Constraints on Marriage Timing and Women's Educational Investments

Asha Bharadwaj*

April 15, 2026

Click here for the latest version

Abstract

The Indian government has been debating raising the minimum legal age of marriage for women from 18 to 21. This paper asks how the timing of marriage affects women's educational attainment in India. Using variation generated by a common social norm that daughters marry in birth order, I estimate the causal effect of age at marriage on schooling. Preliminary estimates from the 2006 Rural Economic and Demographic Survey (REDS) suggest that delaying marriage by one year increases women's probability of completing high school by 4.4 percentage points.

I then develop and estimate an equilibrium model of marriage and schooling decisions to analyze the effects of raising the minimum legal age of marriage. In the model, households choose dowry offers and educational investments over daughters' life cycles in a directed-search marriage market with dowries. They trade off early marriage against delaying marriage to invest in education, recognizing that these choices shape future match opportunities, dowries, and the marriage prospects of younger siblings. The model is used to quantify the effects of the reform on education, dowries, marriage-market returns to education, and welfare, including heterogeneity by household wealth.

1 Introduction

Among women ages 20 to 24 in India, less than half had completed high school in 2019, while a majority were already married. Because married women in India rarely continue schooling

*Department of Economics, University of Wisconsin–Madison.

after marriage, policies that delay marriage may have important effects on educational attainment. This question has become especially salient as the Indian government has debated raising the minimum legal age of marriage for women from 18 to 21.

This paper studies how constraints on marriage timing affect women's educational attainment, age at marriage, dowries, and welfare. I first estimate the causal effect of delayed marriage on women's education by exploiting exogenous variation in age at marriage. This variation is generated by the practice that daughters marry in order of birth within the household. I then develop and estimate an equilibrium model of household decision-making to evaluate the effects of a minimum marriage age law. The key trade-off faced by families is between investing in daughters' education and the timing of marriage.

India is an especially important setting for this question because female labor force participation remains low, making the marriage market a central channel through which returns to women's education are realized (Andrew and Adams, 2022). Moreover, the prevalence of dowries, which are transfers from the bride's family to the groom's family at the time of marriage, provides a measure of how age and education are valued in the marriage market. In addition, about half of women marry before age 21, so a policy that raises the minimum marriage age to 21 could have large effects.

In the first part of the paper, I estimate the causal effect of marriage timing on women's education by exploiting plausibly exogenous variation in age at marriage generated by a common practice in India. A key institutional feature of the setting is that marriages are largely arranged by parents or relatives. In families with multiple daughters, parents tend to arrange marriages in order of birth. In the presence of search frictions in the marriage market, this implies that younger daughters tend to marry later than the oldest daughter in the household. I exploit this exogenous variation in marriage age by comparing women with the same overall birth order in families of the same size, restricting to families with at least two daughters. Specifically, I use an indicator for whether a woman has an older sister as an instrument for age at marriage. The key identifying assumption is that having an older sister affects educational outcomes only through the timing of marriage. Using the 2006 Rural Economic and Demographic Survey, I find that having an older sister increases a woman's age at marriage by about 0.4 years. Two-stage least squares estimates imply that delaying marriage by one year increases women's schooling by about 1.7 years and raises the probability of completing high school by 5.4 percentage points. These results indicate that marriage timing is an important determinant of educational attainment in this setting.

The design-based analysis identifies how an individual-level constraint on marriage timing affects educational outcomes, but it is not sufficient to evaluate a minimum marriage age reform. Such a law changes the set of feasible matches and therefore the returns to age and

education in the marriage market. Moreover, if daughters are required to remain unmarried until age 21, households may need to support them for longer, which could affect dowries and educational attainment. More broadly, evaluating the welfare consequences of the reform requires a framework that quantifies the trade-offs households face between marriage timing, educational investment, and marriage-market outcomes. Finally, because minimum marriage age laws are imperfectly enforced in practice, their effects depend not only on the legal cutoff itself but also on the extent of enforcement. For these reasons, the policy cannot be evaluated using the instrumental variables strategy alone.

To evaluate a minimum marriage age reform under imperfect enforcement, in the second part of the paper, I develop and estimate a dynamic equilibrium model of marriage and education decisions. The model features households with two daughters who interact with men in a directed-search marriage market over the life cycle. Households enter the model when the first-born daughter turns 16, and in each subsequent period the household makes marriage and schooling choices for its daughters.

In each period, the household posts a menu of dowries for the oldest unmarried daughter, specifying the dowry it would pay to each type of man. In doing so, the household trades off the benefit of attracting more proposals against the cost of paying a higher dowry. Men observe these posted dowries and decide which household to approach with a proposal. Because men cannot coordinate their application decisions, some households may receive many proposals while others receive none. Men therefore take into account not only the dowry offered by a household, but also how many other men are likely to apply to that household.

Households then choose among the proposals they receive, if any, and transfer the posted dowry upon marriage. When making this decision, households take into account the utility from the man's type and the dowry to be paid. Because households have multiple daughters and search for them sequentially, they also take into account the continuation value of marrying off the current daughter and being able to start searching for the next daughter in future periods. Unmarried men and women search again in the next period when they are a year older. This structure allows the model to capture the joint evolution of daughters' age, schooling, and marriage prospects.

The model embeds this marriage market in a dynamic household decision problem. After the marriage stage, the household chooses schooling investments for daughters who remain unmarried, taking into account how current schooling affects future marriage-market prospects. Education is therefore valuable because it can improve future match outcomes, but acquiring education requires postponing marriage. Postponing marriage is costly because women's age is negatively valued in the marriage market and the set of potential spouses

shrinks with age. Additionally, delaying one daughter's marriage also delays marriage for her younger sister.

I estimate the model using indirect inference. The results from first-part of the paper help me discipline the model's estimates of marriage-market returns to education. I then use the estimated model to study the effects of raising the minimum legal age of marriage from 18 to 21. Under the reform, more women acquire education, since the opportunity cost of delaying marriage to acquire education is much lower. However, I find sharp differences in the effects of the reform by income. The increase in education is concentrated among high-income women, while low-income women do not acquire additional education. My next set of counterfactuals will allow for imperfect enforcement, so that some underage marriages may still occur, and will compare the effects of the minimum marriage age law with those of a compulsory schooling reform, alongside an analysis of how welfare changes under the law.

This paper relates to several strands of literature. Some existing work shows that delaying marriage causally raises women's schooling in developing countries (Field and Ambrus, 2008; Wilson, 2022). A separate literature documents that parents in South Asia often arrange daughters' marriages in birth order, so that sibling composition shifts the timing of marriage (Vogl, 2013; Keskar et al., 2025; Malhotra and Tsui, 1996). Building on this insight, I use sex-specific birth order, conditional on family size and sibling composition, as an instrument for age at marriage. This provides new causal evidence from India on how constraints on marriage timing affect women's educational attainment. Wilson (2022) also finds evidence of imperfect enforcement of minimum marriage age laws, especially in rural areas, which motivates the need for a structural model to understand the implications of such laws.

This paper also relates to work on dowry and other marriage payments in South Asia. Dowry payments in India are large, have risen over time, and vary systematically with characteristics of the bride and groom (Anderson, 2007; Weaver and Chiplunkar, 2022). Related work shows that the direction of marriage payments shapes household decisions about marriage timing and investment in daughters (Corno et al., 2016; Ashraf et al., 2020). This emphasis on marriage-market incentives is consistent with Andrew and Adams (2022), who show that returns to women's education in India operate largely through the marriage market rather than the labor market. Taken together, these findings motivate the central mechanisms in my model: households choose schooling and the timing of marriage in anticipation of marriage-market returns, and post type-contingent dowries to influence match outcomes.

The paper also contributes to the search and matching literature. Existing models with search frictions typically treat education as fixed at the point of entry into search (Beauchamp et al., 2021), or model education choices before agents enter the marriage market (Chiappori

et al., 2009, 2018; Ciscato, 2021).[1] By contrast, in my setting, households make schooling decisions during the search process, so realized match outcomes affect future schooling decisions, while schooling in turn shapes future marriage-market prospects. This structure allows the model to capture dynamic trade-offs between marriage timing and educational investment over the life cycle.

The remainder of the paper is structured as follows. Section 2 provides a brief overview of the data. Section 3 offers motivation and describes the setting. Empirical findings are discussed in Section 4. In Section 5, I discuss the model setup. Preliminary results using the model are in Section 6.

2 Data

I use two sources of data for this study. First, the National Family Health Survey (NFHS) is a repeated cross-sectional survey of representative households in India, comprising five waves of data available from 1992, 1998, 2005, 2015, and 2019. The survey collects data on demographic characteristics, fertility, reproductive health, and nutrition. Additionally, it includes information on women's age at marriage and educational attainment. The 2019 sample consisted of data from 636,691 households across India. I use the NFHS to provide descriptive facts and figures.

Second, I use data from the Rural Economic and Demographic Survey (REDS), a detailed panel survey of rural households conducted by the National Council of Applied Economic Research across the 17 most populous states in India. Data were collected in 1969-81, 1982, 1999 and 2006. The original sample of villages was drawn in the first round and was meant to be a representative sample of the rural population in India in those states. All original households, including split households, were followed over time, and retrospective information was collected on deceased members. The 17 surveyed states contain roughly 96% of the population of India [2]. The 1999 and 2006 rounds of the survey collected detailed retrospective information on the marriages of the household head, their parents, brothers and sisters, and

[1] A related labor-market literature studies endogenous human-capital investment in environments with search frictions. Moen and Rosén (2004) analyze firms' provision of general training under competitive search with poaching, and Flinn et al. (2017) model cooperative worker-firm investment in general and match-specific ability during ongoing employment. In both cases, the price-setting agent invests in the searching agents' productivity, whereas in my setting the price-setting agent (the household) chooses investments that affect its own daughters' marriage-market prospects. The labor market analog would be the firm investing in its own productivity.

[2] Since this is a rural sample, my results may not extend to urban areas. However, I find comparable results using the Indian Human Development Survey (IHDS), which is nationally representative.

their sons and daughters. Currently, I use only the 2006 round of the survey and use the questionnaire on the children of the household head.

The REDS data also reports the nominal value of gifts and cash transfers from the household of the bride to that of the groom at the time of marriage, as well as from the household of the groom to that of the bride.

