from the 1st of September to the 31st of August in the following calendar year. We will refer to the 1957 academic cohort—which was the first to be affected by the RoSLA—as those individuals born between September 1957 through August 1958. We focus on a set of academic cohorts $\mathcal{C} = \{1953, \dots, 1960\}$ born around the RoSLA threshold.

2.1. Labour Force Survey

The 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.⁵ The LFS contains information on year and month of birth for each household member and on relationships between household members. Detailed information on qualifications held by each person 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 $c \in \mathcal{C}$.⁶ We use the LFS data to estimate the impact of the RoSLA on academic qualification rates and to characterize the marriage patterns in terms of couples' cohort and qualification profiles.⁷

Table 1 provides descriptive statistics for the full LFS sample by gender and marital status. 68 percent of the men and 70 percent of the observed women are married at the time of the interview. The average age is 39 for both men and women.⁸ The average academic cohort is close to 1956.5 as the observed individuals are nearly uniformly drawn from the cohorts in \mathcal{C} .

The delineation of academic qualification levels will be described in further detail below, but, as noted in the introduction, we will work with three ordered 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 (further described in Section 3). For both males and females, the basic qualification is the most common academic attainment, followed by no academic qualification, and then by advanced qualifications.

⁵At this stage the LFS also became a “rotating panel” whereby each household remains in the survey for five quarters before being replaced. We use information provided by individuals in their first interview.

⁶We restrict attention to England and Wales because other data sources used in the analysis are only available for these constituent countries. We are not conditioning on age. This means that our sample will have an age range of 22 (for someone born in 1961 and observed in 1984) to 61 (for someone born in 1953 and observed in 2014).

⁷The LFS only provides information about the respondent’s current marriage. Hence our analysis will be based on the assumption that the marriage pattern—in terms of spousal characteristics—among currently observed marriages is representative of first marriages.

⁸We make no further use of the age variable as we instead use direct estimates at population-level of the fraction never-married by age 45 by gender, cohort and qualification as described below.

Table 1: Descriptive Statistics for the Pooled LFS Sample of Individuals from Academic Cohorts 1953-1960.

Variable	Males			Females		
	All	Single	Married	All	Single	Married
Age in Years	38.86 (9.11)	38.05 (9.56)	39.25 (8.86)	38.99 (9.12)	39.44 (9.41)	38.80 (8.98)
Ac. Cohort	56.60 (2.29)	56.93 (2.28)	56.44 (2.28)	56.60 (2.30)	56.86 (2.30)	56.49 (2.29)
No Qual.	0.35 (0.48)	0.38 (0.48)	0.34 (0.47)	0.32 (0.47)	0.35 (0.48)	0.31 (0.46)
CSE/O-lev.	0.38 (0.49)	0.36 (0.48)	0.39 (0.49)	0.43 (0.49)	0.40 (0.49)	0.45 (0.50)
A-level+	0.27 (0.44)	0.26 (0.44)	0.27 (0.44)	0.24 (0.43)	0.25 (0.43)	0.24 (0.43)
Obs.	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.

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.⁹ The UK experienced a baby boom that started in the mid-1950s and peaked in 1964. Hence, as a general characterization, the cohorts that 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.¹⁰

3. THE 1972 RAISING OF THE SCHOOL LEAVING AGE

3.1. The Reform

The UK Government introduced Statutory Act 444, known as the Raising of School Leaving Age (RoSLA) Order, in March, 1972. The RoSLA came into operation on September 1st, 1972, rais-

⁹We 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 UK birth cohorts also means that we are ignoring migration when calculating the relative population supplies.

¹⁰First marriages past the age of 45 are rare, whereby the adjustments are very small. 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. As these ONS tables do not contain qualification information the calculation assumes that the rate of entry into first marriage beyond age 45 is homogenous across qualification groups.

ing the minimum school leaving age by one year to age 16, affecting individuals born from 1st September, 1957 onwards. The RoSLA prompted a dramatic 25 percentage points increase in the proportion of individuals leaving education after 16 years of age (see e.g. [Chevalier, Harmon, Walker, and Zhu, 2004](#); [Silles, 2011](#); [Clark and Royer, 2013](#)). The effect of the RoSLA reform 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.

The RoSLA not only impacted the duration of schooling, but also the likelihood of leaving school with an academic qualification. In England and Wales there are two levels of examinations sat during school. The first tier, which we label a basic qualification, leading to the Ordinary Level (O-Level) or Certificate of Secondary Education (CSE) qualifications, are not available until the end of the academic year in which an individual turns 16. If an individual chooses to remain in school after the minimum school leaving age, then after a further two years of study a second set of academic examinations, leading to the Advanced Level (A-Level) qualification, a pre-requisite of entry to higher education, can be taken. That is, the mandated increase in the school leaving age required students to remain in school up to the year in which the first level of academic qualifications are conferred. To foreshadow some of the analysis to come, this feature of the RoSLA motivates our interpretation that, for a large number of individuals, the RoSLA effectively and substantially reduced the opportunity cost of sitting the O-level/CSE examination and attaining a basic qualification.

3.2. The Impact of the RoSLA on Qualifications

As exposure to the RoSLA was determined by a threshold date of birth, it is natural to estimate the reform induced shifts in the qualification rates using a regression discontinuity design (RDD), see e.g. [Hahn, Todd, and der Klaauw \(2001\)](#) and [Imbens and Lemieux \(2008\)](#). Let i index a generic individual and let w_i be individual i 's date of birth. As noted above, there is of a deterministic mapping from date of birth w to academic cohort $c \in \mathcal{C}$.

Let $z_i \in \mathcal{Z}$ denote individual i 's (highest) academic qualification level, and let $y_i^z \equiv \mathbb{1}(z_i = z)$ be an indicator for $z_i = z$. In the potential outcomes framework ([Rubin, 1974](#)), the observed realization of y_i^z is one of two potential outcomes, denoted $y_{i,0}^z$ and $y_{i,1}^z$ each indicating whether z would be i 's qualification level if not exposed to the RoSLA and exposed to the RoSLA, respectively. That is,

$$y_i^z = y_{i,0}^z \mathbb{1}(w_i < 0) + y_{i,1}^z \mathbb{1}(w_i \geq 0). \quad (1)$$

The RoSLA treatment effect on individual i holding qualification z is defined as $\varphi_i^z \equiv y_{i,1}^z - y_{i,0}^z \in \{-1, 0, 1\}$. The average RoSLA treatment effect on holding qualification z among individuals born at time w is defined as the population expectation $\varphi^z(w) \equiv E[\varphi_i^z | w_i = w]$.

If the potential outcome regression functions $E[y_{i,0}^z | w_i = w]$ and $E[y_{i,1}^z | w_i = w]$ are continuous functions of w , the RDD identifies $\varphi^z(0)$ from the discontinuity in the rate of holding qualification z at the reform threshold. Specifically, $\varphi^z(0) = \lim_{w_i \downarrow 0} E[y_i^z | w_i] - \lim_{w_i \uparrow 0} E[y_i^z | w_i]$. For estimating

the size of the discontinuity, we rely on the regression

$$y_i^z = \alpha_0 + \alpha_1 w_i + \alpha_2 w_i \mathbb{1}(w_i \geq 0) + \varphi^z(0) \mathbb{1}(w_i \geq 0) + \epsilon_i, \quad (2)$$

for each $z \in \mathcal{Z} = \{z_0, z_1, z_2\}$. The regression equation (2) includes potentially different linear trends in the running variable w on either side of the threshold $w = 0$, as advocated by [Gelman and Imbens \(2019\)](#). We estimate (2) by gender using the LFS data described above (see [Table 1](#)), following the [Lee and Card \(2008\)](#) procedure for correcting the standard errors on the estimated $\varphi^z(0)$ to account for specification errors in (2) arising from the LFS recording date of birth by month, an inherently discrete variable.¹¹ It is straightforward to add controls for observable regressors in (2). If the continuity assumption underlying the RDD holds, this tend to improve precision, with little impact on point estimates. Our preferred specification omit regressors, but, for completeness, we report results both with and without controls below.

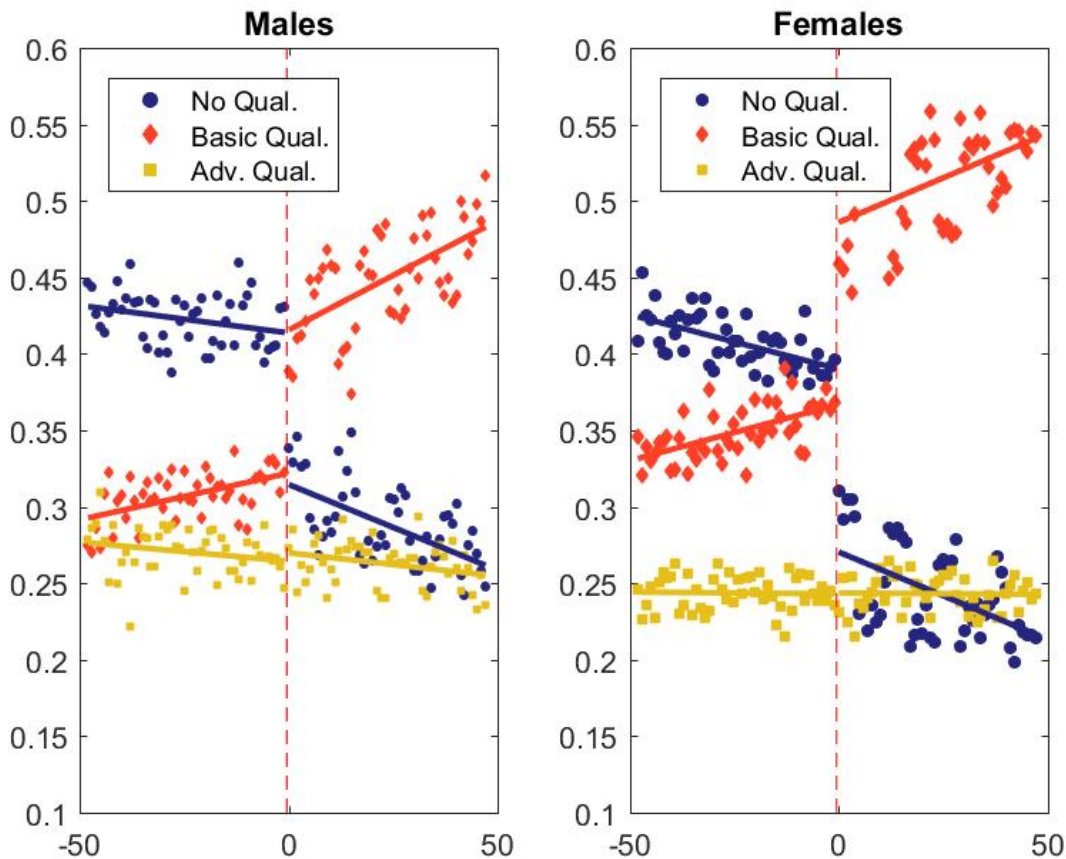


Figure 1: Distribution of Academic Qualifications by Month of Birth Relative to RoSLA Threshold

¹¹[Kolesár and Rothe \(2018\)](#) provide a further thorough analysis of this issue, and suggests an alternative approach. We have confirmed that the two approaches yield very similar results in our case, and that the (clustered) [Lee and Card \(2008\)](#) standard errors represents the conservative choice.

Figure 1 is a graphical rendition of the RDD and the estimated regressions. As our observation period (1953-1960) covers 8 years centered on the RoSLA, there are 48 months of birth on either side to the RoSLA threshold ($w = 0$). For each month of birth the figure plots the observed proportion holding each level of qualification. The figure shows a sharp RoSLA-induced drop in the rate of holding no qualification for each gender with a corresponding discontinuous increase in the proportion holding a basic qualification. The RoSLA did not have any upward spillover effect on the rate of holding an advanced qualification obtained through post-compulsory schooling.

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

	Males		Females	
	(i)	(ii)	(iii)	(iv)
No Qual.	-0.099*** (0.008)	-0.106*** (0.006)	-0.120*** (0.009)	-0.128*** (0.007)
CSE/O-lev.	0.094*** (0.008)	0.106*** (0.005)	0.119*** (0.010)	0.131*** (0.007)
A-Level+	0.006 (0.005)	-0.000 (0.005)	0.001 (0.005)	-0.002 (0.004)
Obs	147,878	147,878	156,549	156,549
Controls	No	Yes	No	Yes

Notes: The sample used in each regression is described in Table 1. Each reported coefficient comes from a separate regression with the dependent variable being a dummy for having that level of academic attainment. The table reports the estimated coefficient on a RoSLA dummy for being born September 1957 or later. Distance of date of birth from the September 1957 threshold measured in months is used as “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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 reports the regression discontinuity estimates of the impact of the RoSLA reform on the qualification distribution, i.e. estimates of $\varphi^z(0)$, by qualification level $z \in \mathcal{Z}$ and gender, and with and without the inclusion of demographic controls. We see that the RoSLA reduced the fraction holding no qualification by about 10 percent for men and by about 12 percent for women, and that the effects are very precisely estimated, with and without controls. These results are in line with previous studies evaluating the RoSLA-effect on qualification rates.¹²

The estimated responses to the RoSLA shows that the reform substantially reduced the proportion of women and men holding no academic qualification. However, even after the reform, a sizeable proportion of both men and women still left school without any formal qualification. In Appendix A we provide evidence of a systematic difference between those who responded to the reform by gaining a qualification and those who did not. Specifically, we use data from the Health

¹²See 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.

Survey for England (HSE) to explore whether the relationship between qualifications and social background changed at the RoSLA threshold. We focus on two markers of social background. First, as there is a strong gradient in mortality, particularly among males (Blane, Smith, and Bartley, 1990), we look at whether the respondent’s natural father had passed away at the time of interview. Similarly, we look at individual own height since, widely regarded as a marker of childhood social conditions (Wadsworth, Hardy, Paul, Marshall, and Cole, 2002), with low achieved height reflecting inadequate nutrition and ill health. We show that, among cohorts born prior to the RoSLA threshold, those who were unqualified were on average shorter and more frequently had fathers who had passed away. More importantly, we show that the gap in both markers between unqualified and qualified individuals increased substantially at the RoSLA threshold. In short, the evidence presented in Appendix A strongly suggests a selective response to the RoSLA whereby those who even after the reform did not obtain any qualification were particularly negatively selected in terms of social background.

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 complied to RoSLA, in the sense of leaving school with an academic qualification, were positively selected in terms of socio-demographic traits. We now document how these population shifts impacted the marriage market.

4.1. Assortative Mating on Qualifications

We confirm that our data features positive assortative mating on qualifications. Figure 2 uses the LFS sample of married individuals. The left panel considers all married males born between 1953 and 1960 and shows the distribution of their wives’ qualification level by the husband’s own qualification level. The right panel provides the corresponding distribution of husbands’ qualifications by the wife’s own qualification for the sample of married women born 1953 to 1960. The high degree of assortative mating is highlighted by the fact that, for each gender and qualification level, the most frequent category is where both spouses have the same qualification.

A more pertinent question is whether the RoSLA affected the degree of assortative mating on qualifications? The first column of Table 3 uses the subsample of married men born between 1953 and 1956 while the second column uses the subsample of married men born between 1957 and 1960. For each subgroup, the table reports the Goodman-Kruskal gamma measure of the rank correlation.¹³ The third and the fourth columns do the same for the subsamples of married women. While Table 3 confirms the high rank correlation in spouses’ qualifications, it does not provide any conclusive evidence that the reform affected the degree of assortative mating on qualifications.

Closer inspection of the data shows that the aggregate rank correlation masks heterogeneous impacts at the various qualification levels. To highlight this, consider the following simple measure

¹³The Goodman-Kruskal gamma measure of the rank correlation is preferred to the Spearman rho and the Kendall tau when the variables in question are ordered categorical and there are many ties as a consequence.

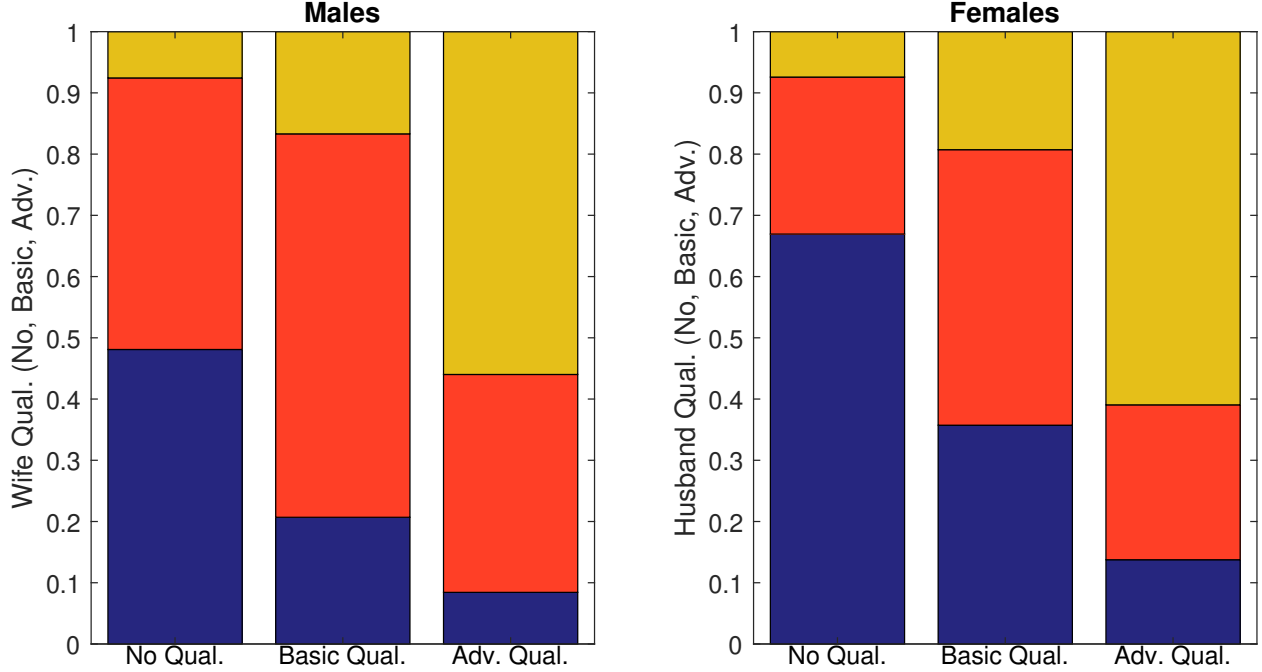


Figure 2: Assortative Matching on Qualifications by Gender

of sorting at qualification level $z \in \mathcal{Z}$ that can be applied in any population of married couples. Let z_i and z_j denote the husband's and the wife's qualification level respectively and define,

$$S(z) \equiv \frac{\Pr(z_i = z, z_j = z)}{\Pr(z_i = z) \Pr(z_j = z)}. \quad (3)$$

The numerator is the probability that, for a randomly drawn couple, both spouses have qualification level z . The denominator is the product of the probabilities of the husband and the wife having qualification level z respectively. If matches were randomly generated, the joint probability would be equal to the product of the marginal probabilities and $S(z)$ would be equal to unity. A value of $S(z)$ above unity thus indicates positive sorting at qualification level z . The advantage of the measure $S(z)$ is that it can be applied for each qualification level separately.¹⁴ We use the LFS sample of married individuals and compute $S(z)$ for each qualification level, by gender and cohort. Specifically, in the left panel of Figure 3, we plot $S(z)$ for each qualification level $z \in \mathcal{Z}$ by the academic cohort $c \in \mathcal{C}$ of the husband.¹⁵ The right panel does the same using the sample of married women.

The figure shows that the strongest assortative mating occurs among those holding an advanced qualification, a pattern that is stable over the cohorts of interest. What is more interesting for our purposes is what happened to the assortative mating among individuals with no qualifications or

¹⁴Indeed, the measure could be applied for any given husband-wife qualification profile. However, our interest here is to explore whether the tendency for married couples to have the same qualification level was strengthened by the RoSLA and, if so, for what qualification level this happened.

¹⁵For the subsample of married couples where the husband is from cohort $c \in \mathcal{C}$ the wife can be from any cohort, including cohorts not in \mathcal{C} .

Table 3: Goodman-Kruskall Rank Correlation in Spouses' Academic Qualification Levels

	Males		Females	
	Pre-Reform (1953-56)	Post-Reform (1957-60)	Pre-Reform (1953-56)	Post-Reform (1957-60)
	0.628*** (0.004)	0.633*** (0.005)	0.638*** (0.004)	0.637*** (0.005)
Obs.	47,419	45,448	50,766	50,057

Notes: The overall sample all includes married couples observed in the Labour Force Survey 1984-2014 with available information on the academic qualification for both spouses. Each column conditions on the individual being born in some academic cohort $c \in \mathcal{C}$, whereas their spouse can be drawn from any cohort. The rank correlation measure provided is the Goodman-Kruskal gamma. Asymptotic standard errors are in parenthesis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

holding a basic qualification. The figure highlights a clear, and permanent, increase in the degree of positive assortative mating among those holding no academic qualification after the RoSLA, both for men and for women. Hence, after the RoSLA, unqualified men and women became increasingly prone to marry each other. In contrast, the degree of assortative mating among those with a basic academic qualification reduced slightly after the reform.

4.2. Never-Married Rates

Figure 4 shows the never-married rate (by age 45) by gender, cohort, and level of qualification calculated from Census data.¹⁶ Overall, the never-married rates exhibit an increasing trend for all qualification levels. This is consistent with gradual decline in marriage rates. We also see that, for each cohort and qualification level, the never-married rate of men is above that of women, again a feature that is consistent with existing empirical evidence.

The most striking feature is how the never-married rate for unqualified individuals increased at the RoSLA threshold. Indeed, for women the first affected academic cohort marks a key turning point. Whereas traditionally, the most qualified women would have been the least likely to marry in their lives, the first RoSLA affected cohort is also the first for which unqualified women were the least likely to ever marry. Among men, the unqualified were already the group least likely to marry in their lives, but at the reform threshold, the gap in the never-married rates for unqualified and qualified rose distinctly.

