The Labor Hub

Job Market Paper Series No. 8

Advisor: Ajinkya Keskar

CROSS-SIBLING LINKAGES IN THE INDIAN MARRIAGE MARKET

Arigna Guha, Ajinkya Keskar, and Sulagna Mookerjee

Cross-Sibling Linkages in the Indian Marriage Market *

Arigna Guha *(Job Market Paper)* Ajinkya Keskar Sulagna Mookerjee

October 22, 2026

[Click here for the latest version]

Abstract

When can unequal spending across children benefit them both? We study this question in the Indian marriage market, where parents allocate scarce dowry resources across daughters who marry sequentially. We model successive marriages as a signalling game: the elder daughter's match conveys information about the family that affects the marriage terms available to her younger sister. We test the model's predictions using a survey experiment with 516 parents in Uttar Pradesh. The design elicits perceived returns to a favourable elder match, responses from parents of potential grooms, and willingness to depart from equal spending when both daughters' marriage outcomes are specified. Holding family resources and elder expenditure fixed, a high-quality elder match raises the younger sister's perceived probability of obtaining a high-quality groom by 47.8 percentage points compared with a low-quality elder match. Parents of daughters expect the dowry required for this fixed target to fall by Rs. 151,600, nearly twice the Rs. 76,500 reduction in stated requirements reported by parents of potential grooms. When displayed allocation outcomes include a benefit to the younger daughter, parents are 30.4 percentage points more likely to choose spending tilted toward the elder. This response is 11.7 points larger among parents whose earlier answers imply a positive cross-sister return. These findings identify a cross-sister return to marriage investment: a favourable elder match can improve the younger sister's opportunities and give parents a reason to invest more in the elder's marriage. This incentive may offset pressures that otherwise disadvantage elder daughters in sequential marriage. More broadly, cross-child returns can rationalize unequal parental spending.

Keywords: Marriage markets; Siblings; Spillovers; Beliefs; Fairness; Dowry; Gender; India

JEL codes: D13, D84, J12, J16, O15

---

*Authors: Guha: Binghamton University (SUNY). Email: aguha2@binghamton.edu. Keskar: Binghamton University (SUNY). E-mail: akeskar@binghamton.edu. Mookerjee: Binghamton University (SUNY). E-mail: smookerj@binghamton.edu. We thank Catalina Herrera Almanza, Prashant Bharadwaj, Palaash Bhargava, Audra J. Bowlus, Christopher Hanes, Subal Kumbhakar, Solomon Polachek, and David Slichter for helpful comments. This study was approved by the Binghamton University Institutional Review Board, Study Number: STUDY00007352. The study's American Economic Association Registry RCT ID: AEARCTR-0019119. While preparing this work, the authors used ChatGPT and Claude. All errors are our own.

1 Introduction

Fairness considerations shape how people allocate resources and which inequalities they accept (Almås et al., 2026). Within families, children can receive different resources, face different opportunities, and ultimately achieve different outcomes. How parents allocate resources shapes the intergenerational transmission of economic opportunity (Becker and Tomes, 1976; Behrman et al., 1982). Parents may sacrifice substantial returns to invest equally in their children (Berry et al., 2020). However, when investment in one child also improves another child's opportunities, unequal spending can benefit both. Whether parents exploit these returns depends on the spillovers they perceive and how they weigh equality in inputs against the resulting outcomes. We study this allocation problem in the Indian marriage market, where investment in one daughter's marriage can change her sister's opportunities.

Three features of the Indian marriage market make this allocation problem salient. First, marriages are predominantly arranged by families, giving parents a central role in selecting their children's spouses (Rubio, 2014).[1] Second, marriages commonly involve a substantial dowry payment from the bride's family to the groom's family (Anderson, 2007; Chiplunkar and Weaver, 2023), and dowry obligations shape households' saving behaviour (Anukriti et al., 2022). Third, daughters typically marry in birth order, linking sisters through the time available to search for grooms (Vogl, 2013).[2] This sequence can also create a cross-sister return. The elder daughter's realised match is visible in the marriage market. If prospective groom families use that match to infer family characteristics they care about, a favourable elder match may impact the terms available to her younger sister.[3] Investment in the elder's marriage therefore uses resources that could otherwise fund the younger's marriage, but can also improve the younger's marriage opportunities. We examine parents' perceived cross-sister returns, parents of potential grooms' responses to the elder daughter's match, and how parents weigh equal dowry spending against improvements in their daughters' marriage outcomes.