3 Motivation and Setting

3.1 Marriages in India

Three features of marriage in India are central to this paper. First, marriages are predominantly arranged by parents or relatives rather than by the bride and groom themselves. In the 2012 IHDS, only 5 percent of women reported arranging their own marriages, while about 72 percent reported that their marriages were arranged solely by parents or relatives. This is important because it implies that parents, rather than women, are the relevant decision-makers in choices about marriage timing and educational investment.

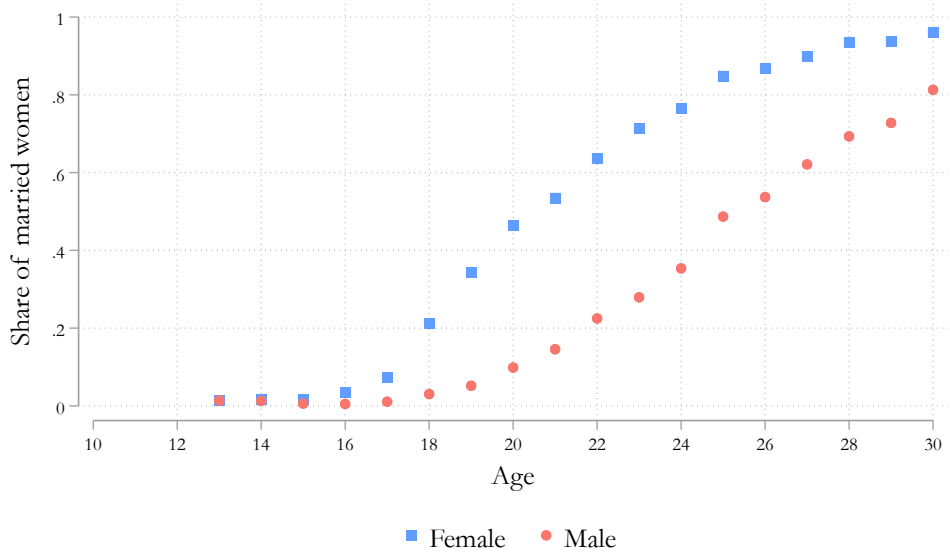
Second, dowry payments are widespread and quantitatively important. In most Indian marriages, the bride's family transfers cash and gifts to the groom's family. Using data from the 1999 Rural Economic and Demographic Survey, Weaver and Chiplunkar (2022) show that around 90 percent of marriages in each wealth category involve a dowry payment. In the 2012 IHDS, the median dowry amount was approximately Rs. 16,500, or around USD 300 in 2012. These transfers are large relative to household resources: the median household spends more than twice its annual income on a daughter's wedding. Dowries therefore represent a major financial commitment for households and provide a direct measure of how women's age and education are valued in the marriage market.

Third, early marriage for women remains common in India. Figure 1 shows that nearly half of women in India were married by age 20 in 2019, while men tend to marry later, with about half married by age 25. The current law prohibits marriage for women under age 18 and for men under age 21. However, enforcement is imperfect: around 10 percent of both men and women marry before the legal minimum marriage age. This makes marriage timing an important policy margin and highlights that the effects of marriage-age laws depend not only on the legal cutoff itself, but also on how strictly the law is enforced.

3.2 Education in India

School enrollment in India has risen substantially in recent decades for both boys and girls. Schooling is compulsory from ages 6 to 14, and enrollment in primary and middle school is

Figure 1: Share ever-married by age



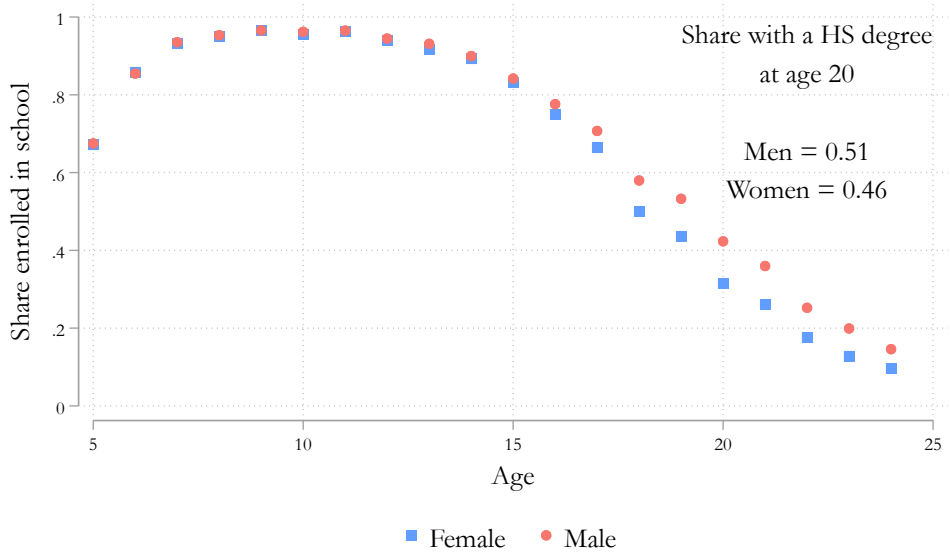
Source: Own calculations using NFHS 2019-21.
Sample: Individuals between 13 to 30 years of age

nearly universal. Figure 2 shows that school enrollment rates for both boys and girls were close to 100 percent at younger ages in 2019. However, enrollment declines sharply at the secondary level, with a somewhat steeper drop for girls. By age 20, only 46 percent of women had completed high school in 2019.[3]

This decline in girls' schooling is unlikely to be explained by strong labor-market opportunities outside school (see Section 13). In the 2019 NFHS, only around 20 percent of women who were not enrolled in school at age 18 reported having worked in the last 12 months, compared with around 80 percent of men of the same age. This suggests that, for many girls, leaving school is closely tied to marriage rather than to labor-market entry.

[3]Because this figure is based on the NFHS, it may partly reflect cohort differences. Using the panel structure of the IHDS, I find similar patterns: among girls who were enrolled in middle school in 2005, 62.5 percent report having completed high school by 2012.

Figure 2: Share enrolled in school by gender

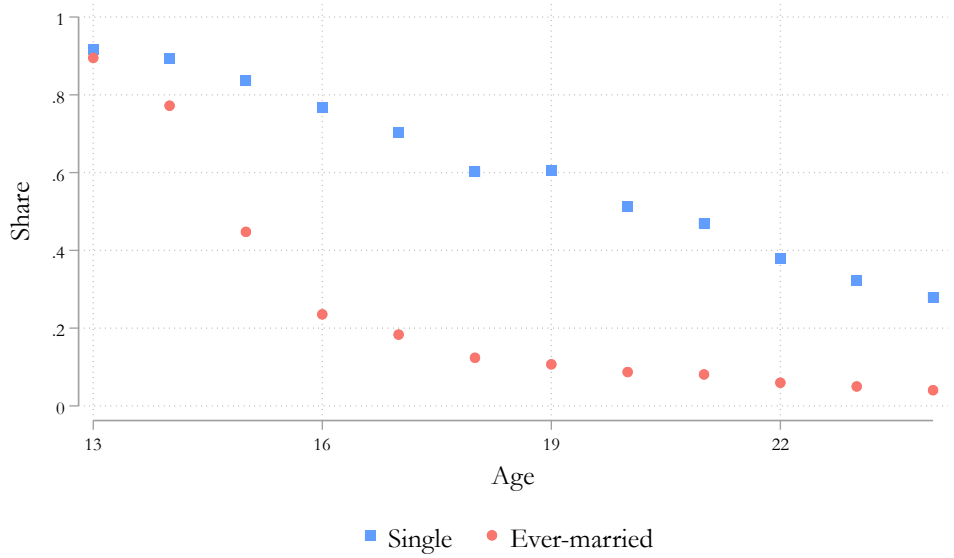


Source: Own calculations using NFHS 2019-21.
Sample: Individuals between 5 to 24 years of age

3.3 Marriage and Schooling

Marriage and schooling are strongly negatively correlated in the data. Figure 3 shows a large gap in school enrollment between married and unmarried women in India. At age 18, the gap in enrollment in high school or college is nearly 50 percentage points. This pattern is consistent with the idea that marriage substantially reduces women's ability to continue their education and motivates the empirical analysis that follows, which seeks to isolate the causal effect of marriage timing on schooling.

Figure 3: Share enrolled in school by marital status



Source: Own calculations using NFHS 2019-21.
Sample: Individuals between 13 to 24 years of age

4 Empirical findings

I now turn to estimating the causal effect of marriage age on educational attainment. Consider the following equation:

$$E_i = \beta_0 + \beta_1 A_i + \beta_2 X_i + \beta_3 Z_i^E + \epsilon_i^E, \tag{1}$$

where E_i is individual i's years of schooling and A_i is i's age at marriage. X_i is a vector of observables such as wealth and caste.

Estimating equation (1) by OLS would not identify the causal effect of age at marriage on education for several reasons. First, schooling and marriage timing are jointly determined within the household. A girl who remains in school longer may mechanically marry later, since married women in the data rarely continue their education. Second, unobserved household characteristics, such as attitudes toward women, religious beliefs, or preferences regarding daughters' schooling and marriage, may affect both educational investments and marriage timing.

To address these concerns, I adopt an instrumental variables strategy. Specifically, I exploit plausibly exogenous variation in age at marriage generated by a common social norm of arranging daughters' marriages in order of birth.

4.1 Presence of an older sister

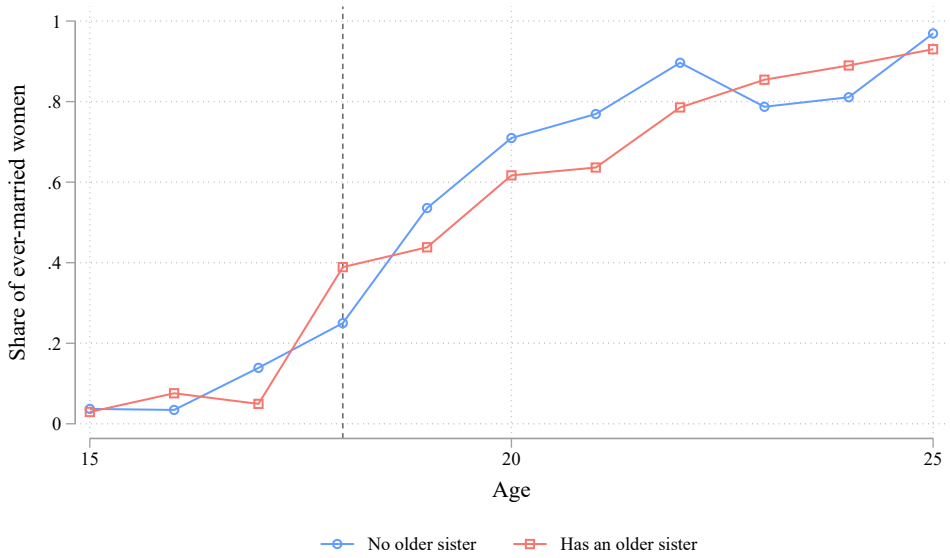
In India, a common practice is to arrange marriages of children of the same sex according to the order of birth. This implies that having an older sister can lead to delays in the younger sister's marriage due to frictions in the marriage market.