4.3. Age Gaps

For our purposes, we define the husband-wife age gap as the difference in their academic cohorts, that is $d_{ij} \equiv c_j - c_i$ where c_i and c_j are cohort of the husband and the wife respectively. The left panel of Figure 5 shows the aggregate husband-wife age gap distribution in the academic cohorts of interest based the LFS sample of married individuals.¹⁷ The figure shows that age gaps of 0,

¹⁶See Section 2 for details.

¹⁷The same underlying sample is used also in the right panels of Figure 5 and in Figure 6 but with indicated conditioning.

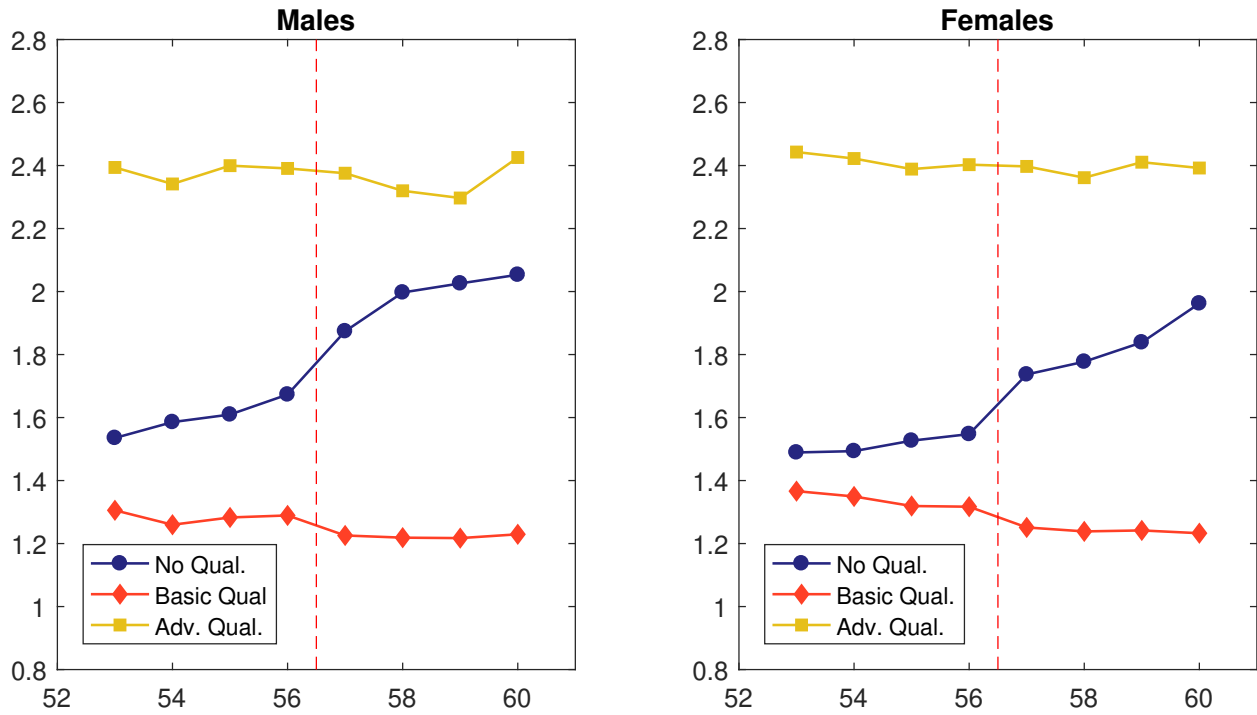


Figure 3: Assortative Matching by Cohort, Gender, and Qualification Level

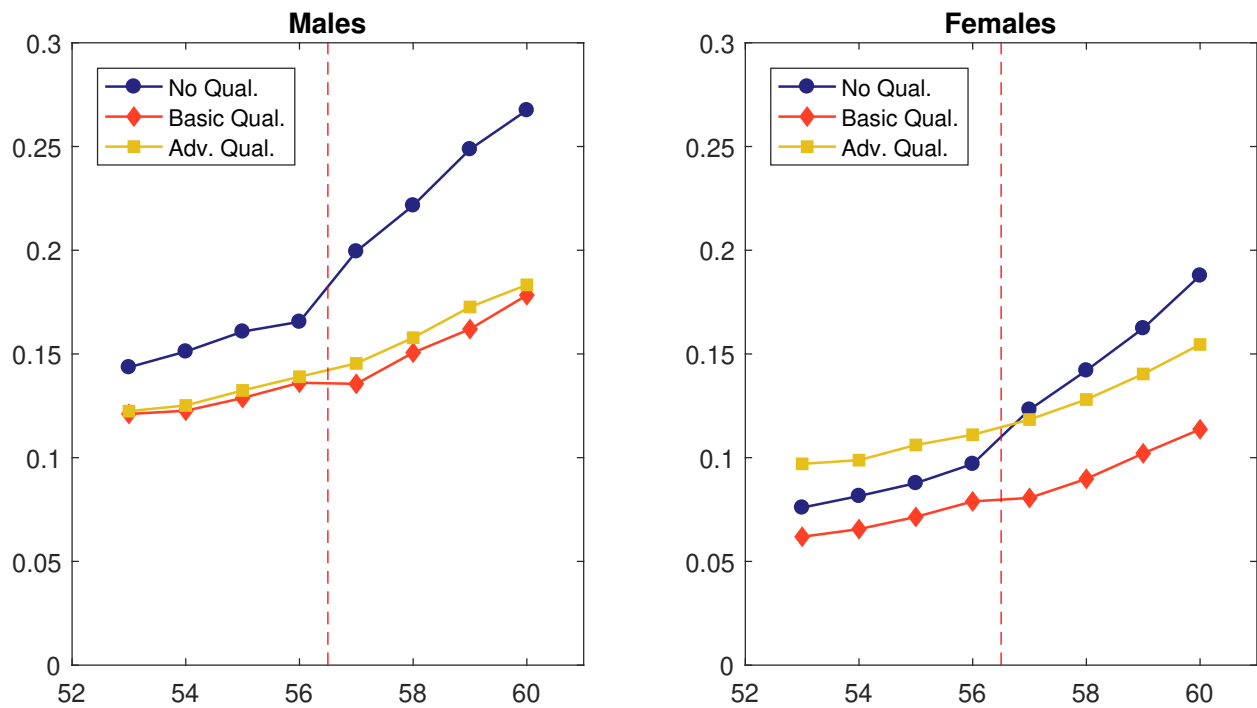


Figure 4: Never-Married (by Age 45) Rates by Cohort, Gender, and Academic Qualification Level

1 and 2 are the most common. There is a sharp drop in frequency when moving to negative age gaps, but a fat right tail for positive age gaps.

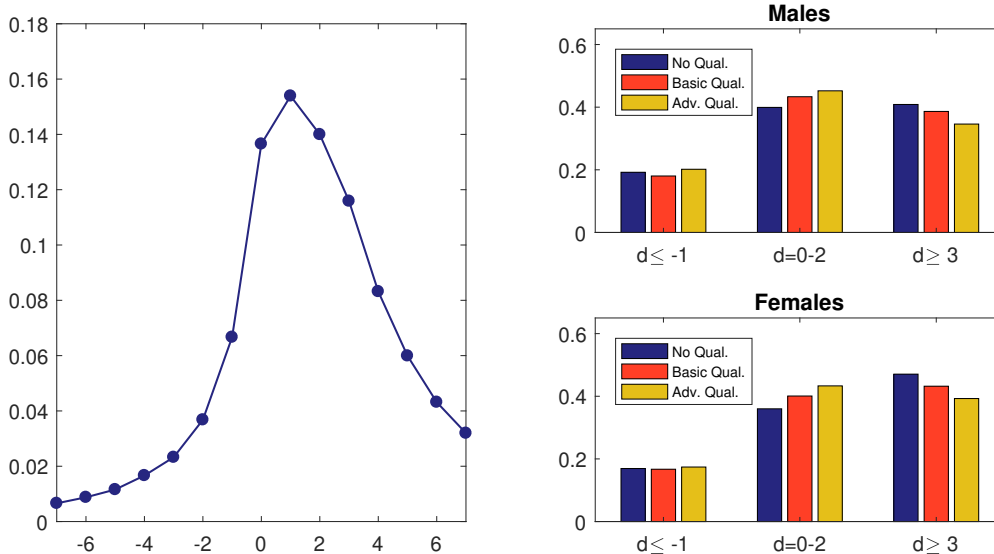


Figure 5: The Aggregate Husband-Wife Age Gap Distribution and Age Gap Distribution by Gender and Qualification

In the right panels of Figure 5 we highlight the age gap distribution, by gender and qualification. The distribution of husband-wife age gaps for married men born in some $c \in \mathcal{C}$ is displayed in the upper panel, while the lower panel does the same for married women.¹⁸ The figure suggests that there is no particular qualification pattern among individuals who choose to marry with a negative age gap. Among those marrying with non-negative age gaps, larger positive age gaps are slightly more common among less qualified individuals.¹⁹

In order to explore whether the age gap distribution was affected by the RoSLA, we consider how the age gap distribution differed in the key cohorts around the reform threshold from the aggregate distribution. For every age gap $d \in \{-3, +3\}$ and for each cohort $c \in \mathcal{C}$ we regress the indicator $\mathbb{1}(d_{ij} = d)$ on a cohort-dummy for being from cohort c . This way we determine, for each d , whether individuals from cohort c had a different likelihood of being married with this particular age gap compared to individuals from all other cohorts in \mathcal{C} .

We graphically present the estimated coefficients from these regressions in Figure 6. 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

¹⁸There are two reasons why the two panels are not identical. First, while marriages are assortative on qualifications they are not perfectly so. Second, the figure does not restrict the spouse to be born in the cohorts of interest as that would have biased the shape of the empirical age gap distribution.

¹⁹Mansour and McKinnish (2014) find evidence that men and women who are married to differently aged spouses are negatively selected on a range of characteristics, including cognitive ability, educational attainment and wages.

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 here thus suggest that the RoSLA temporarily—but only modestly—shifted the age gap distribution, with the early RoSLA-affected men and women more frequently choosing to marry each other.²⁰

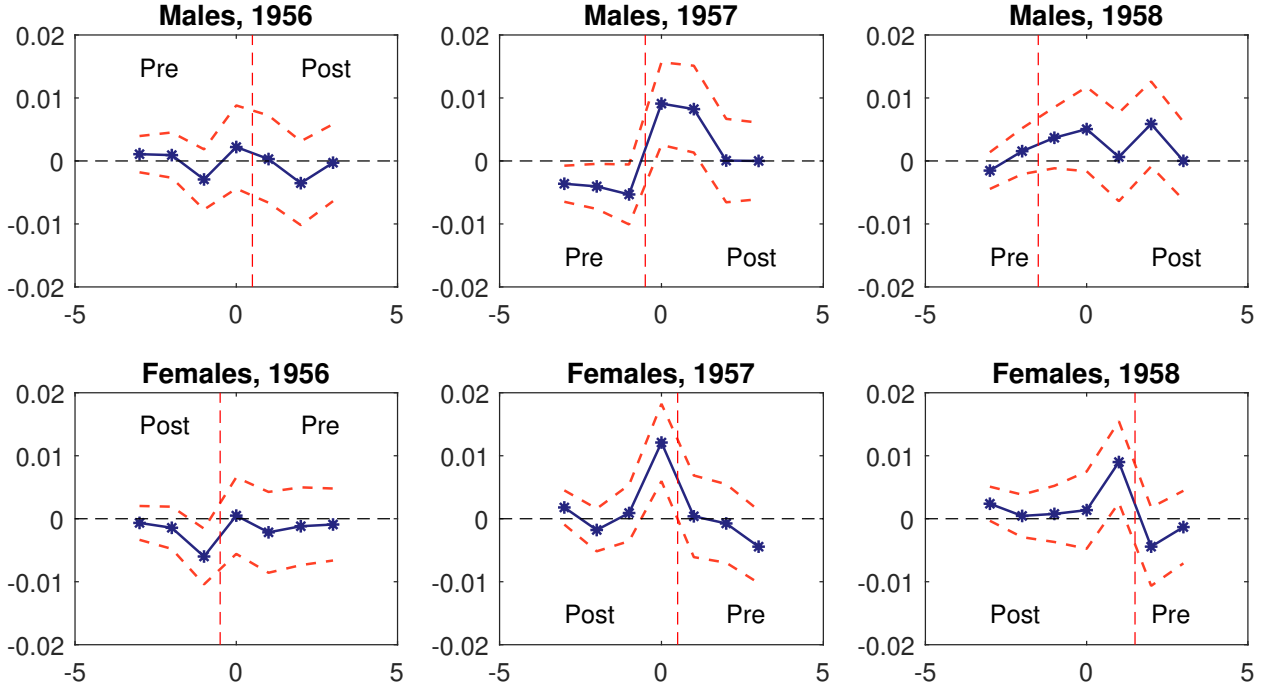


Figure 6: Deviations from the Aggregate Age Gap Distribution by Cohort and Gender

5. MODEL

5.1. The Choo-Siow Framework

We will build on [Choo and Siow \(2006\)](#), who developed a framework for estimating the structural parameters of the transferable utility marriage market model of [Shapley and Shubik \(1971\)](#) and [Becker \(1973\)](#). Let \mathcal{I} and \mathcal{J} be bounded real intervals that denote the continuum sets of men and women respectively. Each individual belongs to one of a finite set of systematic types \mathcal{X} .²¹ The measure of different types of men (women) is denoted $h^m(x)$ ($h^f(x)$). Let $\mathcal{X}_+ = \mathcal{X} \cup \{0\}$ so that an individual has the option of remaining single, by matching with “type 0”. Systematic types are observed by all marriage market participants, and if a type- $x_i \in \mathcal{X}$ of man matches with a type- $x_j \in \mathcal{X}$ woman, this results in a systematic match surplus $\Sigma(x_i, x_j)$. If an individual chooses to remain single, the systematic payoff that he or she obtains is normalized to zero so

²⁰[Geruso and Royer \(2018\)](#) conduct a similar exercise, also in the context of RoSLA, and their findings are largely consistent with ours.

²¹The systematic type space could be different for men and women, but it will not be so in our context. Hence to save notation we ignore this possibility here.

that $\Sigma(x_i, 0) = \Sigma(0, x_j) = 0$. Man $i \in \mathcal{I}$ is also subject to a vector of preference shocks, or preference heterogeneity, $\varepsilon_i = (\varepsilon_i(x))_{x \in \mathcal{X}_+}$. Similarly, woman $j \in \mathcal{J}$ is subject to a vector of shocks $\varepsilon_j = (\varepsilon_j(x))_{x \in \mathcal{X}_+}$. The total surplus generated by a match between man i and woman j is

$$\Sigma(x_i, x_j) + \varepsilon_i(x_j) + \varepsilon_j(x_i). \quad (4)$$

The preference heterogeneity components $\varepsilon_i(x_j)$ and $\varepsilon_j(x_i)$ are separable (Galichon and Salanié, 2015). That is, they enter (4) additively, ruling out interactions between a particular man's and a particular woman's idiosyncratic utilities in shaping marital utility, and depend only on the partner's systematic type, not his or her identity.²² Transferable utility implies that the surplus σ_{ij} can be divided between the two parties in any way the couple chooses.

A matching is a measurable function, μ , from $\mathcal{I} \cup \mathcal{J}$ to itself that satisfies (i) If $i \in \mathcal{I}$ ($j \in \mathcal{J}$) and $\mu(i) \in \mathcal{I}$ ($\mu(j) \in \mathcal{J}$), then $\mu(i) = i$ ($\mu(j) = j$),²³ (ii) $\mu(\mu(i)) = i$ (respectively $\mu(\mu(j)) = j$), (iii) μ is measure-preserving.²⁴ Associated with any matching is a payoff allocation, specifying the payoffs obtained by each individual in $\mathcal{I} \cup \mathcal{J}$.

Chiappori, Salanié, and Weiss (2017) show that, under separability, in any stable matching equilibrium, (i) there exists two mappings $U : \mathcal{X} \times \mathcal{X}_+ \rightarrow \mathbb{R}$ and $V : \mathcal{X}_+ \times \mathcal{X} \rightarrow \mathbb{R}$ such that $U(x_i, x_j) + V(x_i, x_j) = \Sigma(x_i, x_j)$ for any $(x_i, x_j) \in \mathcal{X} \times \mathcal{X}$, and $U(x_i, 0) = 0$ and $V(0, x_j) = 0$, (ii) man i of type x_i achieves utility $u_i = \max_{x \in \mathcal{X}_+} \{U(x_i, x) + \varepsilon_i(x)\}$ and makes the choice that attains the maximum, and (iii) woman j of type x_j achieves utility $v_j = \max_{x \in \mathcal{X}_+} \{V(x, x_j) + \varepsilon_j(x)\}$ and correspondingly makes the choice that attains the maximum.²⁵ $U(x_i, x_j)$ specifies the systematic component of the payoff to a type- x_i man when he matches with a type- x_j woman (while his systematic payoff from remaining single is 0). Similarly, $V(x_i, x_j)$ specifies the systematic component of the payoff to a type- x_j woman when she matches with a type- x_i man.

Let $\mu_{x_j|x_i}^m$ and $\mu_{0|x_i}^m$ denote the probabilities that a type- x_i male marries a type- x_j female and that he remains unmarried respectively. Correspondingly let $\mu_{x_i|x_j}^f$ and $\mu_{0|x_j}^f$ denote the probabilities that a type- x_j female marries a type- x_i male and remains unmarried respectively. Following Choo and Siow (2006), we assume that preference shocks $\varepsilon_j(x_i)$ and $\varepsilon_i(x_j)$ are i.i.d. Type I Extreme Value distributed.²⁶ It is then straightforward to show that

$$\log \left(\frac{\mu_{x_j|x_i}^m}{\mu_{0|x_i}^m} \right) = U(x_i, x_j), \text{ and } \log \left(\frac{\mu_{x_i|x_j}^f}{\mu_{0|x_j}^f} \right) = V(x_i, x_j). \quad (5)$$

Market clearing requires that the measure of type- x_i men who choose marry type- x_j women equals

²²Specifically, man $i \in \mathcal{I}$ contributes random utility $\varepsilon_i(x_j)$ in a marriage with *any* woman of type $x_j \in \mathcal{X}$ and obtains random utility $\varepsilon_i(0)$ remaining unmarried. Similarly, woman j contributes random utility $\varepsilon_j(x_i)$ in a marriage with *any* male of type $x_i \in \mathcal{X}$ and obtains $\varepsilon_j(0)$ remaining unmarried.

²³ $\mu(i) = i$ means that man $i \in \mathcal{I}$ is single and similarly for $j \in \mathcal{J}$ if $\mu(j) = j$.

²⁴That is, if \mathcal{K} is a measurable subset of either \mathcal{I} or \mathcal{J} , the Lebesgue measures of \mathcal{K} and $\mu(\mathcal{K})$ are equal.

²⁵Choo and Siow (2006) obtained this result in the special case where the preference shocks ε_i and ε_j are Extreme Value Type I distributed.

²⁶A notable exception is Galichon and Salanié (2015), who show that the Choo-Siow framework remains tractable with minimum distributional assumptions on ε_i and ε_j .

the measure of type- x_j women who choose to marry type- x_i men, i.e. that

$$h^m(x_i) \mu_{x_j|x_i}^m = h^f(x_i) \mu_{x_i|x_j}^f, \quad (6)$$

holds for each type profile $(x_i, x_j) \in \mathcal{X} \times \mathcal{X}$. Furthermore, adding up implies

$$\sum_{x_j \in \mathcal{X}_+} \mu_{x_j|x_i}^m = 1 \quad \text{and} \quad \sum_{x_i \in \mathcal{X}_+} \mu_{x_i|x_j}^f = 1, \quad (7)$$

for every $x_i \in \mathcal{X}$ and every $x_j \in \mathcal{X}$, respectively.

Under the assumption that the type of every individual is observed by all market participants, existence and uniqueness of a stable matching is guaranteed (Decker, Lieb, McCann, and Stephens, 2012).²⁷ In addition, the equilibrium exhibits natural comparative statics properties, most notably with respect to the population distribution. As an example, consider a decrease in the supply of males of type x_i . As $h^m(x_i)$ decreases men of this type will obtain a larger share of the marital surplus from a marriage to any type of woman; as a result men of type x_i should become more likely to marry relative to remaining unmarried, i.e. $\mu_{0|x_i}^m$ falls.

Separability is key for empirical tractability of the model. With separable preference shocks, competition implies that the (many) type x_j -women who wants to marry a type- x_i man bid up the utility that man i obtains to the point where he fully retains $\varepsilon_i(x_j)$. An analogous argument of course applies to woman j , who fully retains $\varepsilon_j(x_i)$. As a result the stable matching is characterized by individual-level rational choices over a partner's systematic type, which greatly facilitates empirical implementation.²⁸ However, it also implies that marital sorting arises only in terms systematic types. Furthermore, key identification results in Choo and Siow (2006) and Galichon and Salanié (2015) show that the structural parameters $\Sigma(x_i, x_j)$, $(x_i, x_j) \in \mathcal{X} \times \mathcal{X}$ are identified under the assumption that the econometrician observes the systematic types of the market participants, and therefore effectively restrict attention to sorting on observables.²⁹ This is unfortunate, not least because the empirical marriage market literature typically considers only relatively low-dimensional systematic type-spaces, leaving plenty of room for latent traits to contaminate the analysis.³⁰ A key contribution of this paper is to use the RoSLA to identify marital sorting on a latent trait, in our case, ability, without dispensing with the tractability afforded by separability. As we shall see, extending the model to incorporate latent traits is necessary for fitting the observed marriage market responses to the RoSLA documented in section 4.

²⁷Note that existence and uniqueness applies equally to our generalized framework with a latent ability, since we shall assume throughout that market participants observe the types of all potential partners.

²⁸As an aside, note that it does not matter whether ε_i and ε_j are observed by market participants other than the individual or not, since this payoff shock accrues fully to the individual.

²⁹The logic behind these identification results are straightforward. Suppose the econometrician observes a random sample from the population. As the number of observed males of type- $x_i \in \mathcal{X}$ grows, by the Glivenko-Cantelli Theorem, the empirical proportion observed to marry type- x_j women converges almost surely to $\mu_{x_j|x_i}^m$ and the empirical proportion observed to remain single converges almost surely to $\mu_{0|x_i}^m$. It follows immediately from (5) that $U(x_i, x_j)$ is identified. $V(x_i, x_j)$ is similarly identified, as is then $\Sigma(x_i, x_j)$.