[1] In the 2011–12 India Human Development Survey, 95% of women report that their families were involved in selecting their husbands (NCAER National Data Innovation Centre, n.d.).

[2] Having a younger sister waiting to marry can shorten the elder daughter's search, leading her to marry earlier and into a worse match (Vogl, 2013; Guha et al., 2026). These pressures can coexist with positive cross-sister returns: a better elder match may benefit the younger sister even when marriage-order constraints make that match harder to obtain.

[3] Descriptive evidence from REDS 1999 motivates this possibility. Within families, dowry spending is approximately Rs. 19,500 higher (in 2010 rupees) for the eldest daughter than for her younger sisters. One additional year of schooling for the elder sister's husband is associated with 0.44 additional years for the younger sister's husband. These patterns do not establish a cross-sister return: spending and matches also reflect shared family characteristics, preferences, and constraints. Section 2 presents the evidence.

Building on Bloch and Rao (2002); Bloch et al. (2004), we model parents' marriage decisions for two daughters as a signalling game with incomplete information and a direct preference for equal dowry spending. Families differ in an unobserved type that captures characteristics valued by prospective groom families, such as social standing or networks. Parents know their own family's type, but prospective groom families cannot observe it directly. Instead, they observe the elder daughter's realised match and use it to infer the family's type. Parents therefore choose the elder daughter's match considering its cost, the resources remaining for the younger daughter, and how the match will shape prospective groom families' assessment of the family. Obtaining a high-quality elder match is less costly for families with a favourable type. When the benefit of this match is sufficient for favourable-type families but too small to justify imitation by the other type, the game has a unique separating equilibrium outcome satisfying the intuitive criterion (Cho and Kreps, 1987). Favourable-type families obtain a high-quality elder match, while the other families obtain a low-quality match, so the elder daughter's match reveals the family's type. Prospective groom families for the younger daughter update their assessment accordingly. Because they benefit from a connection to a favourable-type family, a more favourable assessment weakly lowers the dowry they require for the younger daughter's marriage.

Our model yields four testable predictions. First, a high-quality elder match improves prospective groom families' assessment of the family. Second, it weakly lowers the dowry required for the younger daughter to marry a given groom. Third, holding elder expenditure and remaining resources fixed, it improves the younger daughter's marriage prospects when better grooms offer at least as large a dowry concession in response to a favourable family assessment. Fourth, when both daughters' marriage outcomes are specified, improving the younger daughter's outcome under an allocation tilted towards the elder increases parents' willingness to choose that allocation. This response can arise even when parents prefer equal dowry spending: the benefit to the younger daughter makes unequal spending more attractive without requiring a stronger preference for the elder daughter.

We test these predictions using a survey experiment with 516 parents across two districts of Uttar Pradesh (Lucknow and Prayagraj). The design separately elicits parents' beliefs about cross-sister returns, groom families' stated assessments and dowry requirements, and parents' choices between equal and unequal dowry allocations. In the bride-side module, 258 parents with at least

two daughters each evaluate eight hypothetical scenarios that vary the elder daughter's groom quality, brother presence, and the age gap between daughters. Across respondents, we vary total marriage savings and the share spent on the elder daughter, holding both fixed throughout each respondent's scenarios. We compare each respondent's answers across scenarios that differ only in the elder daughter's groom quality, holding family composition, sibling spacing, and spending fixed. This ceteris paribus comparison isolates the perceived effect of the elder daughter's match on her younger sister's marriage opportunities. Respondents report how prospective groom families would assess the family, the younger daughter's prospects across groom categories, and the dowry required for her to marry a high-quality groom.[4] In the groom-side module, 258 parents of potential grooms evaluate hypothetical bride families after receiving information about the elder sister's realised match. We elicit their assessment of the family and the minimum dowry they would require. Finally, all 516 respondents complete an allocation task in which they choose between equal dowry spending and an allocation tilted towards the elder daughter. Each alternative specifies spending on both daughters and their resulting marriage outcomes, with variation in whether spending more on the elder also benefits the younger. These specified outcomes allow us to examine how parents trade off equal dowry spending against improvements in their daughters' matches.