Several studies have documented the influence of this social norm on marriage timing. For instance, research in Sri Lanka found that fourth or younger order sisters were only 42 percent as likely to marry at any given age as elder sisters (Malhotra and Tsui, 1996). Similar findings have been observed in Nepal, Bangladesh, India, and Pakistan, with younger daughters leaving home (equivalent to marriage for South Asian women) roughly two years later than the median older daughter (Vogl, 2013).

Several theories may explain the persistence of this norm. First, it could be optimal for families to marry off daughters in order of birth, as men often prefer younger brides (Anukriti and Dasgupta, 2017), leaving older sisters with relatively less time to find a match. Second, adhering to this norm may reduce direct competition between siblings over potential spouses.

The importance of marriage by order of birth is apparent in India from the REDS data (Figure 4). Oldest daughters are more likely to be married than younger daughters of the same age up to age 23. The trend reverses after age 23, supporting the idea that older and younger sisters do not fundamentally differ in their likelihood to ever marry. In the data, only around 5% of younger sisters marry before their older sisters.

Figure 4: Share of married women by birth order



Source: Own calculations using REDS 2006-07.
Sample: All women between ages 15-30 with at least one sister.

Frictions in the marriage market can result in delays in a girl's marriage when she has an older sister. Thus, conditional on the number of siblings, the presence of an older sister provides exogenous variation in a woman's age at marriage. Consequently, I compare outcomes for girls without older sisters with those for girls with older sisters. Previous research by Vogl (2013) has effectively used this exogenous variation in age at marriage and found that the presence of a younger sister decreases the age of marriage and spousal quality, while the presence of an older sister leads to an increase in the age at marriage in Nepal. Similarly, recent work by Keskar et al. (2025) find similar patterns in India.

4.2 Data and sample selection issues

To estimate the impact of having an older sister on age of marriage, I restrict the analysis sample to households with at least two daughters. This allows me to specifically study the impact of being the oldest sister versus having an older sister. Households with only one daughter are not a suitable control group for households with two daughters, as the former do not face any constraints on marriage age.

An important limitation of any marriage data is that I only observe women who are already married. Women who have not married yet will not be a part of my analysis, even though these might precisely be the women affected by the presence of older sisters. I plan to address this limitation with Heckman correction and duration analysis. As a robustness check, I will also restrict the sample to households where all daughters are already married.

4.3 First Stage Results

Table 1 presents results from regressing women's age at marriage on the presence of an older sister. Column (1) presents results using the presence of an older sister in the household, column (2) presents results using the number of older sisters, and column (3) presents results using an indicator for whether the next-older sibling is female, controlling for the age gap and birth order. The results suggest that having an older sister increases a woman's age at marriage by around 0.4 years, which is around 5 months.

Table 1: First stage: Regress age at marriage on older sister instruments

Dep. var.: Age at marriage	Has an older sister	Number of older sisters	Next-older sister
	(1)	(2)	(3)
Has an older sister=1	0.675***		
	(0.12)		
Age gap with next-older sister	-0.064***		
	(0.01)		
1 older sister		0.248***	
		(0.09)	
2 older sisters		0.742***	
		(0.12)	
3 older sisters		0.962***	
		(0.16)	
4 or more older sisters		1.481***	
		(0.23)	
Next-older sister=1			0.378***
			(0.07)
Age gap			-0.066***
			(0.01)
Constant	20.000***	18.888***	19.089***
	(3.20)	(3.23)	(3.23)
N	8,145	7,836	7,836
F	31.58	16.18	25.47

* $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$

Source: Own calculations using REDS 2006.

Notes: Column (1) regresses age at marriage on an indicator for having an older sister. Column (2) regresses age at marriage on the number of older sisters. Column (3) regresses age at marriage on an indicator that the next-oldest (i.e. immediately older) sibling is a sister. All regressions control for the number of siblings, year of birth. Regressions 1 and 3 also control for age gaps. Regression 1 and 2 control for birth order.

Sample: Women between ages 20 and 40

Column 1: Women with at least one sister.

Columns 2 & 3: Women with at least one older sibling.

4.4 Two-stage least squares estimates

Table 2 presents results from a two-stage least squares regression of educational attainment on age at marriage, using the presence of an older sister as an instrument for age at marriage. The dependent variables are three education outcomes, namely, the years of schooling, tenth grade completion and high school completion.

Table 2: IV: Regress schooling outcomes on age at marriage

Dep. var.:	Years of education	Tenth grade or more	HS or more
	(1)	(2)	(3)
Age at marriage	1.710***	0.118***	0.044***
	(0.32)	(0.03)	(0.02)
Age gap with next-older sister	0.021	0.002	0.002
	(0.02)	(0.00)	(0.00)
Constant	-30.774***	-2.087***	-0.718*
	(8.72)	(0.70)	(0.42)
Mean of dep variable	4.77	0.23	0.10
N	8,145	8,145	8,145

* $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$
Source: Own calculations using REDS 2006.
Notes: Each column represents a 2SLS regression, where age at marriage is instrumented by an indicator for having an older sister. Column (1) regresses years of education on age at marriage. Column (2) regresses an indicator for 10th grade completion or more on age at marriage. Column (3) regresses an indicator for high school completion or more on age at marriage. All regressions control for the number of siblings, year of birth, birth order, and age gap with the next-oldest sister
Sample: Women between ages 20 and 40 with at least 1 sister

The results indicate that delaying marriage by a year increases years of schooling by around 1.7 years. It also increases tenth grade completion by 11 percentage points, and high school completion by 5.4 percentage points.

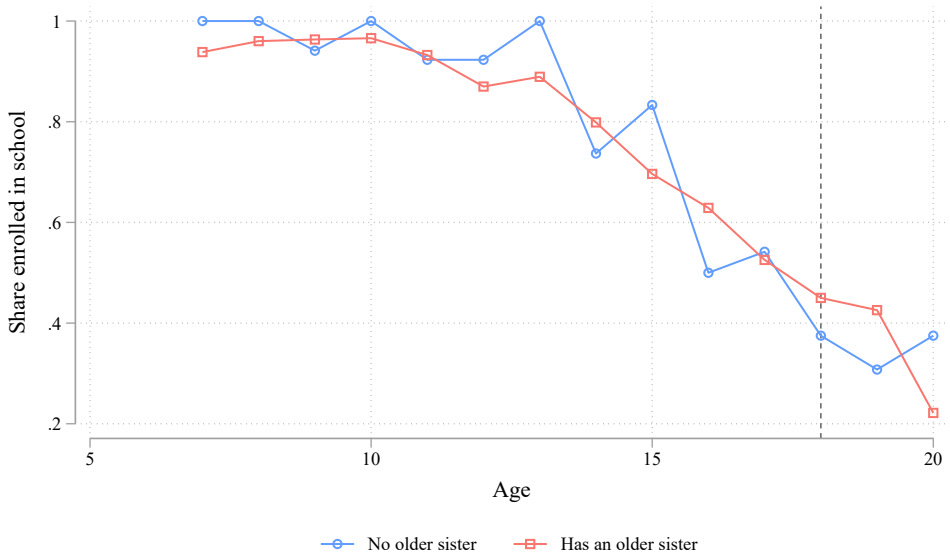
4.5 Threats to exogeneity

4.5.1 Parental investment by birth order

A threat to exogeneity involves differential investments by parents in their children's education based on birth order. In India, parents often depend on their eldest son for old-age support, leading to uneven investments in sons according to their birth order. Jayachandran and Pande (2017) find that this preference for the eldest son is reflected in physical outcomes, with the oldest sons being taller than their younger siblings. In the case of daughters, one rationale for unequal investment could be the expectation that the eldest daughter bears more responsibility for household chores and caregiving for younger siblings. Nevertheless, Figure 5 offers suggestive evidence countering this notion. The data show no statistically significant differences in school enrollment between older and younger daughters up to the age of 17. This is consistent with birth order impacting girls' education primarily when marriage becomes a consideration.

Furthermore, when I control for birth order in Column (2) of Table 1, the results are very similar to the results in Column (1).

Figure 5: Share of women in school by birth order



Source: Own calculations using REDS 2006-07.
Sample: Households with at least two daughters between ages 7-16

4.5.2 Older sibling composition

It is possible that when older sons start working, household wealth increases, leading to greater investment in the education of younger siblings. One way to account for this is to control for the composition of older siblings. When I add controls for the number of older brothers, effect sizes decrease and are less precisely estimated, but are still in the same direction.

Table 3: IV: Controlling for number of older brothers

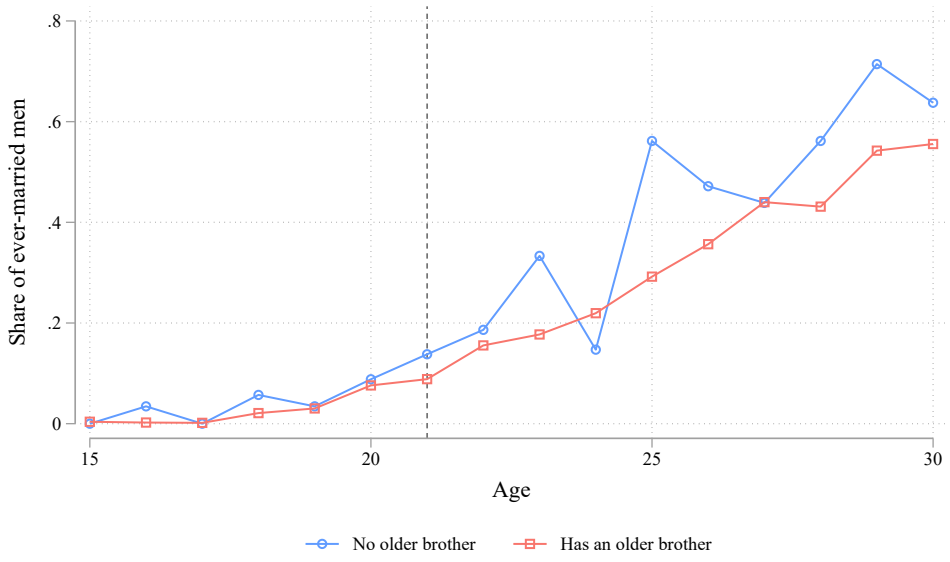
Dep. var.: Years of education	Original	Number of older brothers
	(1)	(2)
Age at marriage	1.564***	0.820*
	(0.52)	(0.43)
Age gap	0.093*	0.036
	(0.06)	(0.05)
Birth order	0.078	
	(0.09)	
Constant	-22.566**	-6.107
	(10.94)	(9.35)
N	2,880	2,880

Source: Own calculations using REDS 2006-07.
Sample: Women whose youngest sister is between 25-50 years of age, and have at least one older sibling
Notes: All regressions control for the number of children in the household and year of birth. Columns 1 is the original specification, and column 2 controls for the composition of older siblings.