³⁰Notable exceptions include Orefice and Quintana-Domeque (2010), Chiappori, Orefice, 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.

5.2. The Model with Latent Ability

In our empirical setting, an individual's type x has three dimensions. First, by her date of birth, an individual belongs to an academic cohort $c \in \mathcal{C}$. Second, she has some ability level $a \in \mathcal{A}$. Finally, she holds some academic qualification level, $z \in \mathcal{Z}$. Hence an individual's type is a triple $x = (c, a, z) \in \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$. In our empirical implementation we define three ability levels $\mathcal{A} = \{a_0, a_1, a_2\}$ which we refer to as low, medium, and high ability, respectively. Similarly, as above there are three qualification levels, $\mathcal{Z} = \{z_0, z_1, z_2\}$, referred to as no qualification, a basic qualification, and an advanced qualification, respectively. The set of academic cohorts \mathcal{C} is split into the pre-RoSLA cohorts $\mathcal{C}_0 = \{1953, \dots, 1956\}$ and the post-RoSLA cohorts $\mathcal{C}_1 = \{1957, \dots, 1960\}$.

We assume that any individual type, including ability, is observed by prospective partners in the marriage market. However, contrary to the conventional specification of the Choo-Siow model, ability is a latent trait, not observed by the econometrician. As a result, the standard identification arguments referenced above no longer applies.

A key step in establishing identification in the extended model relates to how the RoSLA affected the mapping from ability to academic qualifications. In order to obtain a basic qualification z_1 , an individual has to do two things: attend school until age 16 and pass the exams held at the end of that school year. While passing the basic qualifying exams relies on ability, attending school entails an opportunity cost in terms of forgone labor earnings. Pre-RoSLA, an individual would forgo labor market earnings if attending school beyond age 15 while post-RoSLA, there is no option of working for a 15-year old. Hence, the RoSLA dramatically reduced the opportunity cost of acquiring a basic qualification without altering the ability required to pass the exam and obtain qualification-level z_1 . This implies that, pre-RoSLA, an individual could fail to obtain an academic qualification either due to low ability or because she had alternative opportunities. Post-RoSLA, low ability is the only barrier to gaining a basic academic qualification. This logic justifies our Assumption 1.

Assumption 1. Any individual of low ability a_0 is always unqualified z_0 , and any individual with high ability a_2 always holds an advanced qualification z_2 , whether treated or not by the RoSLA. Any individual of medium ability a_1 , if treated by the RoSLA, holds a basic qualification z_1 , regardless of gender. Any individual of gender g and of medium ability a_1 , if not treated to the RoSLA, holds a basic qualification z_1 with probability γ^g and is unqualified z_0 with probability $1 - \gamma^g$.

The assumption is illustrated in Figure 7. By Assumption 1, only four individual ability-qualification profiles can exist: (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 v. medium) among the unqualified, and there is variation in qualifications (unqualified v. basic) among medium ability individuals.³¹ Let $\mathcal{X} \subset \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$ be the set of all full types that can arise. Note that, since four ability-qualification profiles exist in the pre-reform cohorts \mathcal{X}_0 but only three exist in the post-reform cohorts \mathcal{C}_1 , we have that $|\mathcal{X}| = 4 \times |\mathcal{C}_0| + 3 \times |\mathcal{C}_1| = 28$ types exist for each gender.

³¹We are thus implicitly assuming that, in the pre-RoSLA regime, the choice of medium-ability individuals between staying in school to gain a basic qualification and leaving early without one was triggered by some further unobserved heterogeneity, which is unrelated to potential marital surplus, such as short-term local labor market conditions.

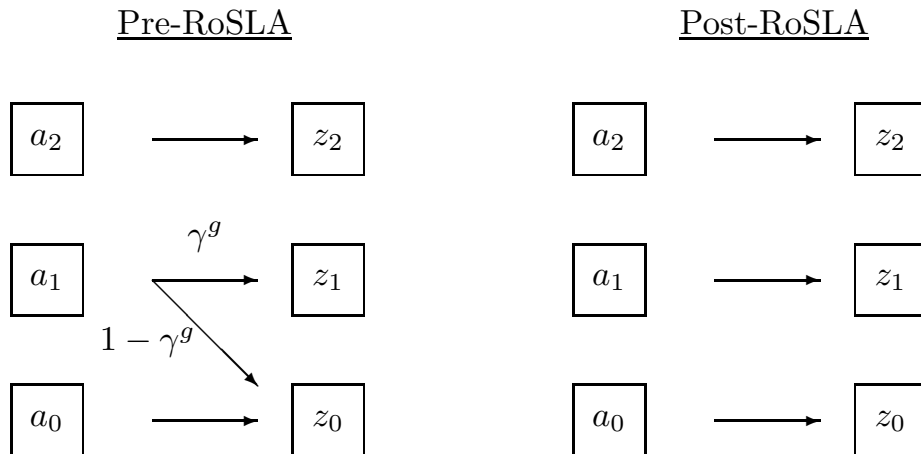


Figure 7: Assumed Relation between Ability and Qualification by Education Regime

Furthermore, we assume that the share of the different ability types in the population does not change discontinuously across the RoSLA threshold.³² In that case, Assumption 1 allows us to infer the total numbers of low ability men (and women) in the pre-RoSLA cohorts. It also allows us to infer the full type of any individual who belongs to a post-RoSLA cohort, given that the mapping between ability to qualifications is one-to-one in the post-RoSLA period. However, we still do not observe the ability of an individual who is unqualified and who belongs to a pre-RoSLA cohort. In consequence, we cannot observe the number of matches involving, for example, a low-ability man in a pre-RoSLA cohort, and nor can we observe the number of such men who remain single. Similarly, we cannot observe the number of medium-ability women who are unqualified, and nor can we observe their singles rate. To illustrate, if there were only two cohorts, one pre-RoSLA and one post-RoSLA, and there were no individuals with high ability and advanced qualifications, there would be 5 feasible types for each gender, and thus possible 25 couple type-profiles, plus 10 singles rates. However, we would not be able to observe 20 of these 35 matching- and singles rates.

A key result of this paper is to establish identification by invoking the following assumption regarding the structure of the systematic component of the marital surplus $\Sigma(x_i, x_j)$. In particular, we will assume additive separability between a couple's ability and qualification profile on the one hand and their cohort profile on the other.

Assumption 2. The systematic surplus function $\Sigma(x_i, x_j)$ is additively separable,

$$\Sigma(c_i, a_i, z_i; c_j, a_j, z_j) = \zeta(a_i, z_i; a_j, z_j) + \lambda(c_i, c_j). \quad (8)$$

In the presence of unobserved ability, we are unable to observe the matching pattern in the marriage market. In particular, in any match involving an unqualified individual pre-RoSLA, we cannot infer his or her ability. Nonetheless, by using the separability assumption, we can infer the

³²Whilst 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. Assumption 1 is then also naturally consistent with a widening of the gap in ability/social background between qualified and unqualified individuals at the RoSLA threshold as suggested by the evidence put forward in Appendix A.

full matching pattern. That is, in the case with only two cohorts previously discussed, we would be able to infer all 35 elements of the matching matrix.

Proposition 1. *Suppose that (i) there are two cohorts, one pre-RoSLA cohort c_0 and one post-RoSLA cohort c_1 , (ii) there are two ability levels $\{a_0, a_1\}$ and two qualification levels $\{z_0, z_1\}$, and (iii) the ability distribution is the same in both cohorts. Under Assumptions 1 and 2, we can infer the full matching pattern, and thereby identify $\Sigma(x_i, x_j)$.*

The intuition for the proof (provided in Appendix B) is as follows. First, we show that the cohort-profile preferences $\lambda(c_i, c_j)$ are identified from the cohort-profiles of married couples where both spouses have a basic qualification. Using $\lambda(\cdot)$ we then show that it is possible to uniquely recover the number of singles of the types where this number cannot be directly observed – that is, the (a_0, z_0) and (a_1, z_0) -types in the pre-reform cohort. Having recovered the never-married rates we then show that the number of marriages between any two types can be inferred. Once the matching pattern has been recovered, it follows from Choo and Siow (2006) that the surplus function is identified.

Since the RoSLA reform induced a marked change in the qualification distribution a potentially instinctive approach would be to treat individuals from the pre- and post-RoSLA regimes as belonging to separate marriage markets and to estimate the marriage surplus function using a multi-market approach in the spirit of, say, Chiappori, Salanié, and Weiss (2017). However, it is key to our identification result to treat pre and post-RoSLA individuals as belonging to a single marriage market. If we treat them as belonging to separate marriage markets, then identification is not possible, as we see in Appendix B.

5.3. Empirical Specification

Our full empirical model amends the $2 \times 2 \times 2$ version of the model considered in Proposition 1 in two ways: we include (i) the high ability type, a_2 , who by Assumption 1 always acquires an advanced academic qualification, z_2 , and (ii) not just two but eight cohorts, four on either side of the reform threshold.

Including a higher ability/qualification type is empirically important as they are a sizeable group and are frequently marriage partners to unqualified and basic qualified individuals from the sample cohorts: Figure 1 highlights that the proportion of the population holding an advanced qualification is about 25 percent in every cohort while Figure 2 shows that about 40 percent of (married) advanced qualified individuals are married to partners who are either unqualified or basic qualified. As a type they are a distinct, but stable, group who acquired an advanced qualification through post-compulsory education. Furthermore, including the advanced qualified type in the analysis does not raise any identification issues as they are, per assumption, always directly distinguishable in the data.

Including further cohorts raises two issues. First, the RoSLA reform will only identify γ^m and γ^f at the reform threshold. Hence any potential cohort variation in γ^m and γ^f (within \mathcal{C}_0) could not be empirically separated from cohort variation in the frequency of ability level a_1 versus a_0 .

Assumption 1 makes it clear that we treat γ^m and γ^f as time-invariant within the four pre-reform periods included in the empirical analysis. Supporting this is the observation from Figure 1 that the trends in the rate of holding a basic qualification are similar before and after the reform both for men and for women, increasing by on average 0.5-1 percentage point per year. This suggests a slow but steady trend towards improved ability at the lower tail end both for men and women.

The second issue raised by the inclusion of multiple cohorts is that there is evidence of trends in the never-married (see Figure 4). Ignoring such trends would risk mis-characterizing the discontinuity in the marriage behavior at the RoSLA threshold which is central to the identification of the surplus structure. Furthermore, with $|\mathcal{C}| = 8$ there are $|\mathcal{C}|^2 = 64$ possible couple cohort-profiles. It is then natural to restrict how marital surplus depends on cohort profile (c_i, c_j) . For these reasons we postulate a surplus structure of the following form,

$$\Sigma(c_i, a_i, z_i; c_j, a_j, z_j) = \zeta(a_i, z_i; a_j, z_j) + \lambda(c_j - c_i) + \tau^m(c_i, a_i) + \tau^f(c_j, a_j). \quad (9)$$

Compared to equation (8) this imposes one restriction and one generalization. It restricts $\lambda(\cdot)$ to depend on the husband-wife age-gap, d_{ij} , rather than on their full cohort profile. Specifically we estimate the following form,

$$\lambda(d_{ij}) = \sum_{d=-3}^3 \beta^d \mathbb{1}(d_{ij} = d) + (\beta_0^- + \beta_1^- d_{ij}) \mathbb{1}(d_{ij} \leq -4) + (\beta_0^+ + \beta_1^+ d_{ij}) \mathbb{1}(d_{ij} \geq 4). \quad (10)$$

$\lambda(\cdot)$ is fully non-parametric at age-gaps around zero but then, for simplicity, imposes linear terms at age gaps outside the range $\{-3, -2, \dots, 3\}$.³³ We normalize $\beta^0 = 0$, implying that $\lambda(0) = 0$, such that there are 10 parameters to be estimated in $\lambda(\cdot)$.

As a generalization we allow for trend terms to vary by gender and ability. With only a small number of cohorts on either side of the threshold, we model these in the simplest possible way as piece-wise linear,

$$\tau^g(c_i, a_i) = \sum_{a \in \mathcal{A}} \left[\beta_a^g (c_i - 1953) \mathbb{1}(a_i = a) + \beta_{a, \mathcal{C}_1}^g (c_i - 1956) \mathbb{1}(a_i = a) \mathbb{1}(c_i \in \mathcal{C}_1) \right], \quad g = m, f. \quad (11)$$

While the inclusion of ability-specific trends formally violates the additive structure in Assumption 2, the terms in (11) will be identified from the observable trends by qualification and gender. The inclusion of the 12 trend terms will help fit the marriage data around the reform threshold and avoid confounding trends with threshold discontinuities.

Since only four (a, z) combinations exist, $\zeta(\cdot)$, the part of the marital surplus function of central interest, can be represented as a 4×4 matrix, adding 16 parameters to be estimated. Overall the postulated surplus function has a total of 38 parameters to be estimated.

Before turning to the estimation, we test the additive structure imposed in (9) using data on cohorts in \mathcal{C}_1 . Consider a man from cohort $c_i \in \mathcal{C}_1$ with qualification z_i and a woman from cohort

³³Note that the set of age-gaps that can occur between two spouses with $c_i, c_j \in \mathcal{C}$ is $d_{ij} \in \{1 - |\mathcal{C}|, \dots, |\mathcal{C}| - 1\}$. In Section 5.4 we outline how marriages to spouses from non-sample cohorts are handled.

$c_j \in \mathcal{C}_1$ woman with qualification z_j . As both are from the post-RoSLA cohorts, by Assumption 1, the qualifications imply aligned ability levels, a_i and a_j , as illustrated by Figure 7. Then, given the linear trend specification in (11), the marital surplus function (9) implies a zero double-difference in cohorts.³⁴ Specifically,

$$\begin{aligned} & [\Sigma(c_i, a_i, z_i; c_j, a_j, z_j, c_j) - \Sigma(c_i - 1, a_i, z_i; c_j - 1, a_j, z_j)] \\ & - [\Sigma(c_i - 1, a_i, z_i; c_j - 1, a_j, z_j) - \Sigma(c_i - 2, a_i, z_i; c_j - 2, a_j, z_j)] = 0. \end{aligned} \quad (12)$$

In post-reform cohorts, this double-difference is observable: Using that the terms in (5) sum up to the systematic surplus, we have that

$$\log \left(\frac{\mu_{c_j, a_j, z_j | c_i, a_i, z_i}^m}{\mu_{0 | c_i, a_i, z_i}^m} \right) + \log \left(\frac{\mu_{c_i, a_i, z_i | c_j, a_j, z_j}^f}{\mu_{0 | c_j, a_j, z_j}^f} \right) = \Sigma(c_i, a_i, z_i; c_j, a_j, z_j), \quad (13)$$

where the marriage and never-married rates on the left-hand side are observed. Hence, the double-differenced marital surplus can be computed by double-differencing the corresponding empirical sum of log-ratios, thus rendering the restriction (12) testable.

As all cohorts included in the double-difference have to be in \mathcal{C}_1 , the possible base years are 1960 and 1959. This leaves four possible male-female base year combinations, and as there are nine possible husband-wife qualification profiles, a total of 36 post-reform double differences can be computed. All 36 estimates are plotted in Figure 8 along with 95 percent confidence intervals.³⁵ Despite fairly high precisions in some categories, the hypothesis of a zero double-difference in marital surplus is rejected in only two cases out of the 36 and with no particular pattern emerging. We conclude that the marriage data from the post-reform cohorts does not reject our specification with additive age-gap-based preferences and linear trends.

5.4. Estimation

The model is parameterized such that the unknown parameters are γ^g , $g = m, f$, the share of gender- g , ability- a_1 individuals who attain basic qualifications in the pre-RoSLA regime (see Assumption 1), and the parameter vector θ , which contains the 10 age gap parameters from (10), the 12 trend parameters from (11), and the 16 parameters in the ability-qualification marital surplus function, $\zeta(a, z)$. Estimation proceeds in two steps. In the first step, we estimate γ^g , $g = m, f$, which allows us to recover the full population distribution of systematic types from the data. In light of Assumption 1, identification of γ^g obtains from the RoSLA regression discontinuity design estimator discussed and presented in section 2. The second step proceeds with estimation of θ , by way of a maximum likelihood estimation procedure. Details of the second step of the estimation procedure are presented in Appendix C.

³⁴In the more general case where $\lambda(\cdot)$ depends on the cohort profile, $\lambda(c_i, c_j)$, the double difference would equal $\lambda(c_i, c_j) - 2\lambda(c_i - 1, c_j - 1) + \lambda(c_i - 2, c_j - 2)$ and should thus be constant across qualification profiles.

³⁵Within each subfigure, the husband-wife base year combinations, reading from top to bottom, are $\{(1960, 1960), (1960, 1959), (1959, 1960), (1959, 1959)\}$.

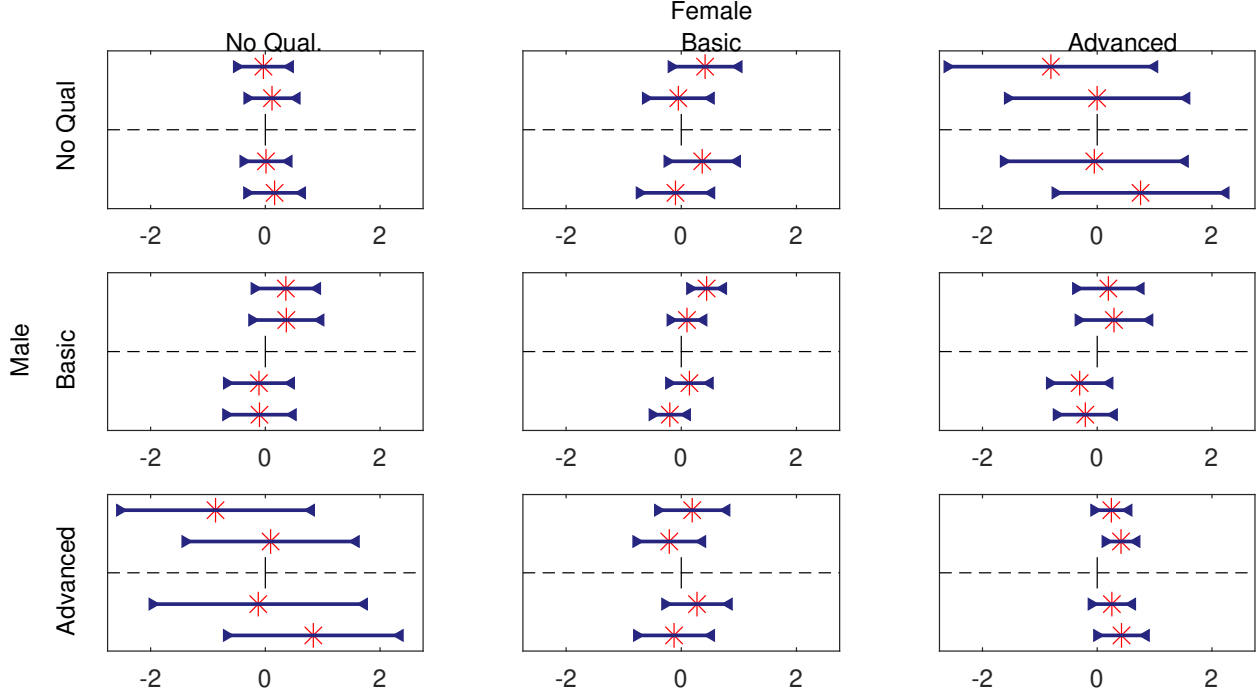


Figure 8: Specification Tests Based on Marital Surplus Double Differencing

5.4.1. *First step: Population distribution of systematic types.* We first give a structural interpretation to the regression discontinuity estimates of the RoSLA effects on qualifications presented in Section 3.2. This allows us to recover the population type distribution. Recall that w_i denotes an individual's date of birth, normalized to zero on the RoSLA threshold, and that y_i^z is an indicator for individual i holding qualification level $z \in \mathcal{Z}$, with $y_{i,0}^z$ and $y_{i,1}^z$ being the corresponding potential outcomes; that is, indicators for whether individual i would hold qualification z if not exposed and exposed to the RoSLA, respectively.

Assumption 1 specifies the ability-conditional population mean of each potential outcome—that is, for each $z \in \mathcal{Z}$ and each gender, conditional on each ability level $a \in \mathcal{A}$ —as some constant. I.e., among males (m) and for a basic qualification (z_1), $E^m[y_{i,0}^{z_1}|a_i = a_1] = \gamma^m$ whilst $E^m[y_{i,1}^{z_1}|a_i = a_1] = 1$, and analogously for women. Note that the regression discontinuity estimation required that the population expectations of the potential outcomes be continuous with respect to w (see Section 3.2). Hence we consider what further model assumption is required to ensure this. Since each ability-conditional potential outcome expectation (for any $z \in \mathcal{Z}$, gender $g = m, f$, ability $a \in \mathcal{A}$) is independent w_i it also follows that the (male) population expectations of the potential outcomes as functions of time of birth can be written as

$$E^m[y_i^z|w_i = w] = \sum_{a \in \mathcal{A}} E^m[y_i^z|a_i = a] \Pr^m(a_i = a|w_i = w), \quad (14)$$

for $z \in \mathcal{Z}$ and $k = 0, 1$, where $\Pr^m(a_i = a|w_i = w)$ is the proportion of males born at time w who have ability level $a \in \mathcal{A}$. It therefore also immediately follows that a sufficient condition for

$E^m [y_i^z | w_i = w]$ to be a continuous function of w is that the distribution of ability among men is continuous in w . A similar expression and logic holds for females. Hence we impose the following continuity assumption.

Assumption 3. $\Pr^g(a|w)$ is a continuous function of w for each ability level $a \in \mathcal{A}$, and for $g = m, f$.