We find a large perceived cross-sister return. First, replacing a low-quality elder match with a high-quality match raises the probability parents assign to the younger daughter obtaining a high-quality groom by 47.8 percentage points, from a reference mean of 13.7 percent (statistically significant at the 1% level). It also raises the probability that prospective groom families are expected to view the household more favourably by 85.1 percentage points and lowers the stated dowry required for a high-quality younger groom by Rs. 198,600 (in 2010 prices, statistically significant at the 1% level), a 22 percent reduction relative to the reference mean. Second, parents of potential grooms report responses in the same direction: a high-quality rather than low-quality elder match raises favourable family assessments by 71.7 percentage points and lowers the minimum dowry they would require by Rs. 76,500 (in 2010 prices, both statistically significant at the 1% level). However, the two sides differ in the magnitude of the dowry response. Averaging the bride-side response across family environments, parents of daughters expect a reduction of

[4] We assess respondents' familiarity with the vignette setting using their own marriage-transfer histories, reported involvement in marriage decisions, and prevailing sibling-marriage norms (see Section 7 for details).

Rs. 151,600—about twice the reduction reported by parents of potential grooms. The difference of approximately Rs. 75,200 is statistically significant at the 5% level. Parents of daughters thus expect a larger dowry concession than parents of potential grooms report.

Third, cross-sister returns change parents' willingness to depart from equal dowry spending. Introducing a benefit to the younger daughter from the elder-tilted allocation raises the probability of choosing that allocation by 30.4 percentage points (statistically significant at the 1% level). Holding both elder-daughter outcomes and the younger daughter's outcome under equal spending fixed, giving the younger daughter a better groom under tilted spending raises the probability of choosing the elder-tilted allocation by 56.9 percentage points (statistically significant at the 1% level). The response is also related to previously elicited beliefs: parents who expect a favourable elder match to reduce the younger daughter's required dowry respond 11.7 percentage points more strongly to the supplied cross-sister benefit (statistically significant at the 5% level). Cross-sister returns give parents a reason to invest more in the elder daughter's marriage because her match benefits both daughters. This incentive may counteract birth-order pressures that shorten her marriage search and worsen her match (Vogl, 2013; Guha et al., 2026). Observed marriage outcomes therefore reflect both the constraints families face and how parents allocate resources in response.

Our first contribution is to the literature on intra-household inequality and parental resource allocation. Prior work has examined how resources are distributed among children and how these allocations affect consumption and poverty (Penglase, 2021; Brown et al., 2021; Bhattarai et al., 2026). Related literature documents inequalities across children in educational attainment, breast-feeding, and child growth, and examines how these patterns reflect birth order, fertility decisions, and son preference (Black et al., 2005; Jayachandran and Kuziemko, 2011; Jayachandran and Pande, 2017). Related work examines how parents respond to differences in children's endowments, how investments spill over across siblings, and how fairness preferences shape allocation choices (Rosenzweig and Schultz, 1982; Behrman et al., 1994; Bharadwaj et al., 2018; Berry et al., 2020; Giannola, 2024). We extend this literature to marriage investment, where prospective groom families' responses to one daughter's match can improve her sister's opportunities. Our allocation experiment shows that these cross-sister benefits increase parents' willingness to depart from equal dowry spending.