4.5.3 Dowry payments for older daughters

Another concern is that having recently paid a large dowry, households might delay the younger daughter's marriage to rebuild savings.

To address this concern, I compare trends in marriage timing for sons and daughters. If saving for the dowry were the primary driver, we would not expect analogous delays for younger sons. Yet, as shown in Figure 6, men with an older brother are less likely to be married at the same age than men without an older brother.



Source: Own calculations using REDS 2006-07.
Sample: Households with at least 2 sons between ages 15-30

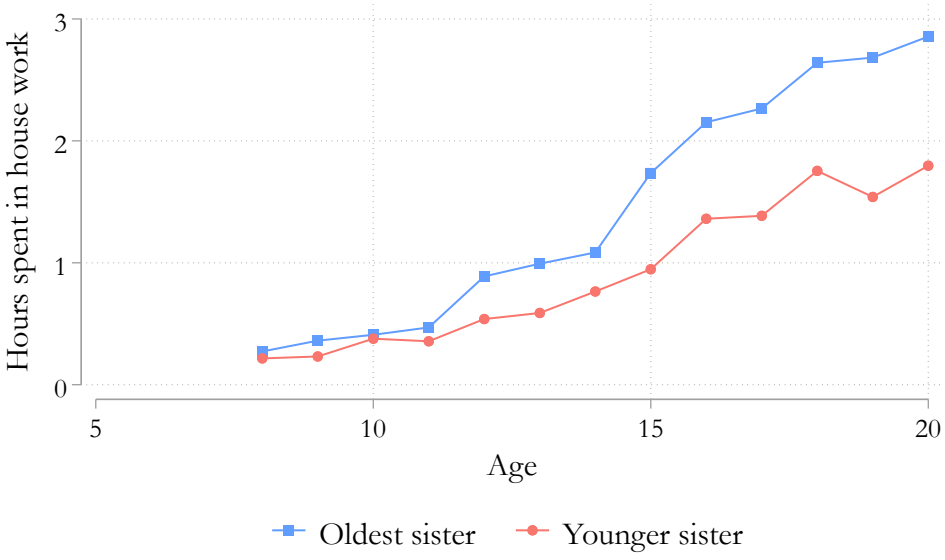
Figure 6: Share of married men by birth order

4.5.4 Home production time for older daughters

An important threat to exogeneity is that older sisters may supply more home production than younger sisters, thereby relaxing time constraints for the younger sister and raising her educational investment. In this scenario, the older-sister instrument influences schooling through home production rather than (or in addition to) marriage timing, violating the exclusion restriction.

Figure 7, based on the 2019 India Time Use Survey, shows that older daughters spend about one hour more in home production than younger daughters. To mitigate this concern, I will implement the following robustness checks:

- Exclude oldest daughters from the analysis
- Restrict the sample to households with a co-resident grandmother or sister-in-law—individuals who plausibly substitute for older daughters' home production.



Source: Own calculations using India Time Use Survey 2019
Sample: Households with two resident daughters between ages 5-20

Figure 7: Hours spent in home production by birth order

5 Model

5.1 Overview

The empirical results indicate a strong causal relationship between the age at marriage and educational attainment. I now develop an equilibrium framework to achieve two broad goals. First, the model formalizes the trade-offs faced by women when they (or their families) delay marriage to acquire more education. Second, it allows me to study the implications of increasing the minimum marriage age from 18 to 21.

The model is designed to capture the central trade-off between education and marriage timing. Delaying marriage to acquire more education can be beneficial because women's education is positively valued in the marriage market and may therefore improve match quality. However, delay is also costly along four dimensions. First, schooling requires a direct monetary investment. Second, because a woman's age is negatively valued in the marriage market, postponing marriage may worsen the quality of available matches. Third, as a woman ages, the pool of unmarried men older than her becomes smaller. Fourth, in households with multiple daughters, delaying marriage for one daughter may delay marriage for her younger sisters when search occurs in order of birth.

To capture these mechanisms, I model households with two daughters who make marriage and schooling choices in a directed search marriage market. The two-daughter structure

is important because it allows the model to capture within-household dynamic trade-offs: because daughters must marry in order of birth, delaying marriage for the older daughter affects the timing and opportunities available to the younger daughter.

5.2 Environment

Time is discrete. In each period, a new cohort of young men enters the marriage market, together with a new cohort of households whose first-born daughter is approaching marriageable age. Households are heterogeneous in income: in each period, an equal mass of low-income and high-income households enters the model. Each household has two daughters, a first-born daughter and a last-born daughter, with an age gap of $g \in \{0, 1, \ldots, \bar{g}\}$. The household enters the model when the first-born daughter turns 15, which I denote as period $t = 1$, and lives for $T = 15 + \bar{g}$ periods. In each period, the household makes marriage and schooling choices for its daughters.

A daughter is characterized by age $a \in \{15, 16, \ldots, 24\}$ and years of education $e \in \{10, 11, ..., 19\}$. Schooling is an annual binary decision for each unmarried daughter, and remaining in school increases her years of schooling. Consistent with the strong empirical pattern that married women do not continue schooling [4], I assume that schooling choices are available only before marriage.

Men are active in the marriage market between ages 18 and 27. In each period, every unmarried man participates in the marriage market. Men are characterized by age and education, where education is fixed and exogenous. If a man does not marry in the current period, he remains active in the marriage market in the following period and ages by one year.

Marriage takes place in a directed-search environment. In each period, households post a menu of dowries for the daughter who is currently eligible to marry. Men observe all posted dowries and decide which household to approach with a proposal. Households then choose among the proposals they receive, if any, and pay the corresponding dowry upon marriage. After the marriage stage is completed, households make schooling decisions for any daughters who remain unmarried.

I assume that households participate in the marriage market in every period until both daughters are married or have aged out, and men participate until they marry or age out. Given the high marriage rates observed in India, this assumption is reasonable. Importantly, this does not prevent agents from effectively delaying marriage: a household that prefers to keep a daughter in school can post a sufficiently unattractive dowry, making proposals

[4] Only around 15% of married women are enrolled in school at age 18, relative to 60% of single women (NFHS 2019).

unlikely. Similarly, a man who prefers to delay marriage can apply to a highly competitive household, reducing the probability that his proposal is accepted.

5.2.1 Types

A household's state is $\boldsymbol{h} = \{m^f, m^\ell, s^f, s^\ell, a^f, a^\ell, y, t\}$. For each daughter $b \in \{f, \ell\}$, where f denotes the first-born and ℓ the last-born, m^b is her marital status, s^b her years of schooling, and a^b her age. y is household income and t is the period within the household life cycle. Let $\nu_{\boldsymbol{h}}$ denote the mass of households of type $\boldsymbol{h}$. In each period, a fixed mass of households with two unmarried daughters, where the oldest has just turned 15, enters the model. The masses of households in each subsequent state evolve endogenously in the model, depending on the household's match outcomes and schooling choices.

A man is characterized by his type $j \in \mathcal{J}$, where j denotes his age and education. Since men's education is fixed and exogenous, the type space $\mathcal{J}$ is the set of male ages and education levels. Let μ_j denote the mass of unmarried men of type j. In each period, a fixed mass of men of age 18 enter the model with either low education or high education. The masses of men of each subsequent age then evolve endogenously, depending on men's match outcomes.

5.2.2 Preferences

Households have log utility over consumption, $u(c) = \log(c)$. Upon matching daughter b with a man of type j, the household receives a match payoff $\tau_{j,\boldsymbol{h}}$ that depends on the man's age and education.

Men receive utility from the dowry as well as an additional match payoff $\gamma_{j,(a^b,s^b)}$ that depends on the woman's age, education and the age gap between the couple. Both sides discount the future at rate β.

Households in the same state are ex-ante identical, and post identical dowries. After receiving proposals, however, each household draws independent Type-I extreme value shocks over the arriving male types, so that households in the same state may accept different proposals. Households also draw independent shocks over schooling options following the matching stage, so that households of the same type may make different schooling decisions for their unmarried daughters. On the men's side, after observing posted dowries, each man draws Type-I extreme value shocks over *individual* households, so that men of the same type may propose to different households.

5.2.3 Costs and Income

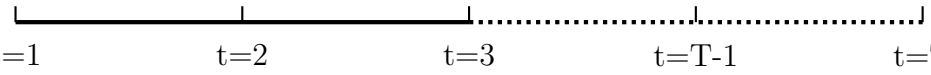
Schooling costs ϕ per period per daughter. Each unmarried daughter imposes a per-period consumption cost of κ on the household. Household income y is fixed and exogenous.

5.2.4 Birth-order search restriction

Following the birth-order patterns observed in the data, households face the restriction that daughters must marry in order of birth. Accordingly, households must search for the first-born daughter first, and only after she marries or ages out of the marriage market can the household begin searching for the younger daughter. Thus, households can search for at most one daughter in any given period. In the rest of the model, I will refer to the oldest unmarried daughter in the household as the "eligible daughter".

This feature is central to the model. Because daughters must marry in sequence, delaying marriage for the older daughter may be costly not only because it affects her own marriage prospects, but also because it delays the younger daughter's entry into the marriage market.

The figure below illustrates this constraint for two daughters with an age gap of 1 year. In period 1, when the older daughter turns 16, she is the only daughter old enough to search. If she remains unmarried, the household must continue to search for her in subsequent periods until she marries or ages out. Once she marries, the household switches to the younger daughter.