As noted in Section 3.2, given continuity of the population expectations of the potential outcomes with respect to w the RDD identifies the average RoSLA treatment effect on the holding of qualification $z \in \mathcal{Z}$ among individuals born at $w = 0$. Specifically, for males and a basic qualification z_1 , using Assumption 1 and equation (14), we have that $\varphi_m^{z_1}(0) = (1 - \gamma^m) \Pr^m(a_i = a_1 | w_i = 0)$. It then follows that γ^m can be recovered from the RDD estimate and the rate of holding qualification z_1 after the reform: rearranging, and using that, by Assumption 1, $\Pr^m(a_i = a_1 | w_i = w) = E^m [y_i^{z_1} | w_i = w]$ for any $w \geq 0$, it follows that

$$\gamma^m = 1 - \frac{\varphi_m^{z_1}(0)}{\lim_{w \downarrow 0} E^m [y_i^{z_1} | w_i = w]}. \quad (15)$$

Identification of γ^f is established in an analogous way.³⁶

Having recovered γ^m and γ^f , and using that, by Assumption 1, each γ^g is constant during our pre-RoSLA observation period, we can back out the population measures of (partly latent) systematic types, $h^g(c, a, z)$ for every $(c, a, z) \in \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$ and $g = m, f$ from the observed cohort sizes and qualification distributions by gender and cohort.

5.4.2. Second step: Structural marriage market model parameters. In the second step, we employ a maximum likelihood estimator of θ . The likelihood function must be built from observable systematic type-specific marriage and never-married probabilities. Latent ability implies that the observable type-space is an aggregation of the full systematic type-space. Proposition 1, the empirically relevant extensions, and the first stage identification and estimation of the population distribution of systematic types, shows that the RoSLA allows recovery of the full set of type-specific marriage and never-married probabilities from the smaller set of observed marriage and never-married probabilities. That is, the likelihood function contains enough information to estimate the full set of structural parameters, including all entries in the ability-qualification marital surplus matrix.

The empirical model has $|\mathcal{X}| = 28$ partly unobserved full types making choices from the set \mathcal{X}_+ . To account for marriages involving a partner not belonging to our cohorts of interest \mathcal{C} , we extend the definition of \mathcal{X}_+ slightly to be $\mathcal{X}_+ \equiv \mathcal{X} \cup \{0, \text{pre}, \text{post}\}$, where “pre” and “post” indicate marriages to partners born prior to 1953 and after 1960 respectively. Emphasizing the dependence on θ , let $\mu_{x_j|x_i}^m(\theta)$ be the equilibrium probability that a male of type $x_i \in \mathcal{X}$ makes the marriage choice $x_j \in \mathcal{X}_+$ and, similarly, let $\mu_{x_i|x_j}^f(\theta)$ be the equilibrium probability that a female of type x_j makes the marriage choice $x_i \in \mathcal{X}_+$. Given the choice probabilities afforded by the logit structure, it

³⁶We bootstrap the standard errors on γ^g to account for estimation errors in $\varphi_g^{z_1}(0)$ as well as sampling and estimation errors in $\Pr^g(a_1|w = 0) = \lim_{w \downarrow 0} E^m [y_i^{z_1} | w_i = w]$.

is straightforward to compute the equilibrium choice probabilities for any value of θ by application of a simple Newton algorithm. For every candidate θ , our algorithm approximates pre- and post-spouse marriage probabilities using the relative weight put on out-of-sample marriages by the age gap function (10) implied by θ .

The equilibrium in full types $x \in \mathcal{X}$ is then aggregated up to characterize the implied choice frequencies in terms of observable types. Let $\tilde{x} \in \tilde{\mathcal{X}}$ denote an observable type, where $\tilde{\mathcal{X}} \equiv \mathcal{C} \times \mathcal{Z}$. Observable choices are in the set $\tilde{\mathcal{X}}_+ \equiv \tilde{\mathcal{X}} \cup \{0, \text{pre}, \text{post}\}$. Hence we correspondingly let $\tilde{\mu}_{\tilde{x}_j|\tilde{x}_i}^m(\theta)$ denote the equilibrium probability that a male of observable type $\tilde{x}_i \in \tilde{\mathcal{X}}$ has the observable marital outcome $\tilde{x}_j \in \tilde{\mathcal{X}}_+$. Female probabilities in terms of observables are similarly denoted $\tilde{\mu}_{\tilde{x}_i|\tilde{x}_j}^f(\theta)$. The following equations describe the aggregation from equilibrium choices in full type to observable spouse types,

$$\tilde{\mu}_{c_j, z_j | c_i, z_i}^m(\theta) = \sum_{a_i \in \mathcal{A}} \sum_{a_j \in \mathcal{A}} \mu_{c_j, a_j, z_j | c_i, a_i, z_i}^m(\theta) \Pr^m(a_i | c_i, z_i), \quad (16)$$

$$\tilde{\mu}_{c_i, z_i | c_j, z_j}^f(\theta) = \sum_{a_j \in \mathcal{A}} \sum_{a_i \in \mathcal{A}} \mu_{c_i, a_i, z_i | c_j, a_j, z_j}^f(\theta) \Pr^f(a_j | c_j, z_j), \quad (17)$$

where $\Pr^g(a|c, z)$ is the population proportion of gender- g individuals of cohort c and with qualification z who are of ability a . These objects are obtained from the first step in the estimation procedure, making the aggregations in (16) and (17) empirically operational. Probabilities for observable types for choosing singlehood and a pre- or post-sample spouse are obtained in analogous fashions.

Our data on marital choices can be represented by gender-specific matrices,

$$\tilde{\mathbf{M}}^g = \{M_{\tilde{x}, \tilde{x}'}^g : (\tilde{x}, \tilde{x}') \in \tilde{\mathcal{X}} \times \tilde{\mathcal{X}}_+\} \text{ for } g = m, f. \quad (18)$$

Here, $M_{\tilde{x}, \tilde{x}'}^g$ is the number of observed individuals who are of gender g and observable type $\tilde{x} \in \tilde{\mathcal{X}}$ and having made the observable marriage choice $\tilde{x}' \in \tilde{\mathcal{X}}_+$. The log-likelihood of the parameter vector θ given the observable marriage data $\tilde{\mathbf{M}}^m$ and $\tilde{\mathbf{M}}^f$ is given by

$$\ell(\theta | \tilde{\mathbf{M}}^m, \tilde{\mathbf{M}}^f) = \sum_{g \in \{m, f\}} \sum_{\tilde{x} \in \tilde{\mathcal{X}}} \sum_{\tilde{x}' \in \tilde{\mathcal{X}}_+} M_{\tilde{x}, \tilde{x}'}^g \log \tilde{\mu}_{\tilde{x}'|\tilde{x}}^g(\theta), \quad (19)$$

where the $\tilde{\mu}_{\tilde{x}'|\tilde{x}}^g(\theta)$ is the probability that an observable type- $\tilde{x}^g \in \tilde{\mathcal{X}}$ chooses observable partner-type $\tilde{x}' \in \tilde{\mathcal{X}}_+$, i.e. including singlehood. For partner-choice $x' \in \tilde{\mathcal{X}} \subset \tilde{\mathcal{X}}_+$, $\tilde{\mu}_{\tilde{x}'|\tilde{x}}^g(\theta)$ are given by (16) and (17); the remaining equilibrium marriage choices are described in Appendix C. Let $\hat{\theta}$ be the Maximum Likelihood estimator. Regularity conditions and standard arguments implies that $\hat{\theta}$ is consistent and asymptotically Normal distributed. We report standard errors based on the asymptotic variance-covariance matrix of $\hat{\theta}$.³⁷

³⁷Although the reported standard errors do not take account of the two-step estimation procedure, we argue that the influence of first step estimation errors in the second step is likely to be small as the first step estimates are obtained with high precision. Furthermore, as the second step parameters are very precisely estimated, correction for the first step estimation errors is unlikely to impact inference.

6. RESULTS AND MODEL FIT

The first step estimates of γ^m and γ^f were obtained from the estimates of the impact of the RoSLA on the rate of holding an academic qualification presented in Table 2. The point estimates for men and women were $\hat{\gamma}^m = 0.755$ (0.011) and $\hat{\gamma}^f = 0.738$ (0.008) respectively. We use these estimates to recover $h^g(x)$ from the observed cohort sizes and distributions of qualifications by cohort and gender displayed in Figure 1.³⁸ As γ^m and γ^f are very precisely estimated we take them, and the implied distributions of full type, as parametrically given.³⁹

Table 4: Estimates of Contribution of Ability-Qualification Profile to Marital Surplus and the Marginal Contribution of Ability and Qualification to Marital Surplus

Panel A: Marital Surplus by Ability-Qualification Profile				
	Females	Low Ab.,	Medium Ab.,	High Ab.,
Males	No Qual.	No Qual.	Basic Qual.	Adv. Qual.
Low Ab.,	-0.580***	-3.061***	-2.030***	-5.042***
No Qual.	(0.222)	(0.544)	(0.139)	(0.159)
Medium Ab.,	-2.697***	-1.472**	-2.234***	-4.921***
No Qual.	(0.559)	(0.669)	(0.407)	(0.443)
Medium Ab.,	-2.500***	-2.501***	-1.133***	-3.064***
Basic Qual.	(0.197)	(0.338)	(0.083)	(0.090)
High Ab.,	-4.656***	-4.271***	-2.310***	-0.498***
Adv. Qual.	(0.208)	(0.355)	(0.089)	(0.090)

Panel B: Value of Ability and Qualification				
	Male:		Female:	
Spouse	Ability	Qual.	Ability	Qual.
Low Ab.,	-2.117***	0.197	-2.481***	1.031*
No Qual.	(0.633)	(0.522)	(0.681)	(0.529)
Medium Ab.,	1.589*	-1.029*	1.225	-0.762
No Qual.	(0.943)	(0.597)	(0.821)	(0.545)
Medium Ab.,	-0.204	1.101***	-0.001	1.368***
Basic Qual.	(0.514)	(0.397)	(0.511)	(0.327)
High Ab.,	0.122	1.857***	0.385	1.961***
Adv. Qual.	(0.565)	(0.435)	(0.532)	(0.344)

Notes: The estimation sample pools all individuals observed in the 1984-2014 LFS from academic cohorts 1953-1960 with non-missing information on age, qualification and marital status to characterize the qualification distribution by gender and cohort, and to characterize the marriage frequencies by husband-wife cohort- and qualification profile. Cohort sizes are based on birth statistics and the never-married rates by gender and qualification level are based on Census data as outlined in Section 2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Turning to the second step estimates, the top panel of Table 4 gives the estimates of the

³⁸Specifically, we use that, for any $c \in \mathcal{C}_0$, $\Pr^g(a_0|c) = \Pr^g(z_0|c) - [(1 - \gamma^g)/\gamma^g] \Pr^g(z_1|c)$ and $\Pr^g(a_1|c) = \Pr^g(z_1|c)/\gamma^g$ to recover the conditional distribution of unobserved ability from the observed conditional distributions of qualifications.

³⁹As robustness we have conducted our empirical analysis for values of each γ^g between 0.715 and 0.775. This affects the point estimates of the elements of the marital surplus matrix, but does not qualitatively affect the conclusions. Details are available on request from the authors.

$\zeta(\cdot)$ matrix defined over ability-qualification profiles. The upper left 2×2 sub-matrix gives the estimated surpluses from marriages where both spouses are unqualified, but with low or medium ability. This matrix trivially exhibits increasing differences. Among the unqualified there is strong complementarity with respect to ability. Similarly, the centre 2×2 sub-matrix gives the estimated surpluses from marriages where both spouses have medium ability, but with or without a basic qualification. This submatrix too exhibits increasing differences, indicating strong complementarity with respect to holding a qualification. If we then disregard the 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.

As the model is estimated using maximum likelihood we can formally test whether the inclusion of ability as a latent characteristic that contributes to systematic marital surplus is statistically justified using a standard likelihood ratio test. Specifically, consider the restriction where the ζ -matrix in (9) is invariant with respect to ability and thus only vary with the couple’s qualification profile. This would impose equality between the first and second column the matrix in top panel of Table 4 and also equality between the first and second row, a total of seven restrictions on $\zeta(\cdot)$. Estimating the restricted model gives a Likelihood Ratio test statistic of $LR = 215.6$. Using that LR is chi-square distributed with seven degrees of freedom, the restrictions can be rejected at any standard level of statistical significance ($p < 0.001$).

It should be noted that the “restricted” model used in this LR -test still has trend terms defined on ability as in (11). In Appendix D we hence re-estimate the restricted model after further redefining also the trend terms to depend on qualifications rather than on ability, thus leaving us with a version of the model where types are defined exclusively in terms of cohort and qualification. We show that this further reduces the fit to the data as measured by the maximized likelihood value. Moreover, the so-restricted model systematically mispredicts the behavior of the marriage market equilibrium at the RoSLA threshold, as we discuss below.⁴⁰

From the estimates in the top panel of Table 4 we can obtain estimates of the contribution of ability of unqualified men to marital surplus by taking the difference between the second and the first row. The contribution of ability of unqualified women is correspondingly given by the difference between the second and first column. Similarly, estimates of the contribution of a basic qualification to marital surplus—given medium ability—are obtained by taking the difference between the third and the second row for men, and between the third and second column for women. In doing so, we obtain estimates of the contribution of (medium) ability and a (basic) qualification by gender, and by the ability-qualification type of the spouse. These estimated values are highlighted in the lower panel of Table 4.

Some distinct patterns emerge among these estimates. First, the results are fairly similar for

⁴⁰As an alternative approach we have estimated a simple “before and after” version of the model where we define the marital surplus on qualification profile only, $\xi(z_i, z_j)$, but allowed this function to change from before to after the reform. Using observed marriages where both partners are born before and after the reform respectively to estimate such regime-specific surplus functions suggests a distinct and unattractive change in $\xi(z_i, z_j)$, particularly a lowering of the surplus for marriages involving individuals with no qualifications.

Table 5: Estimates of Marital Surplus: Age Gap and Trend Terms.

Part A: Age Gap Function, $\lambda(c_j - c_i)$					
β_{-3}	β_{-2}	β_{-1}	β_{+1}	β_{+2}	β_{+3}
-3.579***	-2.628***	-1.435***	0.249***	0.055***	-0.300***
(0.042)	(0.032)	(0.024)	(0.019)	(0.020)	(0.022)
β_0^-	β_1^-	β_0^+	β_1^+		
-3.147***	0.286***	0.088	-0.284***		
(0.163)	(0.030)	(0.074)	(0.013)		
Part B: Trend Functions, $\tau^g(c; a), g = m, f$					
$\beta_{a_0}^m$	β_{a_0, C_1}^m	$\beta_{a_0}^f$	β_{a_0, C_1}^f	$\beta_{a_1}^m$	β_{a_1, C_1}^m
-0.076***	-0.109***	-0.187***	-0.054	0.066***	-0.136***
(0.023)	(0.032)	(0.036)	(0.044)	(0.015)	(0.028)
$\beta_{a_1}^f$	β_{a_1, C_1}^f	$\beta_{a_2}^m$	β_{a_2, C_1}^m	$\beta_{a_2}^f$	β_{a_2, C_1}^f
0.014	-0.136***	-0.039**	-0.065*	-0.085***	-0.087**
(0.017)	(0.031)	(0.017)	(0.036)	(0.019)	(0.039)

Notes: See notes to Table 5 for sample used and text for the specifications of estimated functions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

men and women. A basic qualification always increases marital surplus when the spouse has at least a basic qualification, but does not necessarily do so when the spouse is unqualified. Ability by itself does not increase marital surplus when the spouse is academically qualified; in contrast, when the spouse is unqualified, the complementarity in ability means that it increases surplus when the spouse also has medium ability but decreases it when the spouse only has low ability, though only the latter difference is statistically significant.

The estimated contributions of age gaps and trends to the marital surplus are presented in Table 5. The top panel reports the estimated parameters of the age gap function $\lambda(\cdot)$ while the lower panel reports the estimated parameters of the trend functions $\tau^g(c; a)$, $g = m, f$. The estimated $\lambda(\cdot)$ is closely related to the empirical age gap distribution (see Figure 11). Most estimated trend parameters in Panel B are negative, which is consistent with never-married rates increasing across cohorts.

Next we consider model fit. First, the model replicates the overall assortative mating on qualifications well: Figure 9 shows the model-predicted version of the empirical distributions in Figure 2. Second, the model also replicates the increase in assortative mating among unqualified individuals after the reform: Figure 10 shows the predicted versions of the sorting measure $S(z)$, defined in (3), by gender and cohort as hatched lines and with solid lines still representing the empirical data. The estimated model predicts increases in $S(z_0)$ for both men and women at the reform threshold and of empirically reasonable values.

Next, with respect to the estimated model's fit to the age gap distribution, the left panel of Figure 11 plots the predicted age gap distribution alongside the empirical one for the central values of -4 to $+4$, showing a very close fit. The right panel plots the estimated $\lambda(\cdot)$ in exponential form, highlighting the tight connection between $\lambda(\cdot)$ and the observable age gap distribution.⁴¹

Turning to never-married rates, we see in Figure 12 that the estimated model replicates the

⁴¹Note that the base category of zero age gap has $\exp(\lambda(0)) = 1$ since $\lambda(0)$ is normalized to zero.

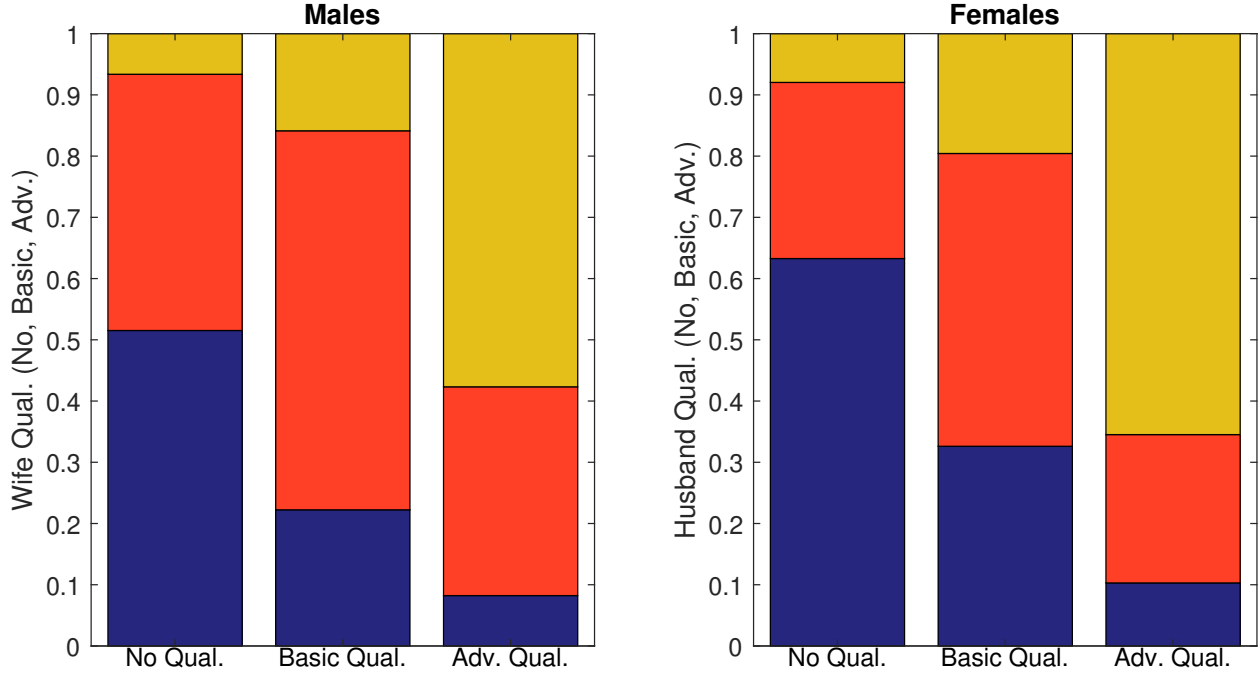


Figure 9: Model Predicted Assortative Matching on Qualifications

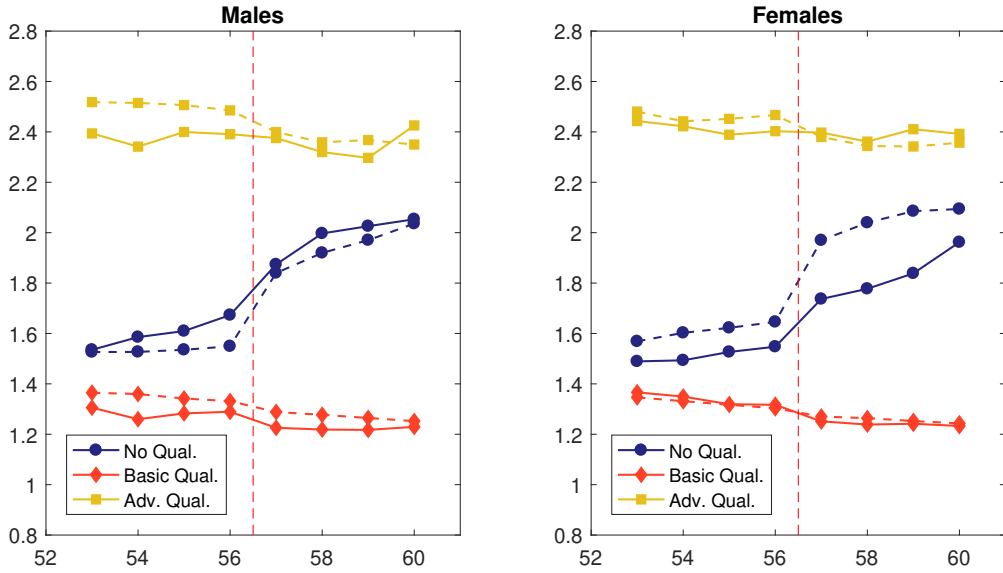


Figure 10: Model Predicted and Empirical Assortative Matching by Cohort, Gender, and Qualification Level

key qualitative patterns in the never-married rates across cohorts and qualifications. Specifically, the model predicts increasing never-married rates for both unqualified men and women at the reform threshold. Allowing for latent ability is central in this respect. To highlight this, we present an estimated version of our model without ability in Appendix D, the “standard” Choos-Siow specification alluded to previously in the paper. Looking specifically at the model fit from this constrained model (Figure D.1), it predicts a sharp reduction in the never-married rates of