Second, we contribute to research on marriage markets in developing countries and marital transfers. Existing work examines caste preferences in marital matching (Banerjee et al., 2013), the interaction between matching, migration, dowries, and spouse choice (Beauchamp et al., 2023), and the joint determination of wealth, dowries, and sex selection (Borker et al., 2026). Chiplunkar and Weaver (2023) document the emergence and evolution of dowry payments, while Anukriti et al. (2022) show how anticipated dowry obligations shape parental saving.[5] We examine how parents divide marriage resources across daughters. Our model combines the informational role of the elder daughter's match with a direct preference for equal dowry spending. Investing more in the elder's marriage reduces the resources available for the younger daughter, but can improve her marriage opportunities by signalling family attributes to prospective groom families. Parents therefore weigh the cross-sister return against the resource cost and their preference for equal spending.

Third, we contribute to research on perceived marriage-market returns and investment decisions. Attanasio and Kaufmann (2017) and Wiswall and Zafar (2021) examine how expected labour- and marriage-market returns relate to schooling choices. Andrew and Adams (2025) and Calvi et al. (2024) study parents' beliefs about daughters' marital returns to education. We extend this work to returns across sisters and compare what parents of daughters expect with what parents of potential grooms report. Both sides report that a favourable elder match lowers the younger daughter's required dowry, but parents of daughters expect a concession about twice as large. This comparison reveals agreement about the direction of the cross-sister return alongside a substantial difference in its perceived magnitude.

Fourth, we contribute to research on siblings and marriage outcomes.[6] Existing work documents how marriage-order norms constrain elder daughters' search and how sibling composition affects dowries and spousal quality (Vogl, 2013; Guha et al., 2026; Ahsan and Jain, 2026). We study a benefit operating alongside these pressures: a favourable elder match can improve the younger sister's marriage opportunities. This cross-sister return gives parents an incentive to invest more in the elder daughter's marriage and may counteract the disadvantages imposed by

[5] A large theoretical literature examines the origins and functions of dowry, interpreting it as a premortem bequest to daughters, a market-clearing transfer, or a combination of the two (Becker, 1981; Botticini and Siow, 2003; Anderson, 2003, 2007; Anderson and Bidner, 2015).

[6] A broader literature studies how siblings affect one another's human capital. Evidence includes older sisters' contributions to early childhood development in Kenya (Jakiela et al., 2025), returns to older sisters' schooling for younger brothers in Pakistan (Qureshi, 2018), spillovers in school achievement in England (Nicoletti and Rabe, 2019), and effects on college and programme choices across four countries (Altmejd et al., 2021).

marriage-order constraints. Observed differences between sisters can consequently reflect both market frictions and parents' allocation responses.