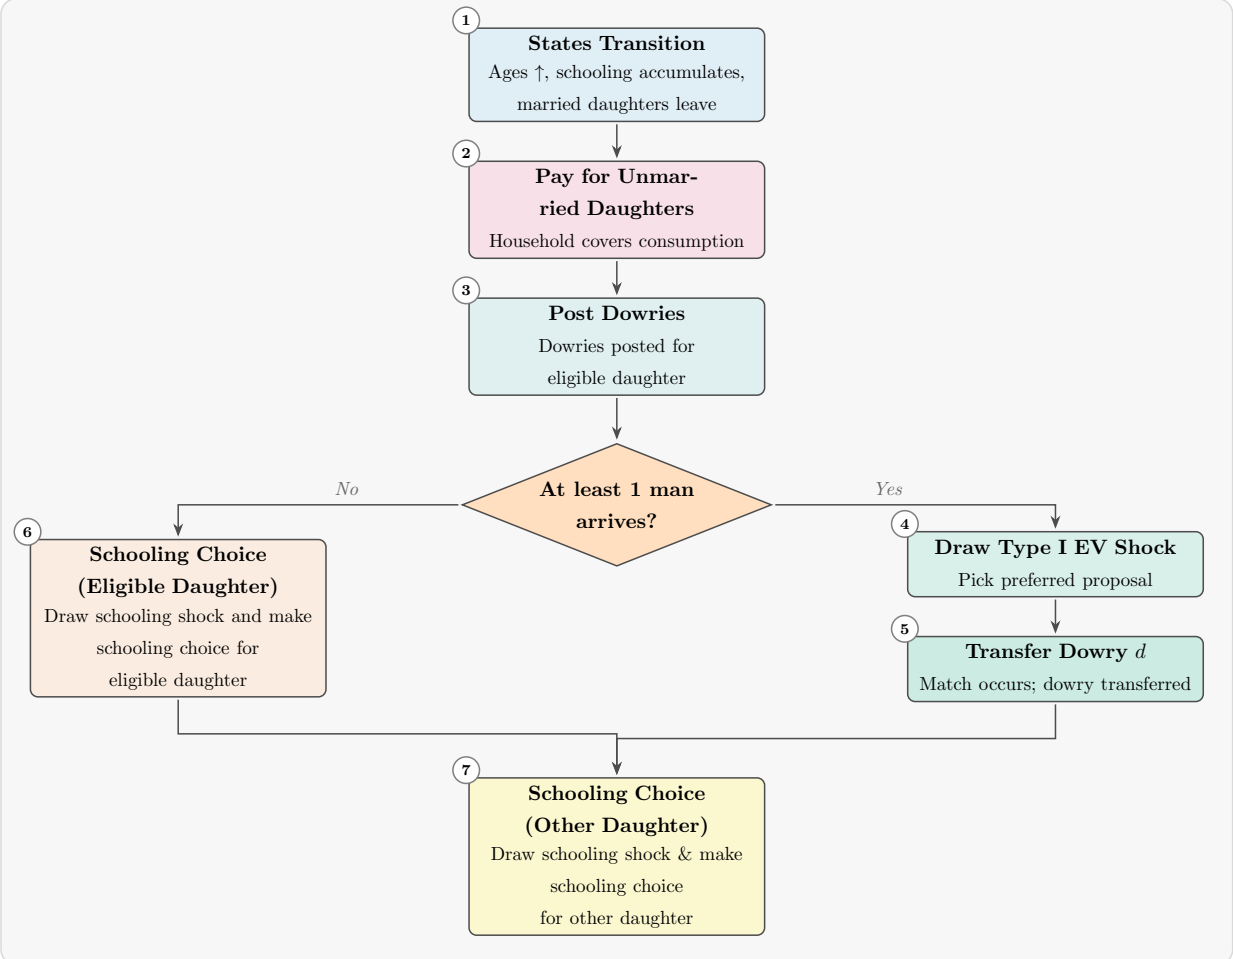
	t=1	t=2	t=3	t=T-1	t=T
Older	16	17	18 ·····	30	
Younger		16	17 ·····	29	30

5.3 Timing

The figure below summarizes the timing within a period. At the beginning of the period, any daughter who married in the previous period leaves the household, and the household state is updated accordingly. Unmarried daughters age forward, and previously chosen schooling determines current education status. The household also pays the cost for supporting its unmarried daughters.

A period then has three broad stages: dowry posting, matching and schooling choice. First, the household posts a menu of dowries for each type of man for the daughter who is currently eligible to marry. Second, men observe posted dowries and draw idiosyncratic Type I extreme value shocks over households. Then, each man chooses a household to bring a proposal to. The household observes the set of proposals it receives. If it receives at least one proposal, it draws idiosyncratic Type I extreme value shocks over the proposing male types and chooses the proposal that yields the highest value. The match is then formed and the household transfers a dowry to the chosen suitor. If the household receives no proposals, no marriage occurs in that period. After the matching stage is completed, the household makes a binary schooling decision for each unmarried daughter remaining in the household. For each such daughter, the household draws an idiosyncratic shock over the schooling option, chooses whether she continues in school and pays the corresponding schooling cost if applicable.

Period Timing



In the following sections, I will first describe the marriage market, then describe the household's schooling decision and then present the household's Bellman equation.

5.4 Marriage market

I will now describe the 4 pieces of the marriage market: the household's dowry posting decision, men's problem, the arrival rates of men's proposals and the household's acceptance decision.

5.4.1 Dowry posting

In each period, households post a menu of dowries for the *eligible* daughter, that is, the oldest unmarried daughter currently permitted to search. Dowries can depend on the age and education of the man, but not on his individual identity. Let $d_{j,\boldsymbol{h}}$ denote the dowry offered by a household of type $\boldsymbol{h}$ to a man of type j.

Men observe the dowries posted by all households and choose which household to approach. Each man can bring a proposal to only one household in a given period. A household that receives at least one proposal chooses one suitor and, if a match is formed, pays the corresponding dowry.

5.4.2 Men's proposal decision

In each period, a man of type j chooses one household to approach with a proposal until he either marries or ages out of the market. Let $P_{j,(\boldsymbol{h},i)}$ denote the probability that a proposal from a particular type-j man to a household i of $\boldsymbol{h}$ is accepted. Let $\gamma_{j,(a^b,s^b)}$ denote the utility he obtains from marrying the eligible daughter, which depends on her age and schooling. Then the deterministic part of his value from choosing a household i of type h is:

$$Z_{j,(\boldsymbol{h},i)} = P_{j,(\boldsymbol{h},i)} \left(\gamma_{j,(a^b,s^b)} + z(c_{j,(\boldsymbol{h},i)})\right) + \left(1 - P_{j,(\boldsymbol{h},i)}\right) \beta Z(j') \tag{2}$$

s.t.

$$c_{j,(\boldsymbol{h},i)} = d_{j,(\boldsymbol{h},i)}$$

where $Z(j')$ denotes the value of searching in the next period, one year older, with the same education level, and $z(c)$ denotes his utility from consumption.

The acceptance probability $P_{j,(\boldsymbol{h},i)}$ depends on two factors. First, it depends on which other male types propose to household $(\boldsymbol{h}, i)$. Second, conditional on the household selecting type j, it depends on how many other type-j suitors have proposed to that household (see Section 5.4.3 for the exact formulation). Men take this probability as given when making their proposal decision.

After observing all posted dowries, a particular man of type j draws i.i.d. Type I extreme value shocks $\omega_{(\boldsymbol{h},i)}$ over households and applies to household $(\boldsymbol{h}, i)$ that maximizes Z_j:

$$Z_j = \left[\max_{(\boldsymbol{h},i))}\{Z_{j,(\boldsymbol{h},i)} + \omega_{\boldsymbol{h},i}\}\right] \tag{3}$$

Men also have an outside option of not proposing to any household. Let $Z_{j,0}$ denote the value of this option, under which the man remains single in the current period, ages by one period, and re-enters the marriage market next period.

5.4.3 Proposal arrival rates

From the logit structure of men's idiosyncratic shocks, the probability that an individual man of type j proposes to household i of type $\boldsymbol{h}$ is

$$\pi_{j,(\boldsymbol{h},i)} = \frac{\exp\big(Z_{j,(\boldsymbol{h},i)}/\sigma_j\big)}{\exp(Z_{j,0}/\sigma_j) + \sum_{(\boldsymbol{h}',i')\in\mathcal{H}} \exp\big(Z_{j,(\boldsymbol{h}',i')}/\sigma_j\big)}. \tag{4}$$

Here, $\pi_{j,(\boldsymbol{h},i)}$ is an *individual proposal probability*: it gives the probability that a single type-j man chooses household $(\boldsymbol{h}, i)$.

Given a mass μ_j of unmarried men of type j, the expected number of proposals that household $(\boldsymbol{h}, i)$ receives from type-j men is

$$q_{j,(\boldsymbol{h},i)} = \mu_j \pi_{j,(\boldsymbol{h},i)}. \tag{5}$$

Thus, $q_{j,(\boldsymbol{h},i)}$ is the *expected queue length* from male type j at household $(\boldsymbol{h}, i)$.

Assuming a continuum of men and independent proposal decisions, the realized number of type-j proposals received by household $(\boldsymbol{h}, i)$ is Poisson with mean $\delta q_{j,(\boldsymbol{h},i)}$ [5] where δ is a scaling parameter governing proposal arrival rates. Therefore, the probability that household $(\boldsymbol{h}, i)$ receives at least one proposal from a type-j man is [6]:

$$Q_{j,(\boldsymbol{h},i)} = 1 - \exp\big(-\delta q_{j,(\boldsymbol{h},i)}\big)\,. \tag{6}$$

In words, men's logit choice over households determines individual proposal probabilities $\pi_{j,(\boldsymbol{h},i)}$, which aggregate into expected queue lengths $q_{j,(\boldsymbol{h},i)}$, which in turn determine arrival probabilities $Q_{j,(\boldsymbol{h},i)}$.

[5] Let $q_{j,h} \equiv \mu_j \cdot \pi_{j,h}$ denote the queue length: the expected number of type-j men applying to a household of type h. Then $A_{j,h} \sim \text{Binomial}(\mu_j, \pi_{j,h})$, and in the limit $\mu_j \to \infty$, $\pi_{j,h} \to 0$ with $\mu_j \cdot \pi_{j,h} = q_{j,h}$ fixed, $A_{j,h}$ is approximately Poisson$(q_{j,h})$. The parameter δ scales the meeting technology, so the realized number of proposals received is Poisson$(\delta q_{j,h})$.