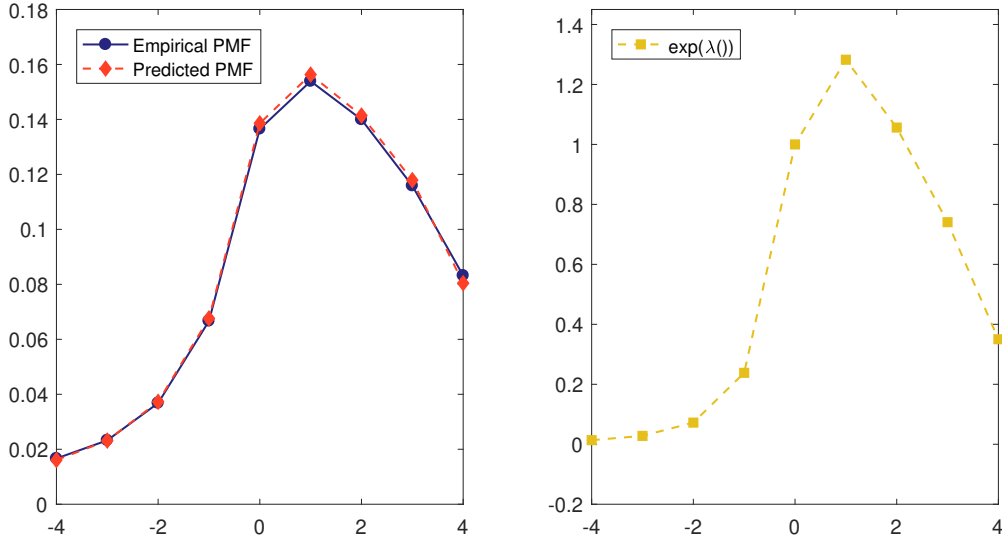


Figure 11: Model Predicted and Empirical Age Gap Distribution and Estimated Contribution of Age Gaps to Marital Surplus

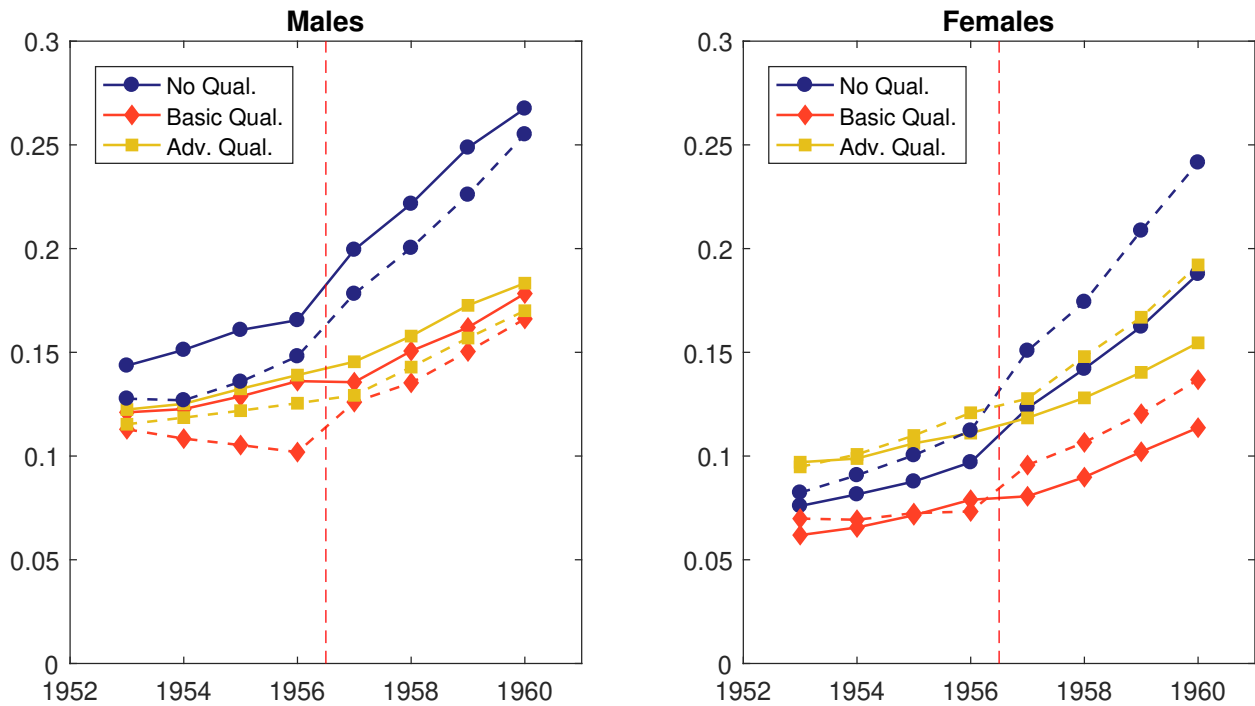


Figure 12: Model Predicted and Empirical Never-Married Rates by Cohort, Gender, and Qualification Level

unqualified individuals at the reform threshold driven by the sudden reduction in their supply.

Our model also fits 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 rates at the RoSLA of individuals holding a basic

qualification due to their increased supply.⁴² However, there is little evidence of such an effect in the data.

7. MARRIAGE MARKET EFFECTS OF THE ROSLA

In this section we show that the RoSLA affected the marital outcomes of a wider set of people, including individuals whose educational choices and outcomes were not directly affected by the reform. Specifically, focusing on the probability of ever-marrying and on marital sorting, we contrast predicted marriage outcomes simulated from our estimated model to those obtained from a counterfactual simulation where the RoSLA reform was never implemented.

In the counterfactual scenario we assume that the pre-reform mapping from ability to qualifications illustrated in Figure 7 continued to apply across all the sample cohorts $c \in \mathcal{C}$, providing us with a counterfactual distribution of qualifications. We then use the counterfactual distribution of full types along with the estimated marriage surplus parameters to compute the predicted counterfactual marriage market equilibrium. We emphasize ability types as individual ability was not affected by the RoSLA reform, and is in this sense a more fundamental individual characteristic than the correlated academic qualification.

Consider first how the reform affected the proportion never-married by ability type, gender and cohort. Figure 13 plots the difference in never-married rates with and without the RoSLA induced population-shifts. Hence, a positive number for a particular ability-type, gender and cohort in Figure 13 implies that the RoSLA increased the never-married rates of individuals of that particular ability-type, gender and cohort.

Consider first the medium ability individuals. As this was the only ability type directly affected by the reform in terms of their academic qualifications, this is also the only type for which the impact of the reform was discontinuous at the threshold. The model suggests that the reform significantly increased the never-married rates for medium ability individuals of both genders. This partly reflects that the reform created a positive supply shock for basic qualified individuals. But it also partly reflects a pre-existing marital qualification premium whereby, even before the reform, a basic qualification given medium ability was associated with a reduced probability of marriage for men but for not women. We return to these qualification premia in Section 8.

The other two ability types were not directly affected by the reform in terms of their academic attainment, but may have been indirectly affected in terms of their marital outcomes. Before the reform, low ability men and women frequently married unqualified medium-ability partners. When this type disappeared as a result of the RoSLA-reform, the low ability individuals lost a natural choice of marital partner and, as a result, their never-married rate increased due to the reform. Note that the typically positive husband-wife age gap meant that the reform increased the never-married rate for low ability males born even before the threshold whereas for women the effect was concentrated among the post-reform cohorts.

For high ability individuals, who always hold an advanced qualification, the medium-ability type when unqualified was never an attractive marriage partner. Instead, when more individuals of

⁴²This feature is shared also by the model without unobserved ability as illustrated in Figure D.1.

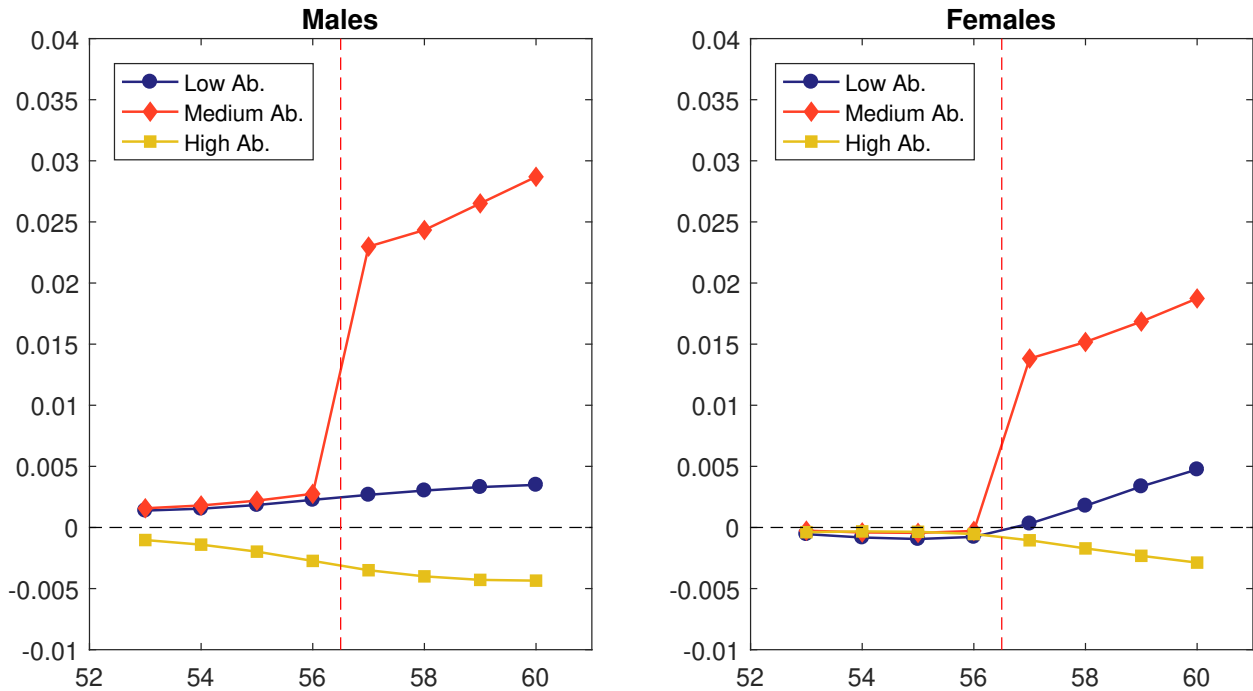


Figure 13: The Effect of the RoSLA on the Proportion Never-Married by Ability Type, Cohort and Gender

the opposite gender gained basic qualifications as a result of the reform, the high ability individuals benefited from the increased supply of academically qualified potential partners, leading them to marry more frequently. While the effect of the reform, due to the marital age gaps, affected earlier cohorts of men than women, within a couple of years, both low and high ability types of both genders were affected in terms of their probability of ever-marrying at a rate of close to half a percentage point, negatively for the low ability type and positively for the high ability type. While economically significant, it is worth noting that these effects are still relatively small compared to the trend in never-married rates over the sample cohorts.

Figure 14 illustrates the effect of the reform on marital sorting on ability. In particular, 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 low ability males, as a consequence of the reform, married low ability women more frequently and medium ability women less frequently. A corresponding effect is highlighted for the low ability women in the bottom left panel. The result thus indicates that the reform unambiguously worsened the marital prospects for the low ability types, reducing their chances of ever marrying and making them more prone to marry among themselves rather than to marry up in terms of ability. The two right panels in contrast show that the high ability individuals, as a consequence of the reform, married less frequently among themselves and married more frequently medium ability, reflecting that the reform increased the academic qualification rate of the medium ability individuals.

This is also reflected in the two middle panels which highlight the effect of the reform on the ability distribution among spouses to medium ability individuals. Here discontinuities naturally

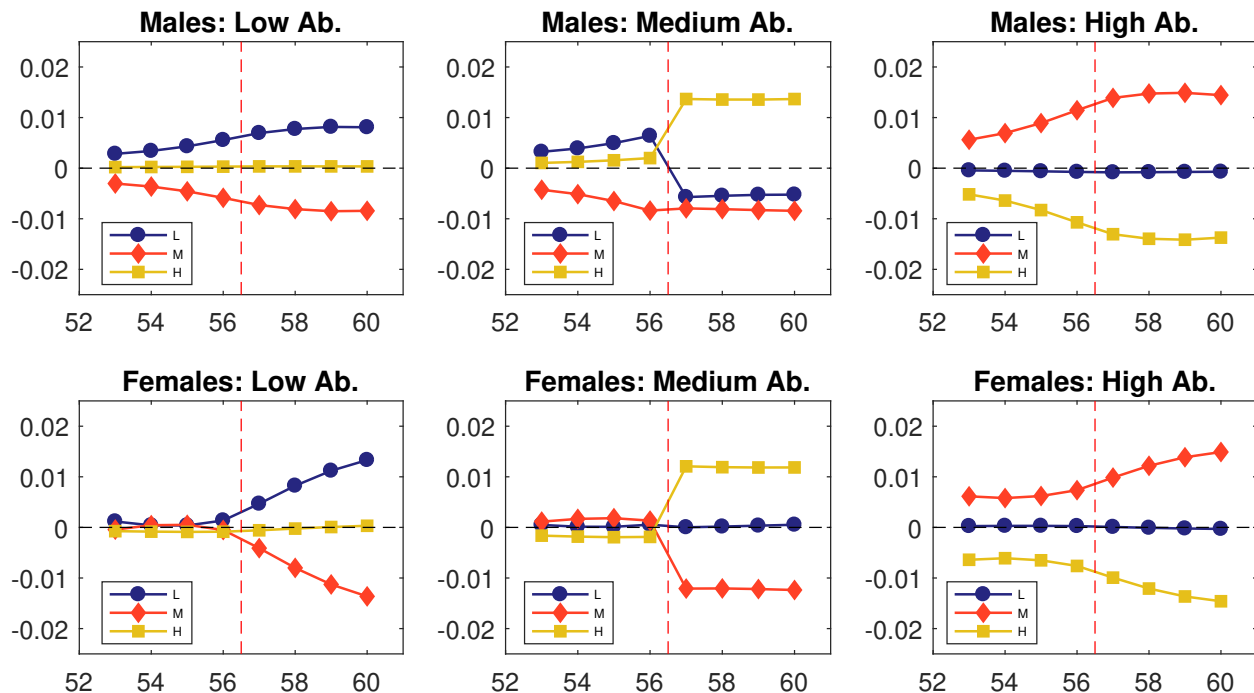


Figure 14: The Effect of the Reform on the Distribution of Spouse Ability

occur as the reform directly affected the academic qualifications held by the medium ability type. For instance, for medium ability males from the pre-reform cohorts—whose qualification rate was thus still low—the reform increased the probability of marrying low ability women, but decreased the probability of marrying medium ability women. This naturally reflects that their wives would often have been from the post-reform cohorts where the medium ability women were now more frequently obtaining qualifications. For medium ability males from the post-reform cohorts—whose own qualification rate was increased—the reform generated a switch away from low ability spouses towards high ability spouses. For the medium ability women, the reform induced more mixing with high ability men and relatively less marriages to medium ability men.

Hence as a broad characterization, based on Figures 13 and 14, the RoSLA (i) increased never-married rates of low and medium ability individuals, while increasing it for high ability individuals, (ii) increased assortative mating among low ability individuals by lowering their rate of marriage to medium ability spouses, and (iii) increased the marital mixing between high and medium ability types. Effectively, equilibrium adjustments in the marriage market in response to the RoSLA induced population shifts left low ability individuals isolated in the marriage market. The equilibrium effects are similar across genders, although different cohorts of men and women are impacted differently due to the estimated surplus contribution of a positive husband-wife age gap.

8. THE MARRIAGE GAP AND THE SPOUSAL QUALIFICATION GAP

The estimated ability-qualification marital surplus matrix revealed that both ability and qualifications are valued in the marriage market, and that it exhibits complementarities both with respect

to ability when both spouses are unqualified, and with respect to qualification when both spouses are of medium ability. In this section we demonstrate how the structural parameters manifest in observable marriage market outcomes. We focus on the difference in marriage probabilities between unqualified individuals and individuals with basic qualifications—the marriage gap—and the difference in the likelihood of being married to a qualified spouse between individuals with no and basic qualifications—the spousal qualification gap.

Our analysis of the marriage gap and spousal qualification gap relates to the literature that seeks to estimate the causal effects of education on marital outcomes, typically using instrumental variables approaches. A problem that has hampered this literature is that most available strong instruments for education are based on institutional changes that operate at the cohort level, either nationally or locally.⁴³ Natural experiments that affect entire cohorts will, however, due to general equilibrium effects, fail to identify the marriage market effects of individual educational attainment.⁴⁴ Only a few contributions provide instrumental variables estimates of the causal effect of education on marital outcomes without relying on cohort-based educational reforms.⁴⁵ These instrumental variables tend to be markedly weaker, making it difficult to draw definite inference from the regressions. [Lefgren and McIntyre \(2006\)](#) study marital outcomes of women in the US using quarter of birth as instrument for educational attainment. Their estimate of the causal effect of an extra year of education on the probability of being married is negative but small and not statistically significant. Similarly, [Anderberg and Zhu \(2014\)](#) consider the effect of holding a basic (CSE/O-level) qualification on marital outcomes of women in the UK, relying on the previously existing Easter-school leaving rule that split each academic cohort in two parts based on the individual’s date of birth. Their point estimate of the effect on the probability of being married is also negative but, likewise, not statistically significant. These papers also address to what extent the observed strong assortative mating on education can be given a causal interpretation, and provide evidence of a positive causal effect of a woman’s educational attainment on the “quality” of her husband as measured by his educational attainment and/or income.

By embedding our analysis in a structural equilibrium model of the marriage market, we are able to exploit the power of a national, cohort-level education reform—the RoSLA—to identify the causal effect of qualification, as opposed to latent ability, on ever-married rates (and hence on marital premia) and on the likelihood of marrying a qualified spouse.

Our analysis of the qualification marriage gap further connects to the more structural literature on “marital premia” which makes use of the fact that, in the Choo-Siow framework differences in marriage rates across types reflect differences in expected marital utility ([Chiappori, Salanié, and](#)

⁴³This includes not only reform such as the RoSLA, but also other instruments based on compulsory schooling laws, or indeed even large-scale local school building programmes as used for instance by [Duflo \(2001\)](#).

⁴⁴Technically, the use of such instrumental variables violates the Stable Unit Treatment Value Assumption (SUTVA) implicit in the potential outcomes model, see e.g. [Rubin \(1980\)](#). The issue was highlighted early on in the marriage market literature by [Lefgren and McIntyre \(2006\)](#) who noted that compulsory schooling laws do not offer identification “because any change in schooling laws will change the schooling of all potential husbands and competing women” ([Lefgren and McIntyre, 2006](#), p. 807).

⁴⁵Contributions that use compulsory school leaving age reforms to investigate the impact of educational attainment on marital and fertility outcomes include [Black, Devereux, and Salvanes \(2008\)](#), [Silles \(2011\)](#), [Cygan-Rehm and Maeder \(2013\)](#), and [Günes \(2016\)](#).

Weiss, 2017).

8.1. The Marriage Gap

Our model allows us to study (i) the difference in the marriage probability for individuals with medium ability a_1 compared to individuals with low ability a_0 , conditional on being unqualified z_0 , and (ii) the difference in the marriage probability for individuals with a basic qualification z_1 compared to individuals with no qualification z_0 , conditional on medium ability a_1 . Specifically, define the “ability marriage gap” for gender g and cohort c as

$$M_A^g(c) \equiv \mu_{0|c,a_0,z_0}^g - \mu_{0|c,a_1,z_0}^g \quad \text{for } g = m, f. \quad (20)$$

As $\mu_{0|c,a_0,z_0}^g$ and $\mu_{0|c,a_1,z_0}^g$ are functions of structural parameters only, $M_A^g(c)$ can be given a causal interpretation as the effect of (medium) ability on the ever-married rate. Similarly, define the “qualification marriage gap” for gender g and cohort c as

$$M_Q^g(c) \equiv \mu_{0|c,a_1,z_0}^g - \mu_{0|c,a_1,z_1}^g \quad \text{for } g = m, f. \quad (21)$$

Similarly, $M_Q^g(c)$, too can be given a causal interpretation as the effect of (basic) qualification attainment on the ever-married rate.

Of course, the causal 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 is given by the “total marriage gap”,

$$M_T^g(c) = M_A^g(c) + M_Q^g(c) = \mu_{0|c,a_0,z_0}^g - \mu_{0|c,a_1,z_1}^g \quad \text{for } g = m, f. \quad (22)$$

Note that our model only identifies $M_A^g(c)$ and $M_Q^g(c)$ for the pre-RoS LA cohorts $c \in \mathcal{C}_0$. In contrast $M_T^g(c)$ is identified for all sample cohorts, but directly observable only for the post-RoS LA cohorts.

Indeed, as a point of comparison, consider the observed gap in marriage probabilities between individuals with a basic qualification and individuals who are unqualified (of gender g and cohort c), denoted $\tilde{M}_Q^g(c)$. The structural model implies that

$$\tilde{M}_Q^g(c) = M_T^g(c) - \Pr^g(a_1|c, z_0)M_A^g(c) = M_Q^g(c) + [1 - \Pr^g(a_1|c, z_0)]M_A^g(c) \quad \text{for } g = m, f, \quad (23)$$

where $\Pr^g(a_1|c, z_0) = (1 - \gamma^g)\Pr^g(a_1|c)/[\Pr^g(a_0|c) + (1 - \gamma^g)\Pr^g(a_1|c)]$. Hence, prior to the RoS LA, where $\gamma^g < 1$, the observed marriage gap $\tilde{M}_Q^g(c)$ confounds the causal effects $M_A^g(c)$ and $M_Q^g(c)$. Post-RoS LA, $\gamma^g = 1$, and $\tilde{M}_Q^g(c) = M_T^g(c)$, as basic qualified and qualified individuals always also differ in ability.

The above defined marriage gaps directly relate to the notion of “marriage premia” defined as an expected marital utility difference. As noted by Chiappori, Salanié, and Weiss (2017), the negative of the log never-married rate, $-\log(\mu_{0|x}^g)$, among individuals of gender g and type $x \in \mathcal{X}$ is a measure of their average marital utility. This allows us to reinterpret the (proportional) marriage

gaps correspondingly as marriage premia.⁴⁶ One of the central themes in [Chiappori, Salanié, and Weiss \(2017\)](#) is that the “marriage college premium” due to an increasing college wage premium and technological innovations that have reduced the time needed for domestic production, can be expected to have increased over time, particularly for women. This would be reflected in the marriage rate of college-educated women increasing relative to lower qualified women, and the authors use US data on cohorts born over the period 1942-1973 to verify their model’s predictions. As can be seen from [Figure 4](#), even though we only have eight sample cohorts, the same feature appears in our data: over the sample cohorts, the ever-married rate of advanced qualified women increased in relative term compared both to unqualified and basic qualified women.