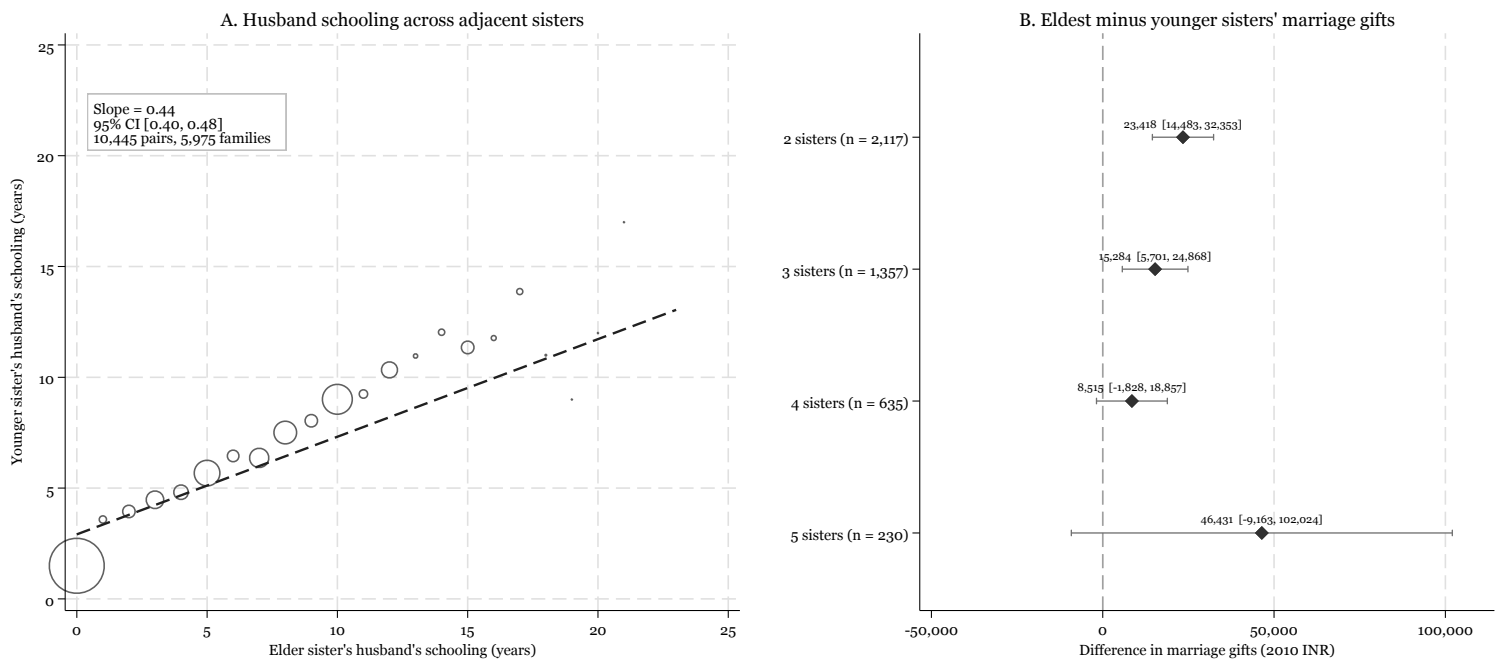
2 Setting and Observational Data

India's marriage market creates unusually strong links between siblings' marriage decisions. Arranged marriage remains the dominant institution, with parents and relatives playing a central role in spouse selection (Anukriti and Dasgupta, 2017; Rubio, 2014). Marriage is nearly universal, separation is rare, and residence is predominantly patrilocal, making a daughter's match consequential for her subsequent economic circumstances and welfare (Jacob and Chattopadhyay, 2016; Khalil and Mookerjee, 2019; Bloch and Rao, 2002; Srinivasan and Bedi, 2007; Calvi and Keskar, 2023). Securing a preferred match is also closely tied to dowry, a large transfer from the bride's family that responds to spouse characteristics and marriage-market conditions (Rao, 1993; Anderson, 2007; Chiplunkar and Weaver, 2023). For families with several daughters, these expenditures are made sequentially under a strong norm of arranging marriages according to birth order. A younger sister can therefore compete with the elder for both household resources and search time. The same sequencing, however, means that the elder daughter's realised marriage is observed before the younger daughter's match is secured. If a favourable elder match is interpreted by the market as a good signal, her match may alter her sister's subsequent opportunities.

To examine whether marriages appear linked across sisters in the data, we turn to the 1999 wave of the Rural Economic and Demographic Survey (REDS), which contains detailed sibling marriage histories, spousal characteristics, and marital transfers for rural Indian households. Two patterns stand out in Figure 1. First, marriage expenditure is tilted towards the eldest daughter. Families with two to four daughters spend approximately Rs∼23,400, Rs∼15,300, and Rs∼8,500 more on the eldest daughter's marriage than on her younger sisters, respectively. Second, match outcomes are strongly correlated across sisters. A one-year increase in the elder sister's husband's schooling predicts 0.44 additional years of schooling for the younger sister's husband, conditional on both sisters' education and marriage characteristics, caste, religion, state of residence, and the younger sister's marriage cohort. These patterns are striking when viewed alongside existing causal evidence. When an elder daughter is followed closely by a sister rather than a brother, she marries earlier and into a worse match and with a lower dowry, as the birth-order queue

compresses search time and places additional pressure on marriage resources (Vogl, 2013; Guha et al., 2026). The reduced-form effect of a nearby younger sister is therefore adverse. The natural interpretation is that several forces operate at once. A younger sister can impose queueing and resource costs on the elder, while a favourable elder match may itself generate a return for the younger sister. Existing work studies the first of these forces. The second, a return running from the elder daughter's realised match to her sister's marriage prospects, has not been measured and is the focus of this paper.