[6] This follows from the Poisson pmf: $\Pr(A_{j,h} = z) = \exp(-\delta q_{j,h})\frac{(\delta q_{j,h})^z}{z!}$. Setting $z = 0$ gives $\Pr(A_{j,h} = 0) = \exp(-\delta q_{j,h})$, so the probability that the household receives at least one proposal from a type-j man is $1 - \exp(-\delta q_{j,h})$.

5.4.4 Household acceptance decision

A household fails to match its eligible daughter only if it receives no proposals. If the household receives at least one proposal, it selects one of them. If proposals from multiple male types arrive, the household draws a Type I extreme value shock associated with each type and chooses the type that maximizes its value.

Denote the type chosen for the eligible daughter as j^b, where $b \in \{f, \ell\}$ denotes the birth order of the eligible daughter, at the posted dowry d_{j^b}. Let $\tau_{j^b,\boldsymbol{h}}$ denote the utility the household gets from accepting a proposal from a type j^b man. This term captures the household's gains from the match and depends on the man's age and education.

Then the value to a household in state $\boldsymbol{h}_t$ from accepting a proposal from a man of type j^b for its eligible daughter b can be written as:

$$W^a(\boldsymbol{h}_t, d_{j^b}, j^b) = u\big(c^a\big) + \tau_{j^b} + \beta \mathbb{E} W^s(\boldsymbol{h}_t, d_{j^b}, j^b)\Big].$$

where $W^s(\boldsymbol{h}_t, d_{j^b}, j^b)$ denotes the household's value from the optimal schooling choice (described in Section 5.5).

Then the probability that the proposal by a particular man k of type j is accepted by a household in state $\boldsymbol{h}$ (suppressing individual subscripts for brevity):

$$P_j(\boldsymbol{h}) = \underbrace{\frac{1 - e^{-\delta q_{j,\boldsymbol{h}}}}{q_{j,\boldsymbol{h}}}}_{\text{within-type selection}} \sum_{\substack{\mathcal{R} \subseteq \mathcal{J} \\ j \in \mathcal{R}}} \underbrace{\left[\prod_{k \in \mathcal{R} \setminus \{j\}} (1 - e^{-\delta q_{k,\boldsymbol{h}}}) \prod_{k \notin \mathcal{R}} e^{-\delta q_{k,\boldsymbol{h}}} \right]}_{\Pr(\text{arrival set} = \mathcal{R} \mid j \text{ present})} \cdot \underbrace{\frac{e^{W^a(\boldsymbol{h}, d_j, j)/\sigma}}{\sum_{x \in \mathcal{R}} e^{W^a(\boldsymbol{h}, d_x, x)/\sigma}}}_{\text{choice among arriving types}} . \tag{7}$$

The first term is the probability that a particular man of type j is the selected suitor within his type (see Section B for a detailed derivation). The summation integrates over all possible subsets of other male types that may also arrive. Let R denote the set of male types from which the household receives at least one proposal in period t. Conditional on the set of available male types, the household chooses among them according to a logit rule. When no other type arrives, the conditional choice probability is one.

5.4.5 Household value in no-match state

A household fails to match its eligible daughter only if it receives no proposals. If the household receives no proposals, it moves on to the schooling stage and makes schooling decisions

for each unmarried daughter. Let $W^n(\boldsymbol{h_t}, j^b = 0)$ denote the household's value in the case that it receives no proposals. Then:

$$W^n(\boldsymbol{h}_t, j^b = 0) = u\big(c^n\big) + \beta \mathbb{E} W^s(\boldsymbol{h}_t, j^b = 0)\Big].$$

where c^n denotes the household's consumption in this state.

5.5 Schooling decision

After the marriage stage, the household chooses schooling decisions for both daughters, (e^f, e^ℓ). This formulation allows the schooling problem to be written uniformly across all household states. However, schooling is feasible only for daughters who remain unmarried after the matching stage. Thus, if the eligible daughter b matches with a man of type $j^b \in \mathcal{J}$, her schooling choice is constrained to satisfy $e^b = 0$. More generally, any married daughter has schooling choice fixed at zero. Schooling is an annual binary decision for each unmarried daughter.

Let $V(\boldsymbol{h}_t)$ denote the household's value in state h_t. Then, conditional on match outcome j^b for the eligible daughter b, the household's schooling problem is

$$W^s(\boldsymbol{h_t}; d_{j^b}, j^b) = \max_{e^f, e^\ell} \left\{ \beta V(\boldsymbol{h}_{t+1} \mid e^f, e^\ell) + \varepsilon(e^f, e^\ell) \right\}, \tag{8}$$

s.t.

$$\begin{aligned} e^b = 0 \quad \text{if} \quad j^b > 0; \quad b \in \{f, \ell\} \\ e^b = 0 \quad \text{if} \quad m^b = 1; \quad b \in \{f, \ell\} \end{aligned}$$

where ε is an i.i.d. Type I extreme value shock. This formulation yields the standard logit representation of the schooling choice. h' denotes the household state in the next period, where unmarried daughters age by a year and any daughter that married this period exits the household.

Thus, schooling is chosen after the matching stage and by construction, only for daughters who remain unmarried.

5.6 Household problem

The household chooses the vector of posted dowries for the eligible daughter, $\mathbf{d} \equiv d_{j\, j \in \mathcal{J}}$, to maximize expected utility, taking as given the mapping from posted dowries to men's proposal behavior. Let R denote the set of male types from which the household receives at

least one proposal in period t. Because the household may receive proposals from multiple male types, R can be any subset of $\mathcal{J}$: for example, the household may receive proposals only from type 1, so that $R = \{1\}$, or from types 1 and 2, so that $R = \{1, 2\}$. More generally, R indexes the subset of male types represented among the proposals received by the household.

$$V(\boldsymbol{h}) = \max_{\mathbf{d}} \Bigg\{ \underbrace{\sum_{\substack{\mathcal{R}\subseteq\mathcal{J} \\ \mathcal{R}\neq\varnothing}} \Pr(\mathcal{R} \mid \mathbf{d})\, \mathbb{E}_{\eta}\left[\max_{j^f \in \mathcal{R}} \left\{ W^a(\boldsymbol{h}, d_{j^f}, j^f) + \eta_{j^f} \right\} \right]}_{\text{match}} + \underbrace{\Pr(\mathcal{R} = \varnothing \mid \mathbf{d})\, W^n(\boldsymbol{h}, j^f = 0)}_{\text{no match}} \Bigg\}.$$

s.t.

$$c^a = y - d_{j^f} - \phi e^f - \phi e^{\ell} - n\kappa \tag{9}$$

$$c^n = y - \phi e^f - \phi e^{\ell} - n\kappa \tag{10}$$

The first term captures the expected utility of the household in the case that it receives at least one proposal from each type in the set R. In that case, the household draws a Type-I Extreme Value shock associated with each type of proposal and accepts the proposal from the type of man that gives it maximum value. If multiple men of the same type propose, as mentioned previously, the household is indifferent between men of the same type and hence chooses a proposal randomly for its optimal type.

The second term captures the expected utility of the household in the case that it receives no proposals.

The budget constraint requires that the household's consumption must be less than its income minus the other costs, including the dowry (if there is a match), the cost of education ϕ, if either daughter goes to school, and the total cost of supporting unmarried daughters, where κ denotes the per-capita cost of supporting an unmarried daughter. n denotes the number of unmarried daughters in the household at the start of the period. The assumption is that the household pays for unmarried daughters in this period, and any daughter who marries in this period leaves the household at the beginning of the next period.

The household's choice of dowries affects queue lengths through men's logit demand. Queue lengths in turn determine proposal arrival probabilities, and through the household's acceptance rule, the probability of matching with each male type.

5.7 Stationary equilibrium

A stationary competitive search equilibrium consists of household value functions $\{V(s)\}_{s\in\mathcal{H}}$, posted dowries $\{d_{m,s}\}_{m,s}$, men's application probabilities $\{\pi_{m,s}\}_{m,s}$, queue lengths $\{q_{m,s}\}_{m,s}$, acceptance probabilities $\{P_{m,s}\}_{m,s}$, and a stationary distribution of male and household types $\{\mu_j, \nu_{\boldsymbol{h}}\}$ such that:

1. For each household type s, posted dowries maximize $V(s)$.
2. For each male type m, application probabilities satisfy the logit rule implied by $\{J_{m,s}\}_{s\in\mathcal{H}}$.
3. Queue lengths satisfy
$$q_{m,s} = M_m \pi_{m,s}$$
for all m and s.
4. Proposal arrival probabilities and acceptance probabilities are consistent with the Poisson arrival process and the household's acceptance rule.
5. The distribution of male and household types is stationary under the transition rules induced by equilibrium matching, aging, and schooling.

In practice, I solve for equilibrium as a fixed point in household policies, men's application behavior, and the stationary distribution of types.

6 Estimation

I estimate the model using indirect inference. The key parameters to estimate are the returns to age and education in the marriage market (entering through both household and men's match utilities), the schooling cost ϕ, the search friction parameter δ, and the baseline utility levels for both sides.

6.1 Mapping the model to data

To bring the model to the data, I discretize the state space as follows. Women's ages are grouped into two bins: young (15–19) and old (20–24). The education choice corresponds to high school completion. Given the two-age-bin structure, the schooling decision is made only during the young period. Men's ages are similarly grouped into young (18–22) and old (23–27), and men's education is classified as low (less than high school) or high (high school or more). The ratio of low-education to high-education men is set to 3:1, based on

the distribution in the data. Household income takes two values: Rs. 26,000 and Rs. 180,000, each multiplied by 4 to scale to model periods.

6.2 Functional forms

The household has log utility over consumption, $u(c) = \log(c)$. The household's match utility from a man of type j is

$$\tau_j = \alpha_0 + \alpha_1 \mathbb{1}\mathbb{1}[a^m = \text{old}] + \alpha_2 \mathbb{1}\mathbb{1}[s^m = \text{high}],$$

where a^m denotes men's age and s^m denotes men's education. Thus α_1 captures the household's gains from the man's age and α_2 captures the household's gains from the man's education.

The man's match utility from marrying a daughter with characteristics (a^b, s^b) is

$$\gamma_{(a^b,s^b)} = \psi_0 + \psi_1 \mathbb{1}[a^b = \text{young}] + \psi_2 \mathbb{1}[s^b = \text{high}].$$

Thus ψ_1 captures men's utility from the woman's age and ψ_2 captures their utility from her education.