The top panels of [Figure 15](#) illustrate $M_A^g(c)$, $M_Q^g(c)$, and $M_T^g(c)$ for men and women, respectively. The figure highlights how the total marriage gap between medium ability basic qualified individuals and low ability unqualified individuals 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. In line with [Lefgren and McIntyre \(2006\)](#) and [Anderberg and Zhu \(2014\)](#), the top-right panel shows that, for women, the estimated effect of holding a basic qualification on the probability of ever-marrying is effectively zero. The top-left panel shows that the estimated qualification effect for men is negative, but relatively small. In contrast, the estimated effects of ability are positive for both men and women. As a result, the total gap in marriage probability $M_T^g(c)$ is mainly driven by the positive effect of ability. Reinterpreted correspondingly as marriage premia, the model suggests positive ability marriage premia for both men and women, but zero and negative (basic) qualification marriage premia for women and men respectively.

To avoid cluttering in [Figure 15](#), we do not plot the observed marriage gap $\tilde{M}_Q^g(c)$. It is clear from [\(23\)](#), however, that $\tilde{M}_Q^g(c)$ is an upward biased estimate of the causal effect on qualification attainment on the ever-married rate, since $M_A^g(c) > 0$. Indeed, pre-RoS LA, a plot of $\tilde{M}_Q^g(c)$ would lie between the plots of $M_Q^g(c)$ and $M_T^g(c)$ in [Figure 15](#). Post-RoS LA, $\tilde{M}_Q^g(c)$ and $M_T^g(c)$ coincides. As a case in point, for pre-RoS LA women in [Figure 15](#), the causal effect of qualification on the probability of ever marrying is zero; yet, the observed marriage gap $\tilde{M}_Q^g(c)$ is positive.

The shift in the never-married rates for the unqualified women is particularly salient because, as noted earlier, it coincides with the moment where unqualified women overtake women with advanced qualifications as the group that married the least in their lives. From [Figure 4](#), in the 1956-cohort women with advanced qualifications had a never-married rate that was about 1.4 percentage points higher than that of the unqualified women. One academic cohort later, the unqualified women had a never-married rate that was instead 0.4 higher than that of the women holding advanced qualifications. Back-of-the-envelope calculations show that the observed 1.8 percentage

⁴⁶Formally, for male type $x_i \in \mathcal{X}$, define $\bar{u}(x_i) \equiv \mathbb{E} [\max_{x_j \in \mathcal{X}_+} \{U(x_i, x_j) + \varepsilon_i(x_j)\}] = -\log(\mu_{0|x_i}^m)$, while for female type $x_j \in \mathcal{X}$, $\bar{v}(x_j) \equiv \mathbb{E} [\max_{x_i \in \mathcal{X}_+} \{V(x_i, x_j) + \varepsilon_j(x_i)\}] = -\log(\mu_{0|x_j}^f)$. Then define the ability and the qualification marriage premium for males of cohort c as the relevant gaps in expected marital utility, $P_A^m(c) \equiv \bar{u}(c, a_1, z_0) - \bar{u}(c, a_0, z_0)$ and $P_Q^m(c) \equiv \bar{u}(c, a_1, z_1) - \bar{u}(c, a_1, z_0)$ respectively. Premia for women, $P_A^f(c)$ and $P_Q^f(c)$, can be defined correspondingly. Note then that, using the standard log approximation $\log(1+x) \approx x$, it follows that $P_A^g(c) \approx M_A^g(c) / \mu_{0|c, a_1, z_0}^g$ and $P_Q^g(c) \approx M_Q^g(c) / \mu_{0|c, a_1, z_1}^g$.

point swing is fully accounted for by the compositional change among the unqualified women. Indeed, the model-implied never-married rate of unqualified individuals of gender g can be written as the weighted average $\mu_{0|c,z_0}^g = \Pr^g(a_0|c, z_0) \mu_{0|c,a_0,z_0}^g + \Pr^g(a_1|c, z_0) \mu_{0|c,a_1,z_0}^g$. The direct effect of the elimination of the medium ability types from the unqualified group caused by the RoSLA would shift the never-married rate to $\mu_{0|c,a_0,z_0}^g$, a shift of size $\Pr^g(a_1|c, z_0) M_A^g(c)$, see (20). Using the values for women in 1956, the last pre-RoSLA cohort, we have that $\Pr^f(a_1|1956, z_0) = 0.33$ and $M_A^f(1956) = 0.049$, implying that the RoSLA-induced compositional change comprises a 1.7 percentage point upward shift in the never-married rate of the unqualified women, sufficient for the unqualified women to overtake women with advanced qualifications in terms of the never-married rate.

8.2. The Spousal Qualification Gap

We define $Q_A^g(c)$ 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 $g = m, f$. Similarly, we define $Q_Q^g(c)$ as the gap in spouse qualification rate between those with qualification z_1 and those unqualified z_0 , conditional on medium ability a_1 and married, for gender $g = m, f$. Finally, define $Q_T^g(c)$ as the total gap in spouse qualification rate between those of ability-qualification type (a_1, z_1) and those of type (a_0, z_0) . As in the case of the marriage probabilities, $Q_A^g(c)$ and $Q_Q^g(c)$ are identified only for pre-RoSLA cohorts, whereas $Q_T^g(c)$, which confounds the effects of ability and qualification, is directly observable and identified for all sample cohorts.

A note of caution is warranted before we turn to the empirical analysis of the spousal qualification gap. In our analysis of the marriage gap in the previous subsection, we stressed that our structural model admitted a causal interpretation of the identified effects. Because the spousal qualification gap is defined only conditional on marriage, such interpretation is not possible. Nonetheless we believe the decomposition of $Q_T^g(c)$ into the quantities $Q_A^g(c)$ and $Q_Q^g(c)$ is novel and informative about the relative importance of ability and qualifications in the marriage market.

The lower panels of Figure 15 illustrate how our estimated model predicts that the probability of being married to an academically qualified spouse increases with own ability and with an own (basic) qualification respectively. The figure suggests that, for both men and women, ability and an own qualification have similar sized 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 well with the observable differences highlighted in Figure 2.

9. CONCLUSIONS

One of the most well-known stylized facts about marital choices is that there is strong assortative mating on education and qualifications. However, as several decades of research in labor and education economics has convincingly demonstrated, educational and qualification attainment is associated with latent individual characteristics, most notably academic ability. Evidence on the

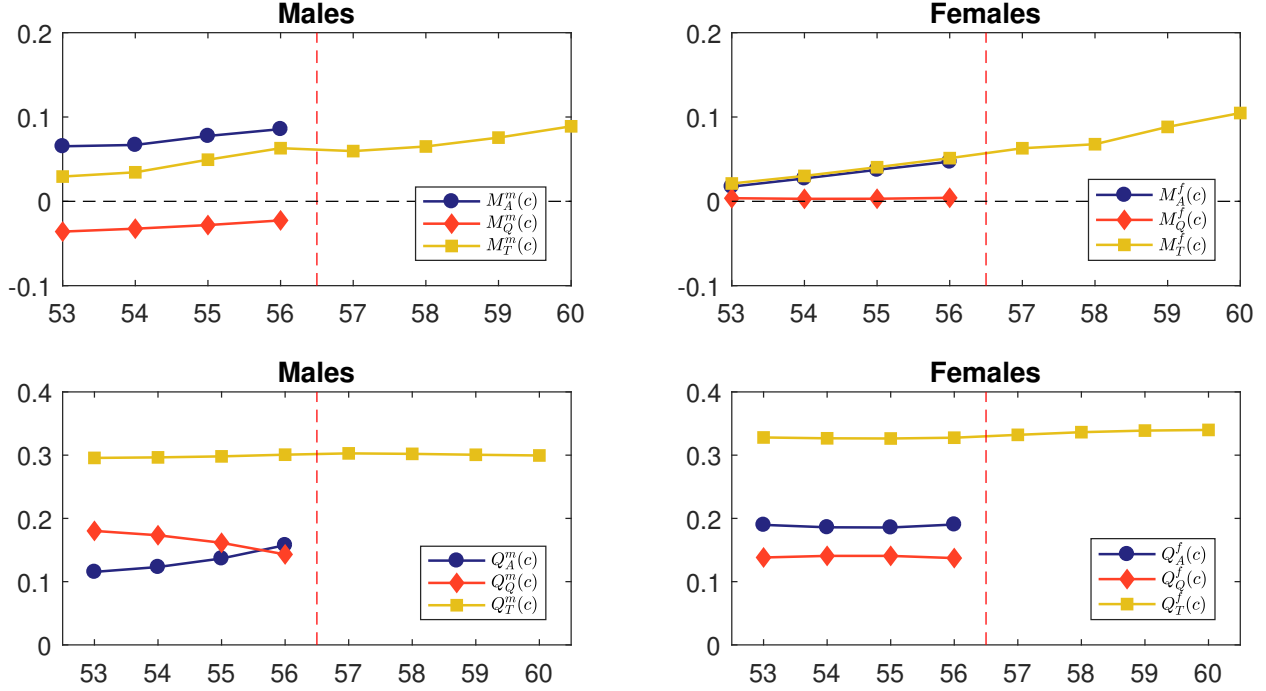


Figure 15: The Model-Predicted Effect of Own Ability and Basic Qualification on the Probability of (i) Ever Marrying (Top Row) by Gender and Cohort and (ii) the Spouse Holding at Least a Basic Academic Qualification by Gender and Cohort (Bottom Row)

extent to which observed assortative mating reflect sorting on latent ability or qualification is scant. A key contributing factor to this state of affairs is that typically used instruments for qualification attainment operate at a cohort-level, either nationally or locally. This issue applies to instruments based on school-leaving laws, large-scale school-building programmes etc. Such instrumental variables are not valid in the context of evaluating marriage market outcomes as they conflate individual and general equilibrium effects as cohorts interact in the marriage market.

Motivated by the marriage market responses to a large UK education reform, the RoSLA, which was implemented nationally in 1972 and therefore operated at the (academic) cohort level, we have tackled this issue by modelling a [Choo and Siow \(2006\)](#) marriage market equilibrium, where individuals match not only on observable age (i.e. cohort) and qualifications, but also on latent ability. Interpreting the RoSLA as aligning the mapping from ability to qualification, an assumption supported by empirical evidence on selective RoSLA compliance, we showed that the Choo-Siow model with a latent systematic type is identified from the large RoSLA-induced shifts in the qualification distributions, and estimated the model’s structural parameters.

We verified that accounting for latent ability is central to fitting the observable marriage market responses to the RoSLA. In particular, we provided empirical evidence that the never-married rate of unqualified individuals increased at the RoSLA threshold despite their reduced population supplies. A benchmark Choo-Siow model, where marriage market types are delineated by the intersection of age and qualifications only, fails to reproduce increasing never-married rates at the RoSLA. Our extended model offers a simple explanation for the observed increases: the educational response to

the RoSLA was selective, with those responding to the reform by gaining an academic qualification having higher ability than those not responding. As our estimated model indicates a significant role for ability as a determinant of marital surplus, this compositional shift raised the never-married rate among those still unqualified after the RoSLA reform.

By estimating a full marriage market equilibrium we are able to determine not only how individual marital outcomes are related to ability and qualifications, but also how large the marriage market general equilibrium effects of the reform are. The effects of the reform were by no means confined to cohorts and ability types directly affected by the reform in terms of their qualification attainment. We decompose observed gaps in marriage probabilities and spousal qualification probabilities between unqualified individuals and those with a basic qualification into causal effects stemming from ability differences and qualification differences. The marriage gap is driven almost entirely by ability, whereas ability and qualifications contribute equally to the observed spousal qualification gap.

REFERENCES

- 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.
- BECKER, G. (1973): “A Theory of Marriage: Part I,” *Journal of Political Economy*, 81(4), 813–846.
- BLACK, S. E., P. J. DEVEREUX, AND K. G. SALVANES (2008): “Staying in the Classroom and Out of the Maternity Ward? The Effect of Compulsory Schooling Laws on Teenage Births,” *Economic Journal*, 118(530), 1025–1054.
- BLANE, D., G. D. SMITH, AND M. BARTLEY (1990): “Social Class Differences in Years of Potential Life Lost: Size, Trends, and Principal Causes,” *British Medical Journal*, 301, 429–432.
- 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.
- 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., S. OREFFICE, AND C. QUINTANA-DOMEQUE (2012): “Fatter Attraction: Anthropometric and Socioeconomic 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.
- CYGAN-REHM, K., AND M. MAEDER (2013): “The Effect of Education on Fertility: Evidence from a Compulsory Schooling Reform,” *Labour Economics*, 25(C), 35–48.
- 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.

- 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.
- DUFLO, E. (2001): “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 91(4), 795–813.
- DUPUY, A., AND A. GALICHON (2014): “Personality Traits and the Marriage Market,” *Journal of Political Economy*, 122(6), 1271–1319.
- FOX, J., C. YANG, AND D. H. HSU (2018): “Unobserved Heterogeneity in Matching Games,” forthcoming in the *Journal of Political Economy*.
- GALICHON, A., AND B. SALANIÉ (2015): “Cupid’s Invisible Hand: Social Surplus and Identification in Matching Models,” New York University.
- 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,” UC Santa Barbara.
- GRENET, J. (2013): “Is it Enough to Increase Compulsory Schooling to Raise Earnings? Evidence from French and British Compulsory Schooling Laws,” *Scandinavian Journal of Economics*, 115(1), 176–210.
- GÜNEŞ, P. M. (2016): “The Impact of Female Education on Teenage Fertility: Evidence from Turkey,” *The B.E. Journal of Economic Analysis & Policy*, 16(1).
- HAHN, J., P. TODD, AND W. V. DER KLAAUW (2001): “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design,” *Econometrica*, 69(1), 201–209.
- IMBENS, G., AND T. LEMIEUX (2008): “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 142, 615–635.
- KNOWLES, J. A., AND G. VANDENBROUCKE (2018): “Fertility Shocks and Equilibrium Marriage-Rate Dynamics,” Simon Fraser University.
- KOLESÁR, M., AND C. ROTHE (2018): “Inference in Regression Discontinuity Designs with a Discrete Running Variable,” *American Economic Review*, 108(8), 2277–2304.
- 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.

- MANSOUR, H., AND T. MCKINNISH (2014): “Who Marries Differently Aged Spouses? Ability, Education, Occupation, Earnings, and Appearance,” *Review of Economics and Statistics*, 96(3), 577–580.
- MARE, R. D. (1991): “Five Decades of Educational Assortative Mating,” *American Sociological Review*, 56(1), 15–32.
- MOURIFIÉ, I., AND A. SIOW (2017): “The Cobb Douglas Marriage Matching Function: Marriage Matching with Peer and Scale Effects,” University of Toronto.
- OREFFICE, S., AND C. QUINTANA-DOMEQUE (2010): “Anthropometry and Socioeconomics among Couples: Evidence in the United States,” *Economics and Human Biology*, 8(3), 373–384.
- RUBIN, D. B. (1974): “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies,” *Journal of Educational Psychology*, 66(5), 688–701.
- (1980): “Randomization Analysis of Experimental Data: The Fisher Randomization Test: Comment,” *Journal of the American Statistical Association*, 75(371), 591–593.
- 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.
- WADSWORTH, M. E. J., R. J. HARDY, A. A. PAUL, S. F. MARSHALL, AND T. J. COLE (2002): “Leg and Trunk Length at 43 Year in Relation to Childhood Health: Evidence from the 1946 National Birth Cohort,” *International Journal of Epidemiology*, 31, 383–390.

A. Evidence of a Selective RoSLA Response

In this appendix we present suggestive evidence that the RoSLA widened the gap in social background between individuals with and without academic qualifications. To do so we use data from the Health Survey for England (HSE), a representative survey of individuals living in England, with approximately 10,000 respondents in each wave. The HSE has been running annually since 1991, however due to the availability of detailed qualifications information, we use data from 1998 to 2014. Respondents complete a core questionnaire containing demographic, lifestyle and health-related questions which is further supplemented with physical measurements, including height, taken by a health-care professional. Each wave contains a supplemental module, focusing on a particular condition or disease.

As outlined in the main text, we focus on two markers of disadvantaged social background: Father’s mortality and individual height. Individual height is available in every HSE wave; in contrast, parental mortality is available only in 1998, 2003, and 2006 as these were years when then HSE’s Cardio-Vascular module was implemented. In order to obtain larger sample sizes, we include all individuals born in the academic cohorts 1947 to 1966 in our analysis of the HSE data.

Our aim is to highlight selective responses to the RoSLA. To that end, we revisit the RDD in equation (2) using the HSE data, and consider three different outcome variables. The results are provided in Table A1.⁴⁷ In column (i), the outcome variable is a dummy for holding some academic qualification—basic or advanced—and where the regression includes a gender dummy. The coefficient on being RoSLA-treated indicates an 8.3 percentage point increase the rate of holding some qualification. This is slightly below the response estimated in the LFS data, but broadly speaking consistent.⁴⁸ The estimates reported in Table A1, columns (ii) and (iv) refer to specifications where the outcome variables are own height (measured in centimeters) and an indicator for the respondent’s father having passed away, respectively. The estimates in column (ii) shows that the average height was 175 cm among men and 162 cm among women, and that the RoSLA, as expected, did not impact average height. From column (iv), about half of the respondents reported their fathers to have passed away at the time of the interview, with only a minor difference between male and female respondents, and again with no effect of the RoSLA-treatment.

Columns (iii) and (v) present estimates from an extended version of the RDD equation (2) that includes an indicator for holding no academic qualification, $\mathbb{1}(z_i = z_0)$, and the same indicator interacted with the RoSLA exposure dummy, $\mathbb{1}(w_i \geq 0)$,

$$y_i = \alpha_0 + \alpha_1 w_i + \alpha_2 w_i \mathbb{1}(w_i \geq 0) + \varphi(0) \mathbb{1}(w_i \geq 0) + \rho_1 \mathbb{1}(z_i = z_0) + \rho_2 \mathbb{1}(z_i = z_0) \mathbb{1}(w_i \geq 0) + \epsilon_i \quad (\text{A1})$$

where y_i is one of the various outcome variables described above, and ϵ_i is an error term. Estimates of ρ_1 and ρ_2 in (A1) are of particular interest. The parameter ρ_1 measures the pre-RoSLA average difference in the outcome variable, i.e. height and paternal mortality, between unqualified and qualified individuals. The

⁴⁷A more detailed RDD analysis demonstrating the robustness of the results presented here are available on request from the authors.

⁴⁸The HSE was recently used by [Clark and Royer \(2013\)](#) to estimate, in a RDD, the effect of education on health outcomes using the same educational reform that we focus on here, along with the earlier 1947 reform which raised the school-leaving age from 14 to 15. Their main finding was that there was little or no effect of these reforms on health outcomes.

parameter ρ_2 measures how this difference was impacted by the RoSLA. A significant parameter estimate of ρ_2 indicates a changing average gap and hence a selective RoSLA response.

Table A1: Regression Discontinuity Estimate of the Impact of the RoSLA on the Academic Qualification Rate and of the Relationship Between Holding an Academic Qualification and Own Height and Father’s Mortality Based on Data from the Health Survey for England

	Ac. Qual.	Own Height		Father Dead	
	(i)	(ii)	(iii)	(iv)	(v)
Constant	0.658*** (0.010)	175.32*** (0.08)	176.01*** (0.09)	0.494*** (0.013)	0.468*** (0.013)
Gender (female)	0.023*** (0.004)	-13.27*** (0.06)	-13.32*** (0.06)	-0.019* (0.008)	-0.017* (0.008)
RoSLA (≥ 57)	0.083*** (0.012)	0.05 (0.11)	-0.02 (0.11)	-0.007 (0.015)	-0.011 (0.0157)
No Qual.			-2.12*** (0.08)		0.082*** (0.010)
No Qual * RoSLA			-0.40** (0.13)		0.037* (0.015)
Obs.	58,456	53,947	53,947	15,661	15,661

Notes: The sample pools all individuals from academic cohorts 1947-1966 with information about academic qualifications observed in the Health Survey for England, 1998-2014. The running variable splits each academic year into three periods (Sept-Dec, Jan-April, May-Aug) and is centred on the first RoSLA treated group (born Sept-Dec, 1957). Each regression includes the running variable and its interaction with the RoSLA indicator. Information on father’s mortality is available only in the 1998, 2003 and 2006 surveys. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The parameter estimate of ρ_1 reported in Table A1, column (iii) therefore suggests that, among individuals born prior to the RoSLA, the difference in height between qualified and unqualified was a little over 2 centimeters on average. The estimate of ρ_2 implies that, at the RoSLA, this gap increased by 0.4 centimeters, a 20 percent increase over the pre-reform gap. Similarly, the estimate of ρ_1 reported in Table A1, column (v) tells us that, among individuals born prior to the RoSLA, the proportion reporting that their fathers are dead was 8 percentage points higher among unqualified than among qualified individuals. From the estimate of ρ_2 , we conclude that, among individuals born after the reform threshold, this gap between unqualified and qualified increased by 3.7 percentage points, an increase of more than 40 percent over the pre-reform gap. Hence we find that for two standard markers of social background, the gap between unqualified and qualified individuals increased at the RoSLA threshold. In particular, those who were unqualified after the RoSLA were found to be shorter on average and more likely to have fathers who were dead at the time of interview. These results are suggestive of a selective response to the RoSLA whereby those who, even after the reform, did not obtain any qualification, were particularly negatively selected in terms of social background.

B. Proof of Proposition 1

B.1. Population Recall that $\mathcal{X} \subset \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$. With just two cohorts, two ability types and two qualification types, and given Assumption 1, only five full types exist

$$x \in \mathcal{X} \equiv \{(c_0, a_0, z_0), (c_0, a_1, z_0), (c_0, a_1, z_1), (c_1, a_0, z_0), (c_1, a_1, z_1)\}, \quad (\text{B1})$$

for each gender $g = m, f$.