6.3 Identification

The IV estimates from the design-based analysis discipline the marriage-market returns to education. The design-based results show that when faced with a constraint on marriage timing, households choose to educate their daughters rather than pursue other options such as home production, suggesting positive perceived returns from education in the marriage market. The first-stage of the IV helps me identify the search friction parameter. Cross-sectional patterns of matches by age, education, and income identify the remaining parameters, including the returns to age, and the schooling cost.

6.4 Estimated parameters

Table 4 reports the estimated parameters.

Men place a large negative value on the wife's age ($\psi_1 = -1.79$) and a large positive value on her schooling ($\psi_2 = 2.5$). This is consistent with the central trade-off in the model: delaying marriage to acquire education is costly because men dislike older brides, but is beneficial because men value educated wives. On the household side, the estimated return to the husband's age is zero ($\alpha_1 = 0$), while the return to the husband's education is positive

Table 4: Estimated Parameters

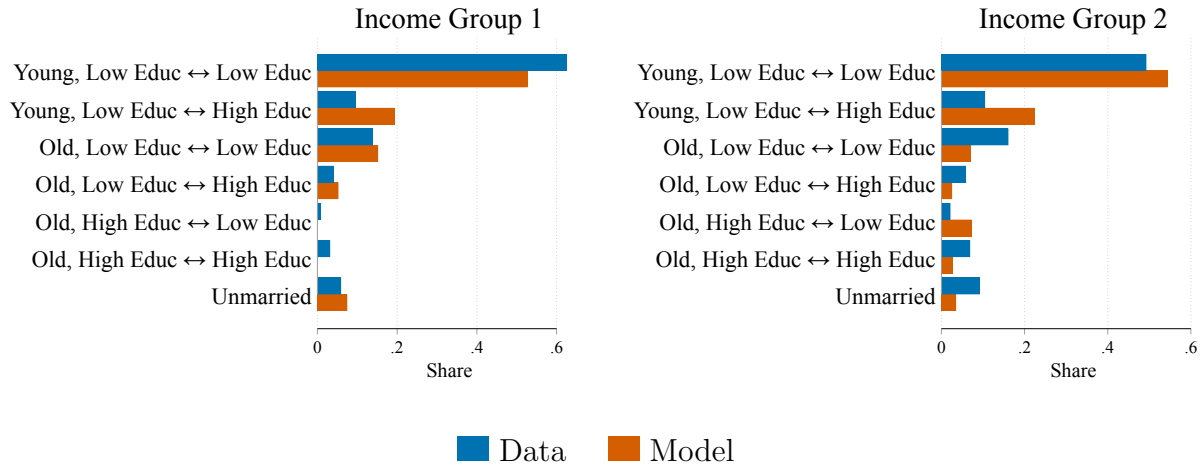
Parameter	Estimate
Men's utility, wife age (ψ_1)	-1.79
Men's utility, wife schooling (ψ_2)	2.5
HH utility, husband age (α_1)	0
HH utility, husband schooling (α_2)	0.5
Search friction (δ)	1.45
Schooling cost (ϕ, in 1000s)	Rs. 144.49
Men's baseline utility (ψ_0)	2.91
HH baseline utility (α_0)	0.93

($\alpha_2 = 0.5$), indicating that households value educated grooms but are indifferent over the groom's age. The estimated schooling cost is Rs. 144,490, and the search friction parameter is $\delta = 1.45$.

6.5 Model fit

Figures 8–9 assess the model's ability to match the targeted moments. Figure 8 compares the model-predicted and data match type shares by income group. Figure 9 compares the model-predicted and data moments for the IV relationship and age-pair distributions.

Figure 8: Match type shares by income



As a validation exercise, Figure 10 compares the model-predicted average dowries to their data counterparts. Dowries are not targeted in estimation, so this provides an out-of-sample check on the model's fit.

Figure 9: IV relationship and age-pair distributions

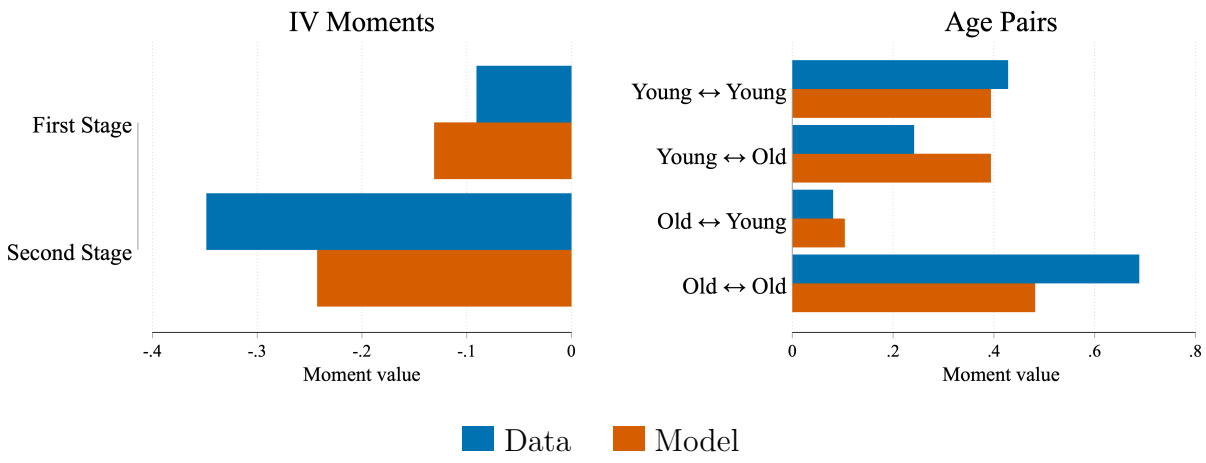
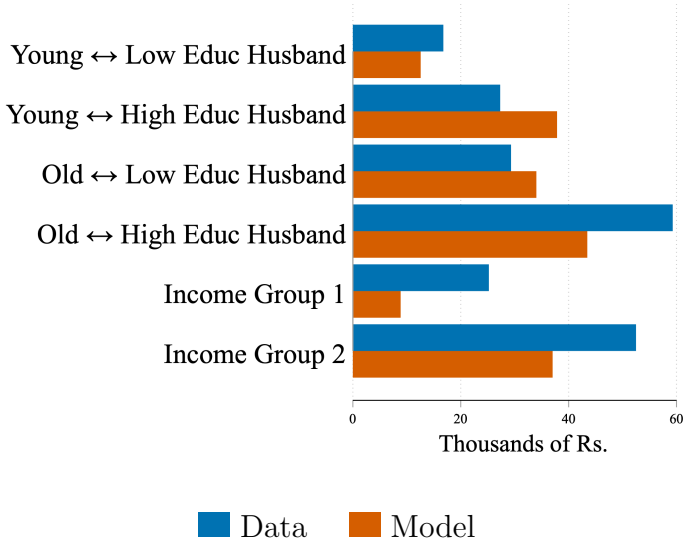


Figure 10: Average dowries (non-targeted)



7 Counterfactual analysis: Minimum marriage age policy

The policy experiment of interest is an increase in the minimum legal age of marriage from 18 to 21. I model this reform as a change in the set of daughters eligible to participate in the marriage market: under the reform, daughters can enter the marriage market only in the "old" stage (ages 20 - 24). In the future, I will also allow for imperfect enforcement by introducing a penalty for violating the law, and study how the effects of the reform vary with the degree of enforcement.

Figure 11 shows the effect of the reform on women's educational attainment. The share of women completing high school increases from 5.7% in the baseline to 22.8% under the reform, a 17.1 percentage point increase. By preventing marriage before age 21, the reform effectively removes the opportunity cost of schooling during the young period: since daughters cannot marry while young, households choose to invest in their education instead.

Figure 11: Effect of minimum marriage age reform on education

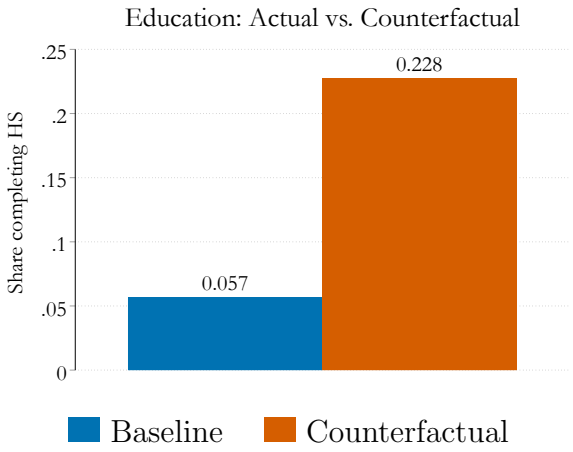
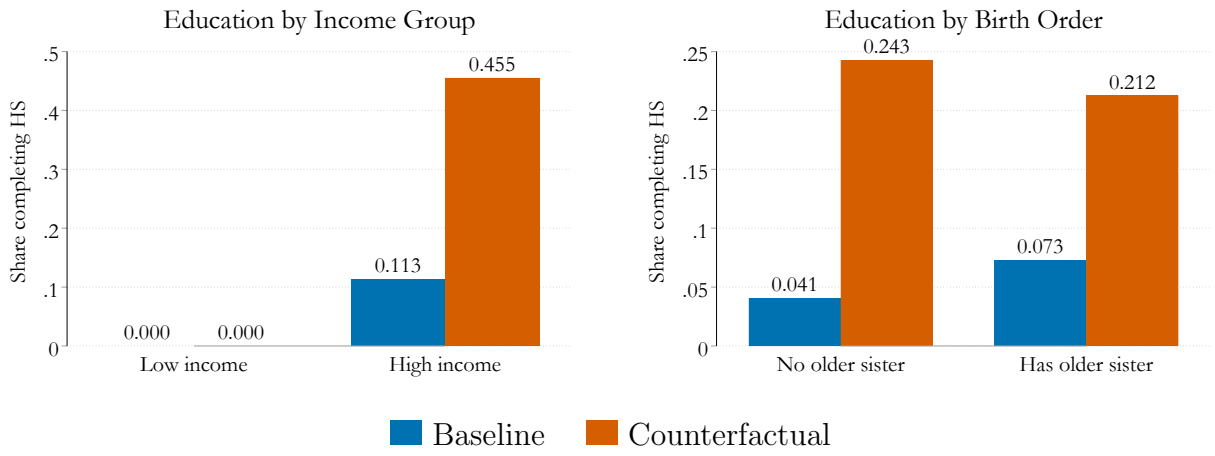


Figure 12 examines heterogeneity in the reform's effects. The left panel shows results by household income. Interestingly, this reveals that the effect is driven entirely by high-income households: their high school completion rate rises from 11.3% to 45.5%, an increase of 34.2 percentage points. Low-income households see no change, with completion rates remaining at zero in both the baseline and the counterfactual. This is consistent with the model's budget constraint: even when the reform delays marriage and frees up the young period for schooling, low-income households cannot afford the direct cost of education.

The right panel shows results by birth order. Both first-born and last-born daughters benefit from the reform, but the effect is larger for first-borns. Daughters without an older sister see their completion rate rise from 4.1% to 24.3% (a 20.2 percentage point increase),

Figure 12: Effect of minimum marriage age reform on education: heterogeneity



while daughters with an older sister see an increase from 7.3% to 21.2% (13.9 percentage points).

8 Conclusion

This paper studies how constraints on the timing of marriage shape women's educational attainment in rural India. Using the presence of an older sister as an instrument for age at marriage, I find that delaying marriage by one year raises women's schooling by 1.7 years and the probability of completing high school by 4.4 percentage points. To evaluate the proposed increase in the minimum legal age of marriage from 18 to 21, I develop and estimate an equilibrium model in which households make decisions about dowries, match acceptance, and schooling over their daughters' life cycles. Under the reform, high school completion rises from 5.7 to 22.8 percent, though the gains are concentrated among higher-income households for whom the direct cost of schooling does not bind. Taken together, the empirical and structural results indicate that marriage-market constraints are a first-order determinant of women's education in this setting, and that raising the minimum marriage age can produce large educational gains — gains that are, however, unevenly distributed across the wealth distribution. In ongoing work, I extend the model to allow for imperfect enforcement, compare the reform to a compulsory schooling alternative, and quantify the welfare consequences of each policy.

References

Anderson, S. (2007). Why the marriage squeeze cannot cause dowry inflation. *Journal of Economic Theory*, 137(1):140–152.

Andrew, A. and Adams, A. (2022). Revealed beliefs and the marriage market return to education. Technical report, Institute for Fiscal Studies.

Anukriti, S. and Dasgupta, S. (2017). Marriage markets in developing countries. *The Oxford handbook of women and the economy*, pages 97–120.

Ashraf, N., Bau, N., Nunn, N., and Voena, A. (2020). Bride price and female education. *Journal of Political Economy*, 128(2):591–641.

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

Chiappori, P.-A., Dias, M. C., and Meghir, C. (2018). The marriage market, labor supply, and education choice. *Journal of Political Economy*, 126(S1):S26–S72.

Chiappori, P.-A., Iyigun, M., and Weiss, Y. (2009). Investment in schooling and the marriage market. *American Economic Review*, 99(5):1689–1713.

Ciscato, E. (2021). The changing wage distribution and the decline of marriage. *Unpublished manuscript*.

Corno, L., Hildebrandt, N., and Voena, A. (2016). Weather shocks, age of marriage and the direction of marriage payments. Technical report, Working Paper.

Field, E. and Ambrus, A. (2008). Early marriage, age of menarche, and female schooling attainment in Bangladesh. *Journal of Political Economy*, 116(5):881–930.

Flinn, C., Gemici, A., and Laufer, S. (2017). Search, matching and training. *Review of Economic Dynamics*, 25:260–297.

Jayachandran, S. and Pande, R. (2017). Why are Indian children so short? The role of birth order and son preference. *American Economic Review*, 107(9):2600–2629.

Keskar, A., Mookerjee, S., and Guha, A. (2025). Role of siblings in marriage and dowry practices: Evidence from india. *Available at SSRN 5227861*.

Malhotra, A. and Tsui, A. O. (1996). Marriage timing in Sri Lanka: The role of modern norms and ideas. *Journal of Marriage and the Family*, pages 476–490.

Moen, E. R. and Rosén, Å. (2004). Does poaching distort training? *The Review of Economic Studies*, 71(4):1143–1162.

Vogl, T. S. (2013). Marriage institutions and sibling competition: Evidence from South Asia. *The Quarterly journal of Economics*, 128(3):1017–1072.

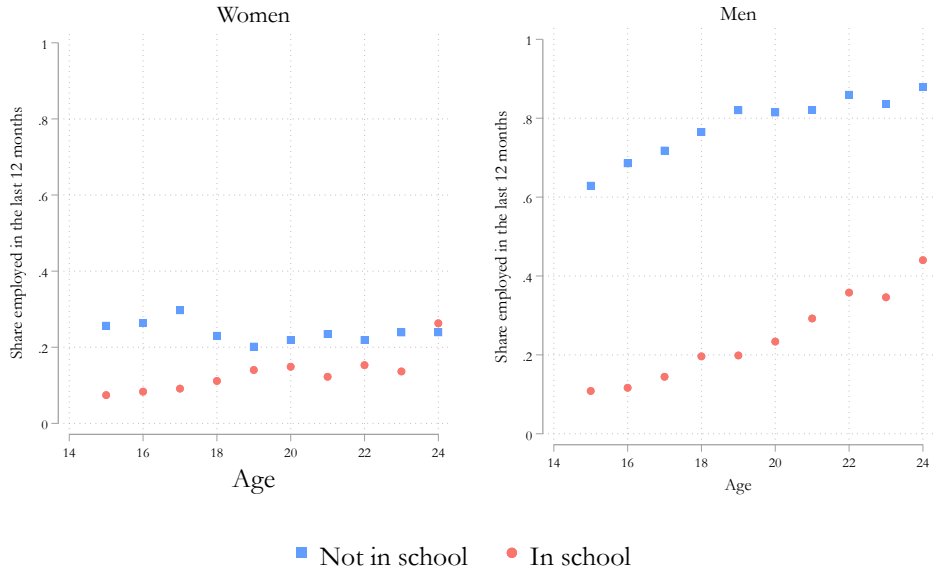
Weaver, J. and Chiplunkar, G. (2022). Marriage markets and the rise of dowry in india. *Available at SSRN 4092332*.

Wilson, N. (2022). Child marriage bans and female schooling and labor market outcomes: Evidence from natural experiments in 17 low-and middle-income countries. *American Economic Journal: Economic Policy*, 14(3):449–477.

Appendix

A Labor force participation

Figure 13: Labor force participation of men and women



Source: Own calculations using NFHS Eligible Women and NFHS Men (2019-21).

B Derivation of household acceptance probability

B.1 Probability that man (j, i) is picked among his own type

Consider a man of type j who has proposed to household $(\mathbf{h}, i)$. Let Z denote the number of *other* type-j men who also propose to $(\mathbf{h}, i)$. Under the Poisson assumption on arrivals, $Z \sim \text{Poisson}(\delta q_{j,\mathbf{h}})$. Conditional on $Z = z$, the household is indifferent between men of the same type and picks (j, i) with probability $1/(z+1)$:

$$\Pr\big((j,i) \text{ picked among own type} \mid (j,i) \text{ proposes to } (\mathbf{h},i),\ Z = z\big) = \frac{1}{z+1}. \tag{11}$$

Taking expectation over Z,

$$\begin{aligned}
&\Pr\big((j,i) \text{ picked among own type} \mid (j,i) \text{ proposes to } (\mathbf{h},i)\big) \\
&= \sum_{z=0}^{\infty} \frac{1}{z+1}\, e^{-\delta q_{j,\mathbf{h}}} \frac{(\delta q_{j,\mathbf{h}})^z}{z!} \\
&= \frac{e^{-\delta q_{j,\mathbf{h}}}}{\delta q_{j,\mathbf{h}}} \sum_{z=0}^{\infty} \frac{(\delta q_{j,\mathbf{h}})^{z+1}}{(z+1)!} \\
&= \frac{e^{-\delta q_{j,\mathbf{h}}}}{\delta q_{j,\mathbf{h}}} \big(e^{\delta q_{j,\mathbf{h}}} - 1\big) \\
&= \frac{1 - e^{-\delta q_{j,\mathbf{h}}}}{\delta q_{j,\mathbf{h}}}.
\end{aligned} \tag{12}$$

B.2 Probability that man (j,i) is selected conditional on proposing

Let $\mathcal{R} \subseteq \mathcal{J}$ denote the set of male types from which the household receives at least one proposal. Conditional on (j,n) proposing, $j \in \mathcal{R}$. For each other type $k \neq j$, the probability that at least one type-k man proposes to $(\mathbf{h},i)$ is $1 - e^{-\delta q_{k,\mathbf{h}}}$, and the probability that no type-k man proposes is $e^{-\delta q_{k,\mathbf{h}}}$. Hence

$$\Pr(\text{arrival set} = \mathcal{R} \mid j \in \mathcal{R}) = \prod_{k \in \mathcal{R}\setminus\{j\}} \big(1 - e^{-\delta q_{k,\mathbf{h}}}\big) \prod_{k \notin \mathcal{R}} e^{-\delta q_{k,\mathbf{h}}}. \tag{13}$$

Conditional on arrival set $\mathcal{R}$, the household draws a Type-I extreme value shock associated with each arriving type and selects the type that maximizes its value. The probability that type j is selected is given by the logit rule

$$\Pr(j \text{ selected} \mid \mathcal{R}) = \frac{\exp\big(W^a(\mathbf{h}, d_j, j)/\sigma\big)}{\sum_{x \in \mathcal{R}} \exp\big(W^a(\mathbf{h}, d_x, x)/\sigma\big)}. \tag{14}$$

Combining the three pieces,

$$\begin{aligned}
P_{j,(\mathbf{h},i)} = \; & \underbrace{\frac{1 - e^{-\delta q_{j,\mathbf{h}}}}{q_{j,\mathbf{h}}}}_{\text{within-type selection}} \sum_{\substack{\mathcal{R} \subseteq \mathcal{J} \\ j \in \mathcal{R}}} \underbrace{\left[\prod_{k \in \mathcal{R}\setminus\{j\}} \big(1 - e^{-\delta q_{k,\mathbf{h}}}\big) \prod_{k \notin \mathcal{R}} e^{-\delta q_{k,\mathbf{h}}} \right]}_{\Pr(\text{arrival set}=\mathcal{R} \mid j \text{ present})} \\
& \cdot \underbrace{\frac{\exp\big(W^a(\mathbf{h}, d_j, j)/\sigma\big)}{\sum_{x \in \mathcal{R}} \exp\big(W^a(\mathbf{h}, d_x, x)/\sigma\big)}}_{\text{choice among arriving types}}.
\end{aligned} \tag{15}$$