As researchers we observe cohorts and qualifications, but not abilities. This means that we cannot directly distinguish between the first two types in \mathcal{X} while the last three types can be immediately distinguished. Hence we can partition \mathcal{X} into \mathcal{X}^I and \mathcal{X}^D where the former contains the two indistinguishable types in \mathcal{X} and the latter the three distinguishable types. As c and z are observed, so is the qualification distribution by cohort, whereby γ^g is trivially identified for each gender g . This in turn implies that the population measure $h^g(x)$ is identified for each gender. Hence for the remainder of the analysis we will treat the population distribution as known.

B.2. Surplus and Preferences As, in this case, $|\mathcal{X}| = 5$, the systematic surplus terms $\Sigma(x_i, x_j)$ can be represented as a 5×5 matrix with male full types on rows and female full types in columns. Let \mathcal{S} denote the set of all possible surplus matrices. We say that a surplus matrix Σ is permissible under Assumption 2 if it has the additive form in (8), and we let $\mathcal{S}^P \subset \mathcal{S}$ denote the space of all permissible surplus matrices. In the below it will be convenient to have a short-hand for the cohort-profile preferences. Hence let $\lambda_{ij} = \lambda(c_i, c_j)$. Note also that one element of either $\zeta(\cdot)$ or $\lambda(\cdot)$ has to be normalized. We set $\lambda_{00} = \lambda(c_0, c_0) = 0$.

As there are, in this case, three ability-qualification types $\{(a_0, z_0), (a_1, z_0), (a_1, z_1)\}$, $\zeta(\cdot)$ can be represented as a 3×3 matrix male ability-qualification types on rows and female types on columns. We thus have that any element $\Sigma \in \mathcal{S}^P$ is characterized by twelve parameters: nine ζ -terms relating to possible ability-qualification profiles and three (non-normalized) λ -terms.

As noted in Section 5, (i) in equilibrium the systematic surplus $\Sigma(x_i, x_j)$ is split into a part that accrues to the male, $U(x_i, x_j)$, and a part that accrues to the female, $V(x_i, x_j)$, and (ii) with the assumed extreme value distributed individual utility components the choice frequencies $\mu_{x_j|x_i}^m$ and $\mu_{x_i|x_j}^f$ (with “0” indicating singlehood) satisfy (5). In general $U(x_i, x_j)$ and $V(x_i, x_j)$ do not inherit the additive separable form from $\Sigma(x_i, x_j)$.

B.3. Permissible Matching in Full Types The equilibrium marriage/singles frequencies for gender g can be collected in a 5×6 marriage matrix denoted $\boldsymbol{\mu}^g$, with own type $x \in \mathcal{X}$ on rows and marital choice from the set $\mathcal{X}_+ = \mathcal{X} \cup \{0\}$ in columns. Letting $\boldsymbol{\mu} = \{\boldsymbol{\mu}^m, \boldsymbol{\mu}^f\}$, $\boldsymbol{\mu}$ is said to be feasible given the population distribution if every element is non-negative and satisfy adding-up (7) and market balance (6). Let \mathcal{M}^F denote the set of feasible $\boldsymbol{\mu}$ given the population type distribution.

The separability property in (8) imposes a set of restrictions on the matching in full types. In particular, using the logit form (5), the balancedness condition (6) (and the assumption of equal-sized cohorts) it is easy to show that the following “projection relations” hold,

$$\frac{\mu_{c_j, a_j, z_j | c_0, a_0, z_0}^m}{\mu_{c_j, a_j, z_j | c_1, a_0, z_0}^m} = \sqrt{\frac{\mu_{0|c_0, a_0, z_0}^m}{\mu_{0|c_1, a_0, z_0}^m}} \exp\left(\frac{\lambda_{01} - \lambda_{11}}{2} \mathbb{1}(c_j = c_1) - \frac{\lambda_{10}}{2} \mathbb{1}(c_j = c_0)\right), \quad (\text{B2})$$

$$\frac{\mu_{c_i, a_i, z_i | c_0, a_0, z_0}^f}{\mu_{c_i, a_i, z_i | c_1, a_0, z_0}^f} = \sqrt{\frac{\mu_{c_0, a_0, z_0}^f}{\mu_{c_1, a_0, z_0}^f}} \exp\left(\frac{\lambda_{10} - \lambda_{11}}{2} \mathbb{1}(c_i = c_1) - \frac{\lambda_{01}}{2} \mathbb{1}(c_i = c_0)\right), \quad (\text{B3})$$

and

$$\frac{\mu_{c_0, a_0, z_0 | c_i, a_i, z_i}^m}{\mu_{c_1, a_0, z_0 | c_i, a_i, z_i}^m} = \sqrt{\frac{\mu_{c_0, a_0, z_0}^f}{\mu_{c_1, a_0, z_0}^f}} \exp\left(\frac{\lambda_{10} - \lambda_{11}}{2} \mathbb{1}(c_i = c_1) - \frac{\lambda_{01}}{2} \mathbb{1}(c_i = c_0)\right), \quad (\text{B4})$$

$$\frac{\mu_{c_0, a_0, z_0 | c_j, a_j, z_j}^f}{\mu_{c_1, a_0, z_0 | c_j, a_j, z_j}^f} = \sqrt{\frac{\mu_{c_0, a_0, z_0}^m}{\mu_{c_1, a_0, z_0}^m}} \exp\left(\frac{\lambda_{01} - \lambda_{11}}{2} \mathbb{1}(c_j = c_1) - \frac{\lambda_{10}}{2} \mathbb{1}(c_j = c_0)\right). \quad (\text{B5})$$

Hence we say that a $\boldsymbol{\mu} \in \mathcal{M}^F$ that satisfies (B2) - (B5) is *permissible* under Assumption 2. The subset of feasible $\boldsymbol{\mu}$ that are also permissible is denoted \mathcal{M}^P .

B.4. Permissible Matching in Observable Types We next define the matching in observable types. The set of observable types for each gender is $\tilde{\mathcal{X}} = \mathcal{C} \times \mathcal{Z}$ with a typical element \tilde{x} . Similar to the full types, we let $\tilde{\mu}_{\tilde{x}_j | \tilde{x}_i}^m$ denote the proportion of observable male type \tilde{x}_i that make the observable marital choice $\tilde{x}_j \in \tilde{\mathcal{X}}_+ \equiv \tilde{\mathcal{X}} \cup \{0\}$. Female observable choice frequencies are given a corresponding notation. The equilibrium marriage/singles frequencies in terms of observables for gender g can be collected in a 4×5 matrix $\tilde{\boldsymbol{\mu}}^g$ with own observable type on rows and observable marital choices on columns. The same feasibility conditions—non-negativity, adding up within each row, and balancedness—also apply to the observable matching matrices. Hence let $\tilde{\mathcal{M}}^F$ denote the set of feasible $\tilde{\boldsymbol{\mu}} = \{\tilde{\boldsymbol{\mu}}^m, \tilde{\boldsymbol{\mu}}^f\}$ given the population distribution.

Assumption 2 imposes restrictions also on $\tilde{\boldsymbol{\mu}}$. In particular, using (5) and (6) it can be shown that, for each gender $g = m, f$, the following four ratios will hold,

$$\frac{\tilde{\mu}_{c_0, z_1 | c_0, z_0}^g}{\tilde{\mu}_{c_1, z_1 | c_0, z_0}^g} = \frac{\tilde{\mu}_{c_0, z_1 | c_0, z_1}^g}{\tilde{\mu}_{c_1, z_1 | c_0, z_1}^g}, \text{ and } \frac{\tilde{\mu}_{c_0, z_1 | c_1, z_0}^g}{\tilde{\mu}_{c_1, z_1 | c_1, z_0}^g} = \frac{\tilde{\mu}_{c_0, z_1 | c_1, z_1}^g}{\tilde{\mu}_{c_1, z_1 | c_1, z_1}^g}, \quad (\text{B6})$$

$$\frac{\tilde{\mu}_{c_1, z_0 | c_0, z_1}^g}{\tilde{\mu}_{c_1, z_1 | c_0, z_1}^g} = \frac{\tilde{\mu}_{c_1, z_0 | c_1, z_1}^g}{\tilde{\mu}_{c_1, z_1 | c_1, z_1}^g}, \text{ and } \frac{\tilde{\mu}_{c_0, z_0 | c_0, z_1}^g}{\tilde{\mu}_{c_0, z_1 | c_0, z_1}^g} = \frac{\tilde{\mu}_{c_0, z_0 | c_1, z_1}^g}{\tilde{\mu}_{c_0, z_1 | c_1, z_1}^g}. \quad (\text{B7})$$

Hence we say that a $\tilde{\boldsymbol{\mu}} \in \tilde{\mathcal{M}}^F$ that satisfies (B6) through (B7) is *permissible* under Assumption 2, and we use $\tilde{\mathcal{M}}^P$ to denote the subset of feasible $\tilde{\boldsymbol{\mu}}$ that are also permissible. It is worth noting that the restrictions in (B6) and (B7) reduce the total number of free moments in $\tilde{\boldsymbol{\mu}}$ to twelve, the same as the number of unknown surplus parameters.

B.5. Identification To prove identification we assume that we have access to data from a representative sample that is large enough that $\tilde{\boldsymbol{\mu}}$ can be treated as known. We show that under our assumptions the equilibrium matching in full types $\boldsymbol{\mu} \in \mathcal{M}^P$ is identified from the matching in observable types $\tilde{\boldsymbol{\mu}} \in \tilde{\mathcal{M}}^P$. Identification of the systematic marital surplus parameters then follows from Choo and Siow (2006).

We proceed in three steps. First we show that the cohort-profile preferences λ_{ij} are identified from the cohort-profiles of married couples where both spouses hold a qualification. Second, we show that, using the separability assumption, the equilibrium single-rates of all full types are recovered by projecting the

marriage choices of the unqualified individuals from the post-reform cohort c_1 . Third, we show that all marriage frequencies in full types can be recovered.

For the first step we use that, by Assumption 1, all qualified individuals z_1 have ability a_1 . Using equations (5) and (8) it then follows that

$$\lambda_{ij} = \log \left(\frac{\tilde{\mu}_{c_j, z_1 | c_i, z_1}^m}{\tilde{\mu}_{0 | c_i, z_1}^m} \right) + \log \left(\frac{\tilde{\mu}_{c_i, z_1 | c_j, z_1}^f}{\tilde{\mu}_{0 | c_j, z_1}^f} \right) - \log \left(\frac{\tilde{\mu}_{c_0, z_1 | c_0, z_1}^m}{\tilde{\mu}_{0 | c_0, z_1}^m} \right) - \log \left(\frac{\tilde{\mu}_{c_0, z_1 | c_0, z_1}^f}{\tilde{\mu}_{0 | c_0, z_1}^f} \right), \quad (\text{B8})$$

for $i, j = 0, 1$. The three non-normalized cohort-profile surplus parameters λ_{ij} are thus initially identified and will be treated as known for the remainder of the analysis.

For the second step we first note that the single rates for all distinguishable male types $x_i \in \mathcal{X}^D$ have directly observable counterparts. Hence we focus on showing that the single rates for the indistinguishable types $x_i \in \mathcal{X}^I$ can also be recovered. To do so we use that (B2) implies that the following two equations hold,

$$\mu_{c_0 | c_0, a_0, z_0}^m = \mu_{c_0 | c_1, a_0, z_0}^m \frac{\sqrt{\mu_{0 | c_0, a_0, z_0}^m}}{\sqrt{\mu_{0 | c_1, a_0, z_0}^m}} \exp \left(-\frac{\lambda_{10}}{2} \right), \quad (\text{B9})$$

and

$$\mu_{c_1 | c_0, a_0, z_0}^m = \mu_{c_1 | c_1, a_0, z_0}^m \frac{\sqrt{\mu_{0 | c_0, a_0, z_0}^m}}{\sqrt{\mu_{0 | c_1, a_0, z_0}^m}} \exp \left(\frac{\lambda_{01} - \lambda_{11}}{2} \right), \quad (\text{B10})$$

where we used the shorthand $\mu_{c_k | x_i}^m \equiv \sum_{\{x_j \in \mathcal{X} | c_j = c_k\}} \mu_{x_j | x_i}^m$ for $k = 0, 1$ to denote the probability that a male of full type x_i marries a female from cohort c_k . Adding up (7) implies that

$$\mu_{c_0 | c_0, a_0, z_0}^m + \mu_{c_1 | c_0, a_0, z_0}^m + \mu_{0 | c_0, a_0, z_0}^m - 1 = 0. \quad (\text{B11})$$

Substituting in (B11) using (B9) and (B10) and replacing the choices made by male type (c_1, a_0, z_0) with their directly observed counterparts yields that,

$$\sqrt{\mu_{0 | c_0, a_0, z_0}^m} \left[\frac{\tilde{\mu}_{c_0 | c_1, z_0}^m}{\sqrt{\tilde{\mu}_{0 | c_1, z_0}^m}} \exp \left(-\frac{\lambda_{10}}{2} \right) + \frac{\tilde{\mu}_{c_1 | c_1, z_0}^m}{\sqrt{\tilde{\mu}_{0 | c_1, z_0}^m}} \exp \left(\frac{\lambda_{01} - \lambda_{11}}{2} \right) \right] + \mu_{0 | c_0, a_0, z_0}^m - 1 = 0, \quad (\text{B12})$$

where $\tilde{\mu}_{c_k | \tilde{x}_i}^m \equiv \sum_{\{\tilde{x}_j \in \tilde{\mathcal{X}} | c_j = c_k\}} \tilde{\mu}_{\tilde{x}_j | \tilde{x}_i}^m$ is the probability that observable male type \tilde{x}_i marries a spouse from cohort c_k , $k = 0, 1$. Note that (B12) defines a quadratic equation in the square root of $\mu_{0 | c_0, a_0, z_0}^m$, with the coefficient defined by the large square bracket term being positive. As the left hand side of (B12) is strictly increasing, less than zero at zero, and above zero at unity it follows that this equation has a unique solution in the unit interval. Hence it follows that $\mu_{0 | c_0, a_0, z_0}^m$ is also identified. Moreover, since the observed single-rate for unqualified males of cohort c_0 , that is $\tilde{\mu}_{0 | c_0, z_0}^m$, is a population weighted average of $\mu_{0 | c_0, a_0, z_0}^m$ and $\mu_{0 | c_0, a_1, z_0}^m$ it follows that $\mu_{0 | c_0, a_1, z_0}^m$ is also identified. By a corresponding argument the single rates for all female full types are also identified.

Once the single rates of all full types have been identified, it is easy to show in a third step—using that choices involving types $x \in \mathcal{X}^D$ have directly observable counterparts—that the equilibrium matching in full type, $\boldsymbol{\mu}$, can be recovered using equations (B2) through (B5). For instance, we can first use (B4) to recover the proportion of any distinguishable male type $x_i \in \mathcal{X}^D$ that marry type female type (c_0, a_0, z_0) . Since, for any such male type we observe the total proportion marrying unqualified women from cohort c_0 it follows

that we can also recover the proportion an any male type $x_i \in \mathcal{X}^D$ marrying type female type (c_0, a_1, z_0) . Having identified the marriage rate of male type (c_1, a_0, z_0) to each female full type $x_j \in \mathcal{X}$, we can then use (B2) to recover the marriage rate of males of type (c_0, a_0, z_0) to each female full types.

A corresponding argument shows that every marriage rate of every female type except type (c_0, a_1, z_0) is identified. In particular, the rate of marriage to male type (c_0, a_1, z_0) is identified for all females except type (c_0, a_1, z_0) . By (6) it then follows that the rate of marriage of male type (c_0, a_1, z_0) to all full female types except (c_0, a_1, z_0) is identified. Finally, since the single rate of male type (c_0, a_1, z_0) as well as their rates of marriage to every female type except (c_0, a_1, z_0) have all been identified it follows from adding up that their rate of marriage to (c_0, a_1, z_0) is also identified. A corresponding argument shows that the marriage rate of every full female type to every full male type is identified.

Having established that $\boldsymbol{\mu} \in \mathcal{M}^P$ is identified from $\tilde{\boldsymbol{\mu}} \in \tilde{\mathcal{M}}^P$, we can now rely on Choo and Siow (2006) to show that the parameters of the systematic surplus are identified. In particular, we know from Section 5 that $\Sigma(x_i, x_j) = \log\left(\mu_{x_j|x_i}^m / \mu_{0|x_i}^m\right) + \log\left(\mu_{x_i|x_j}^f / \mu_{0|x_j}^f\right)$ for every $(x_i, x_j) \in \mathcal{X} \times \mathcal{X}$. Given that $\boldsymbol{\mu} \in \mathcal{M}^P$, the recovered Σ will be in \mathcal{S}^P . Hence it will have the additive form in (8) whereby the nine ζ -terms are readily identified.

C. Estimation Procedure

This appendix contains details on the second step of our estimation procedure.

Our equilibrium marriage market model is indexed by a structural parameter vector $\theta \in \Theta$. There is an unknown true value of the structural parameter vector $\theta_0 \in \Theta$ that generates the observed data. We construct a Maximum Likelihood estimator for θ . The construction of the likelihood function requires the equilibrium marriage rates, the computation of which are also detailed here.

To keep the appendix self-contained, we start by reiterating some model notation. An individual is either male (superscript m) or female (superscript f). The model focuses on a set of cohorts $c \in \mathcal{C} = \{1953, \dots, 1960\}$, that are split halfway into a pre-RoSLA regime \mathcal{C}_0 and a post-RoSLA regime \mathcal{C}_1 . The marriage market participants are endowed with ability $a \in \mathcal{A} = \{a_0, a_1, a_2\}$, where a_0 indicates low ability, a_1 indicates medium ability, and a_2 indicates high ability. Each marriage market participant is further characterized by a qualification-level $z \in \mathcal{Z} = \{z_0, z_1, z_2\}$, with z_0 denoting no qualification, z_1 denoting a basic qualification level, and z_2 denoting an advanced qualification level. The assumed mapping from ability levels to qualifications is regime-specific and outlined in Assumption 1. The structure imposed by the mapping implies that four (a, z) -types exist in \mathcal{C}_0 though only three of these exist in \mathcal{C}_1 . Hence in total the full type-space $\mathcal{X} \subset \mathcal{C} \times \mathcal{A} \times \mathcal{Z}$ features $N = 4 \times |\mathcal{C}_0| + 3 \times |\mathcal{C}_1| = 28$ types.

C.1. Computing Equilibrium Choice Probabilities Here we outline the algorithm used to solve for the equilibrium choice probabilities at any trial value for the parameters.

Out-of-Sample Marriages. Before we turn to the equilibrium computation algorithm, we need to deal with complications arising from out-of-sample marriages. Some sample individuals are observed being married to out-of-sample cohort spouses. Such out-of-sample marriages will occur relatively frequently for males towards the end of \mathcal{C} and for females in the beginning of \mathcal{C} , but will – as large age gaps are quite rare – be quite infrequent for individuals born close to the reform threshold. In order to account for marriages to non-sample-cohort spouses we make use of the age gap preferences.

Indeed, let $pre_{c_i}^m(\theta)$ be the share of c_i -males' marriages that involve pre-sample females ($c_j < 1953$), and let $post_{c_i}^m(\theta)$ be the share of c_i -males' marriages that involve post-sample females ($c_j > 1960$). Denote the analogously defined shares for women by $pre_{c_j}^f(\theta)$ and $post_{c_j}^f(\theta)$. Since $d_{ij} = c_j - c_i$, $pre_{c_i}^m(\theta)$ is the share of c_i -males' marriages that involves age gaps $d_{ij} \leq \min\{\mathcal{C}\} - 1 - c_i$, and $post_{c_i}^m(\theta)$ is the share of c_i -males' marriages that involve age gaps $d_{ij} \geq \max\{\mathcal{C}\} + 1 - c_i$. Similarly, a woman from cohort c_j is married to a pre-sample male if $d_{ij} \geq c_j - \min\{\mathcal{C}\} + 1$ and she is married to a post-sample male if $d_{ij} \leq c_j - \max\{\mathcal{C}\} - 1$.

We approximate $pre_{c_i}^m$ to be the relative weight put on age gaps $d_{ij} \leq \min\{\mathcal{C}\} - 1 - c_i$ by to the age gap component of the (expected) marital surplus function (conditional on $d_{ij} \in \{1 - |\mathcal{C}|, \dots, |\mathcal{C}| - 1\}$). Formally, we define

$$pre_{c_i}^m(\theta) = \frac{\sum_{d=1-|\mathcal{C}|}^{\min\{\mathcal{C}\}-1-c_i} \exp\left(\frac{\lambda(d;\theta)}{2}\right)}{\sum_{d=1-|\mathcal{C}|}^{|\mathcal{C}|-1} \exp\left(\frac{\lambda(d;\theta)}{2}\right)} \quad \text{and} \quad post_{c_i}^m(\theta) = \frac{\sum_{d=\max\{\mathcal{C}\}+1-c_i}^{|\mathcal{C}|-1} \exp\left(\frac{\lambda(d;\theta)}{2}\right)}{\sum_{d=1-|\mathcal{C}|}^{|\mathcal{C}|-1} \exp\left(\frac{\lambda(d;\theta)}{2}\right)}. \quad (\text{C1})$$

The female out-of-sample shares, $pre_{c_j}^f(\theta)$ and $post_{c_j}^f(\theta)$ are defined analogously to (C1). We detail below

how we use $pre_{c_i}^m(\boldsymbol{\theta})$, $post_{c_i}^m(\boldsymbol{\theta})$, $pre_{c_j}^f(\boldsymbol{\theta})$, and $post_{c_j}^f(\boldsymbol{\theta})$ in the computation of the equilibrium and in the estimation procedure.

Equilibrium Algorithm. The equilibrium outcome of interest is a set of never-married rates and marriage rates by gender and type. We compute these using the following algorithm.

0. Initiate the algorithm with structural parameter vector $\boldsymbol{\theta}$ and a N -vector of type-specific single rates for each gender, denoted $\hat{\boldsymbol{\mu}}_0^m$ and $\hat{\boldsymbol{\mu}}_0^f$ for males and females, respectively. The generic elements of $\hat{\boldsymbol{\mu}}_0^m$ and $\hat{\boldsymbol{\mu}}_0^f$ are $\hat{\mu}_{0|x_i}^m$ and $\hat{\mu}_{0|x_j}^f$, respectively. Go to step 1.
1. Using $pre_{c_i}^m(\boldsymbol{\theta})$, $post_{c_i}^m(\boldsymbol{\theta})$, $pre_{c_j}^f(\boldsymbol{\theta})$, and $post_{c_j}^f(\boldsymbol{\theta})$, predict the gender- and type-specific out-of-sample marriage rates as $\hat{\mu}_{pre|x_i}^m = (1 - \hat{\mu}_{0|x_i}^m)pre_{c_i}^m(\boldsymbol{\theta})$, $\hat{\mu}_{post|x_i}^m = (1 - \hat{\mu}_{0|x_i}^m)post_{c_i}^m(\boldsymbol{\theta})$, $\hat{\mu}_{pre|x_j}^f = (1 - \hat{\mu}_{0|x_j}^f)pre_{c_j}^f(\boldsymbol{\theta})$, and $\hat{\mu}_{post|x_j}^f = (1 - \hat{\mu}_{0|x_j}^f)post_{c_j}^f(\boldsymbol{\theta})$. There are $2N$ out-of-sample marriage rates for each gender. Go to step 2.

2. Given $\boldsymbol{\theta}$ and the candidate vectors of single rates, compute the gender-specific candidate marriage rates

$$\hat{\mu}_{x_j|x_i}^m = \left[\hat{\mu}_{0|x_i}^m \hat{\mu}_{0|x_j}^f \frac{h^f(x_j)}{h^m(x_i)} \right]^{1/2} \exp \left[\frac{\Sigma(x_i, x_j; \boldsymbol{\theta})}{2} \right], \quad (\text{C2})$$

and

$$\hat{\mu}_{x_i|x_j}^f = \left[\hat{\mu}_{0|x_i}^m \hat{\mu}_{0|x_j}^f \frac{h^m(x_i)}{h^f(x_j)} \right]^{1/2} \exp \left[\frac{\Sigma(x_i, x_j; \boldsymbol{\theta})}{2} \right], \quad (\text{C3})$$

as implied by the equilibrium allocation. There are N^2 marriage rates for each gender. Stack the type-specific never-married and (in- and out-of-sample) marriage rates in vectors $\hat{\boldsymbol{\mu}}^m$ for males and $\hat{\boldsymbol{\mu}}^f$ for females. The vectors $\hat{\boldsymbol{\mu}}^m$ and $\hat{\boldsymbol{\mu}}^f$ each contain $N + 2N + N^2$ rates. Go to step 3.

3. Given $\hat{\boldsymbol{\mu}}^m$ and $\hat{\boldsymbol{\mu}}^f$, compute the implied gender and type-specific excess supplies, denoted $\hat{\Delta}_{x_i}^m$ and $\hat{\Delta}_{x_j}^f$, and defined as

$$\hat{\Delta}_{x_i}^m \equiv \hat{\mu}_{0|x_i}^m + \hat{\mu}_{pre|x_i}^m + \hat{\mu}_{post|x_i}^m + \sum_{x_j \in \mathcal{X}} \hat{\mu}_{x_j|x_i}^m - 1, \quad (\text{C4})$$

and

$$\hat{\Delta}_{x_j}^f \equiv \hat{\mu}_{0|x_j}^f + \hat{\mu}_{pre|x_j}^f + \hat{\mu}_{post|x_j}^f + \sum_{x_i \in \mathcal{X}} \hat{\mu}_{x_i|x_j}^f - 1. \quad (\text{C5})$$

Stack the gender- and type-specific excess supplies in the $2N$ -vector $\hat{\Delta}$. In equilibrium, excess supply is zero for each gender and all types. Let $\|\cdot\|_\infty$ be the uniform norm. If the candidate equilibrium marriage rates leaves no excess supply, i.e. if $\|\hat{\Delta}\|_\infty \leq \epsilon$, where we take $\epsilon = 10^{-5}$, terminate the algorithm and take $\boldsymbol{\mu}^m(\boldsymbol{\theta}) = \hat{\boldsymbol{\mu}}^m$ and $\boldsymbol{\mu}^f(\boldsymbol{\theta}) = \hat{\boldsymbol{\mu}}^f$ as the equilibrium marriage (and never-married) rates.⁴⁹ If $\|\hat{\Delta}\|_\infty > \epsilon$, go to step 4.

4. Update the candidate equilibrium never-married rates with a simplified and (dampened) Newton-step. Let $\hat{\mathbf{J}}$ be the $2N \times 2N$ -matrix with partial derivatives of the excess supply vector $\hat{\Delta}$ with respect to the log single rate of own (gender and) type on the diagonal, and zeros on the off-diagonals. Hence, $\hat{\mathbf{J}}$ is the Jacobian of $\hat{\Delta}$, taken with respect to the log single rates, but with the off-diagonal elements

⁴⁹Our notation for the resulting equilibrium rates emphasizes their dependence on the unknown structural parameter vector $\boldsymbol{\theta}$.

set to zero. For example, the diagonal entry in $\hat{\mathbf{J}}$ corresponding to a male type- x_i individual therefore contains

$$\frac{\partial \hat{\Delta}_{x_i}^m}{\partial \log \hat{\mu}_{0|x_i}^m} = \left(1 - pre_{x_i}^m(\boldsymbol{\theta}) - post_{x_i}^m(\boldsymbol{\theta}) + \frac{1}{2} \left[\frac{\hat{\mu}_{0|x_j}^f}{\hat{\mu}_{0|x_i}^m} \cdot \frac{h^f(x_j)}{h^m(x_i)} \right]^{1/2} \exp \left[\frac{\Sigma(x_i, x_j; \boldsymbol{\theta})}{2} \right] \right) \hat{\mu}_{0|x_i}^m. \quad (\text{C6})$$

Stack the gender-specific vector of candidate never-married rates $\hat{\mu}_0^m$ and $\hat{\mu}_0^f$ in $\hat{\boldsymbol{\mu}}_0$, and Newton-update $\log \hat{\boldsymbol{\mu}}_0$ according to

$$\log \boldsymbol{\mu}'_0 = \log \hat{\boldsymbol{\mu}}_0 + K \cdot \hat{\mathbf{J}}^{-1} \hat{\boldsymbol{\Delta}}, \quad (\text{C7})$$

where $K < 1$ is a scalar dampening factor. We take $K = 2/3$. Update $\hat{\boldsymbol{\mu}}_0 = \boldsymbol{\mu}'_0$. Go to step 1 with the updated never-married rate candidates $\hat{\boldsymbol{\mu}}_0 = (\hat{\mu}_0^m, \hat{\mu}_0^f)$.⁵⁰

C.2. The Maximum Likelihood Estimator The likelihood function represents the likelihood of the structural model delivering the marriage market choices observed in the data as a function of the structural parameter vector $\boldsymbol{\theta}$. Hence, we must confront the fact that only qualification-levels and academic cohorts are observed, whereas ability is unobserved.

Empirical Marriage Frequencies. In order to conveniently represent our data and derive the likelihood function we slightly expand the notation. There are three observable types (qualification-levels) for each cohort. Let $\tilde{\mathcal{X}} = \mathcal{C} \times \mathcal{Z}$ be the observable type-space, with generic element \tilde{x} . Let $\tilde{\mathcal{X}}_+ = \tilde{\mathcal{X}} \cup \{0, pre, post\}$ be the observable spouse type-space which is the same as the own type-space but extended to non-marriage, and marriage to an out-of-sample spouse. We have $|\tilde{\mathcal{X}}| = 24$ and $|\tilde{\mathcal{X}}_+| = 27$.

Our data can be represented by a matrix of male marriage frequencies, $\mathbf{M}^m = \{M_{(\tilde{x}, \tilde{x}')}^m : (\tilde{x}, \tilde{x}') \in \tilde{\mathcal{X}} \times \tilde{\mathcal{X}}_+\}$, and a matrix of female marriage frequencies, $\mathbf{M}^f = \{M_{(\tilde{x}, \tilde{x}')}^f : (\tilde{x}, \tilde{x}') \in \tilde{\mathcal{X}} \times \tilde{\mathcal{X}}_+\}$, for females. Here, $M_{(\tilde{x}, \tilde{x}')}^m$ is the number of observed type- \tilde{x} males making marriage choice \tilde{x}' . Similarly, $M_{(\tilde{x}, \tilde{x}')}^f$ is the number of observed type- \tilde{x} females making marriage choice \tilde{x}' .

Choice Probabilities. Model identification requires that the observed marriage frequencies on $\tilde{\mathcal{X}} \times \tilde{\mathcal{X}}_+$ contain enough information to uniquely recover the structural parameter vector that governs marriage choices on $\mathcal{X} \times \mathcal{X}_+$, where $\mathcal{X}_+ = \mathcal{X} \cup \{0, pre, post\}$. As described in the main text, RoSLA and the associated gender-specific discontinuities in the qualification distribution (which identify γ^m and γ^f), along with Assumption 1, pin down the gender-, cohort- and qualification-specific distributions of ability, $\Pr^m(a_i|c_i, z_i)$ and $\Pr^f(a_j|c_j, z_j)$.

Let $\mu_{0|c_i, a_i, z_i}^m(\boldsymbol{\theta})$ be the equilibrium probability that a male of full type (c_i, a_i, z_i) remains single, and likewise, let $\mu_{0|c_j, a_j, z_j}^f(\boldsymbol{\theta})$ be the equilibrium probability that a female of full type (c_j, a_j, z_j) remains single. The implied never-married rate for the observable male type (c_i, z_i) is

$$\tilde{\mu}_{0|c_i, z_i}^m(\boldsymbol{\theta}) = \sum_{a_i \in \mathcal{A}} \mu_{0|c_i, a_i, z_i}^m(\boldsymbol{\theta}) \Pr^m(a_i|c_i, z_i), \quad (\text{C8})$$

with an analogous expression for the observable female never-married rates $\tilde{\mu}_{0|c_j, z_j}^f(\boldsymbol{\theta})$. There are $|\mathcal{C}| \times |\mathcal{Z}| =$

⁵⁰Our Newton-step update only uses the diagonal elements of the Jacobian, but nonetheless converges almost instantaneously.

24 unconditional (on ability) never-married rates for each gender.

The equilibrium marriage rates are aggregated across the unobserved ability-levels in a similar fashion. Consider first the out-of-sample marriage rates. Let $\mu_{pre|c_i, a_i, z_i}^m(\boldsymbol{\theta})$, $\mu_{post|c_i, a_i, z_i}^m(\boldsymbol{\theta})$, $\mu_{pre|c_j, a_j, z_j}^f(\boldsymbol{\theta})$ and $\mu_{post|c_j, a_j, z_j}^f(\boldsymbol{\theta})$ be the equilibrium male and female, full-type-specific pre- and post-sample marriage rates. Aggregating across the unobserved ability levels yields out-of-sample marriage rates with observable counter-parts. For males, we have

$$\tilde{\mu}_{pre|c_i, z_i}^m(\boldsymbol{\theta}) = \sum_{a_i \in \mathcal{A}} \mu_{pre|c_i, a_i, z_i}^m(\boldsymbol{\theta}) \Pr^m(a_i | c_i, z_i), \quad (\text{C9})$$

and

$$\tilde{\mu}_{post|c_i, z_i}^m(\boldsymbol{\theta}) = \sum_{a_i \in \mathcal{A}} \mu_{post|c_i, a_i, z_i}^m(\boldsymbol{\theta}) \Pr^m(a_i | c_i, z_i), \quad (\text{C10})$$

where female rates are analogously defined. There are $|\mathcal{C}| \times |\mathcal{Z}| = 24$ observable pre-sample rates and $|\mathcal{C}| \times |\mathcal{Z}| = 24$ observable post-sample rates for each gender.

Finally, let $\mu_{c_j, a_j, z_j | c_i, a_i, z_i}^m(\boldsymbol{\theta})$ be the equilibrium probability that male of full type (c_i, a_i, z_i) marries a woman of full type (c_j, a_j, z_j) , and let $\mu_{c_i, a_i, z_i | c_j, a_j, z_j}^f(\boldsymbol{\theta})$ be the equilibrium probability that a woman of full type (c_j, a_j, z_j) marries male of full type (c_i, a_i, z_i) male. In terms of observable types, aggregating over both the own and the spouse ability we have that

$$\tilde{\mu}_{c_j, z_j | c_i, z_i}^m(\boldsymbol{\theta}) = \sum_{a_i \in \mathcal{A}} \sum_{a_j \in \mathcal{A}} \mu_{c_j, a_j, z_j | c_i, a_i, z_i}^m(\boldsymbol{\theta}) \Pr^m(a_i | z_i, c_i), \quad (\text{C11})$$

with an analogous expression for the female rates. There are $(|\mathcal{C}| \times |\mathcal{Z}|)^2 = 576$ observable marriage rates for each gender.

The Likelihood Function. Having characterized the model implied choice probabilities that a male or female of a given type marries a spouse of a given type, marries out-of-sample, or never marries, as the case might be, the likelihood that the model delivers the observed marriage frequencies \mathbf{M}^m and \mathbf{M}^f is easily obtained. Indeed, the log-likelihood of observing \mathbf{M}^m and \mathbf{M}^f is

$$\ell(\boldsymbol{\theta} | \mathbf{M}^m, \mathbf{M}^f) = \sum_{g \in \{m, f\}} \sum_{\tilde{x} \in \tilde{\mathcal{X}}} \sum_{\tilde{x}' \in \tilde{\mathcal{X}}_+} M_{(\tilde{x}, \tilde{x}')}^g \log \tilde{\mu}_{\tilde{x}' | \tilde{x}}^k(\boldsymbol{\theta}), \quad (\text{C12})$$

where the $\tilde{\mu}_{\tilde{x}' | \tilde{x}}^k(\boldsymbol{\theta})$ s are given by (C8), (C9), (C10), and (C11). The Maximum Likelihood estimator of $\boldsymbol{\theta}_0$ is

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta} \in \Theta} \ell(\boldsymbol{\theta} | \mathbf{M}^m, \mathbf{M}^f). \quad (\text{C13})$$

Regularity conditions and standard arguments implies that $\text{plim } \hat{\boldsymbol{\theta}} = \boldsymbol{\theta}_0$ and that $\hat{\boldsymbol{\theta}}$ is asymptotically Normal distributed. In fact,

$$\sqrt{M}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}), \quad (\text{C14})$$

where $\boldsymbol{\Omega} = \mathbf{E}[-\mathbf{H}(\mathbf{M}^m, \mathbf{M}^f; \boldsymbol{\theta}_0)]^{-1}$, $\mathbf{H}(\mathbf{M}^m, \mathbf{M}^f; \boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}}^2 \ell(\boldsymbol{\theta} | \mathbf{M}^m, \mathbf{M}^f)$ is the Hessian matrix of the log-likelihood function (C12), and the expectation is taken with respect to the distribution of $(\mathbf{M}^m, \mathbf{M}^f)$. $M = \sum_{g \in \{m, f\}} \sum_{\tilde{x} \in \tilde{\mathcal{X}}} \sum_{\tilde{x}' \in \tilde{\mathcal{X}}_+} M_{(\tilde{x}, \tilde{x}')}^g$ is the number of observations employed in the estimation. The asymptotic variance-covariance matrix of $\hat{\boldsymbol{\theta}}$ is therefore $M^{-1} \boldsymbol{\Omega}$. In the empirical implementation, $\boldsymbol{\Omega}$ is

estimated by

$$\hat{\mathbf{\Omega}} = \left[-M^{-1} \sum_{g \in \{m, f\}} \sum_{\tilde{x} \in \tilde{\mathcal{X}}} \sum_{\tilde{x}' \in \tilde{\mathcal{X}}_+} M_{(\tilde{x}, \tilde{x}')}^g \nabla_{\hat{\boldsymbol{\theta}}}^2 \log \mu_{\tilde{x}'|\tilde{x}}^g(\hat{\boldsymbol{\theta}}) \right]^{-1}. \quad (\text{C15})$$

Under standard regularity conditions, $\hat{\mathbf{\Omega}}$ is a consistent estimator of $\mathbf{\Omega}$. That is, $\text{plim } \hat{\mathbf{\Omega}} = \text{E}[-\nabla_{\boldsymbol{\theta}}^2 \log \mu_{\tilde{x}'|\tilde{x}}^g(\hat{\boldsymbol{\theta}})]^{-1}$. We consequently report standard errors and draw inference based on the (consistently estimated) asymptotic variance-covariance matrix of $\hat{\boldsymbol{\theta}}$,

$$\text{A}\hat{\text{var}}(\hat{\boldsymbol{\theta}}) = M^{-1} \hat{\mathbf{\Omega}} = \left[- \sum_{g \in \{m, f\}} \sum_{\tilde{x} \in \tilde{\mathcal{X}}} \sum_{\tilde{x}' \in \tilde{\mathcal{X}}_+} M_{(\tilde{x}, \tilde{x}')}^g \nabla_{\hat{\boldsymbol{\theta}}}^2 \log \mu_{\tilde{x}'|\tilde{x}}^g(\hat{\boldsymbol{\theta}}) \right]^{-1}. \quad (\text{C16})$$

We note that (C16) does not account for estimation errors introduced from the first step in the estimation procedure.

D. The Model without Unobserved Ability

In this Appendix we present estimates from the constrained version of our model that omits ability. In the main text we have referred to this as the “standard” specification. In this case an individual’s type is defined as $x \in \mathcal{C} \times \mathcal{Z}$ with all combinations possible. We impose the same separability assumption on marital surplus as in (9). We also impose the same age-gap preference structure and the same trend structure except that trends are now defined directly over qualification types.

The model estimated here thus takes the form

$$\Sigma(c_i, z_i; c_j, z_j) = \zeta(z_i, z_j) + \lambda(c_j - c_i) + \tau^m(c_i, z_i) + \tau^f(c_j, z_j). \quad (\text{D1})$$

with $\lambda(d_{ij})$ defined as in (10), and

$$\tau^g(c_i, z_i) = \sum_{z \in \mathcal{Z}} \left[\beta_z^g(c_i - 1953) \mathbb{1}(z_i = z) + \beta_{z, \mathcal{C}_1}^g(c_i - 1956) \mathbb{1}(z_i = z) \mathbb{1}(c_i \in \mathcal{C}_1) \right], \quad g = m, f. \quad (\text{D2})$$

The estimated parameters are shown in tables D1 and D2. This loss in the maximized log likelihood value relative to our main model in this case is 163.7 (and larger than the corresponding loss of 107.8 for the restricted model presented in the text where trends were defined on ability). There is no visible difference in the fit to the observed aggregate assortative mating on qualifications compared with our full model and the same applies to the aggregate age gap distribution. (Details of these model predictions are available from authors on request.) The central difference between our main model and the one estimated here without latent ability is in terms of predicted never-married rates. Figure D.1 presents the fit to the empirical never-married rates of the model without latent ability. The main observation is that it starkly mispredicts how the never-married rate of unqualified men and women changed at the reform threshold, predicting a sharp reduction in response to the reduced supply brought about by the RoSLA.

Table D1: Estimates of Marital Surplus by Qualification Profile in Constrained Model

	Females	No	Basic	Adv.
Males		Qual.	Qual.	Qual.
No Qual.		0.244*** (0.076)	-1.413*** (0.079)	-4.350*** (0.091)
Basic Qual.		-1.701*** (0.083)	-1.088*** (0.083)	-3.029*** (0.090)
Adv. Qual.		-3.762*** (0.095)	-2.285*** (0.089)	-0.482*** (0.090)

Notes: See notes to Table 5. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D2: Estimates of Marital Surplus: Age Gap and Trend Terms in Constrained Model

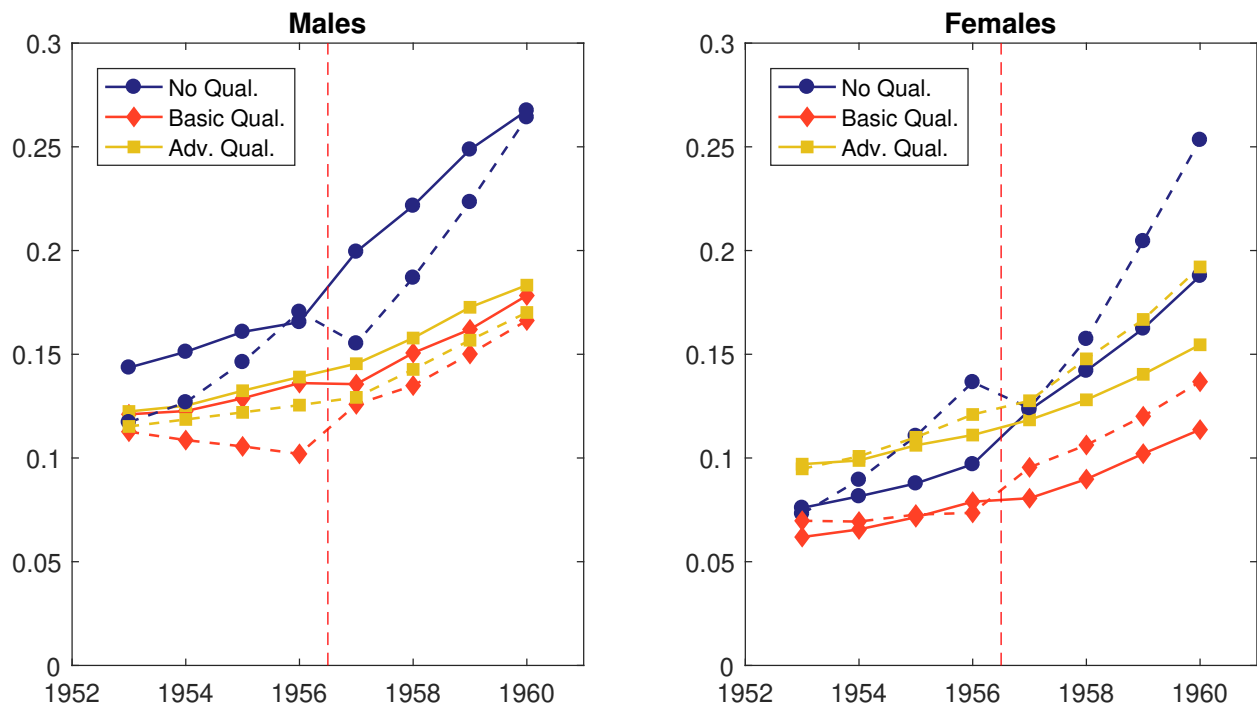


Figure D.1: Model Predicted (Constrained Model) and Empirical Never-Married Rates by Cohort, Gender, and Qualification Level