Figure 1: Sibling-linked marriage outcomes in REDS 1999



Notes: Panel A: Each point is the (unadjusted) mean schooling of younger sisters' husbands at a given number of years of elder sisters' husbands' schooling; marker size is proportional to the number of pairs. The dashed line shows the slope conditional on both sisters' schooling and age at marriage, the birth gap, caste, religion, father's and mother's schooling and log(1 + parents' land in acres) and state and younger sister's marriage-year Fixed Effects, standard errors clustered by state.
Panel B: Each point is the eldest sister's marriage gifts minus the mean for her younger sisters (2010 rupees), from a family fixed-effects regression estimated separately by number of sisters. Labels show the estimate and its 95% confidence interval. Standard errors clustered by family. n is the number of families.

The observational evidence cannot determine the relative importance of these forces. The REDS correlations cannot distinguish a causal cross-sister return from shared family type, caste, networks, wealth, or local marriage-market conditions. Conversely, sibling-composition estimates recover the net effect of a younger sister and therefore combine any positive return to the elder's match with the adverse resource, timing, and search pressures she creates. More generally, observed choices jointly reflect preferences, beliefs, and constraints, and in a matching market the responses of the other side, so realised outcomes alone cannot separate them (Manski, 2004; Wiswall and Zafar, 2018). Equal spending across daughters, for example, may be consistent both with parents who perceive no cross-sister return and with parents who perceive one but forgo it to

treat their daughters equally (Berry et al., 2020). The empirical problem is therefore to separate the return parents believe one daughter's marriage generates for another, the response of the other side of the market, and the allocation decisions parents make in light of those returns.

3 Conceptual Framework

We develop a simple signalling model of sequential marriage, building on the idea that visible marriage-market outcomes can convey otherwise private information (Bloch and Rao, 2002; Bloch et al., 2004). Each family privately knows its persistent type, which captures characteristics such as reputation, networks, or social standing that prospective in-laws value but cannot perfectly observe. Daughters marry in birth order, so the elder daughter's realised match is publicly observed before the younger daughter enters the market. The transfer used to obtain that match is not observed, because dowry is illegal and rarely announced. A high-quality elder match is less costly for favourable-type families to obtain, so the elder daughter's match can reveal the family's type. Prospective groom families update their assessment of the household and, in turn, the transfer they require to marry the younger daughter. Parents therefore choose the elder daughter's match weighing its cost, the resources left for the younger daughter, and its effect on her marriage terms. They also have fairness considerations and place a direct value on spending equally across daughters.

3.1 Setup

A family has resources $W > 0$ and two daughters who marry in birth order. An alternative P specifies the daughters' match qualities (m_1, m_2), the dowries or transfers (d_1, d_2) made at each marriage, resources retained afterwards x, and a non-pecuniary cost c incurred in obtaining the matches. Family context z collects observable composition characteristics such as sibling structure and the age gap between the daughters. Parent i evaluates an alternative according to

$$U_i(P;z) = \alpha_i(z)\,u_i(m_1;z) + [1-\alpha_i(z)]u_i(m_2;z) + x - c - \eta_i(z)\frac{|d_1-d_2|}{W}, \tag{1}$$

subject to $d_1 + d_2 + x = W$ and $x \geq 0$, with $u_i(\cdot;z)$ increasing, $\alpha_i(z) \in [0,1]$, and $\eta_i(z) \geq 0$.

The first two terms capture concern for the daughters' marriage outcomes, with α_i the weight

placed on the elder daughter. The final term captures a distaste for unequal spending. Parents may dislike spending unequally on their daughters, quite apart from what that spending achieves. The ratio $|d_1 - d_2|/W$ is the difference between the shares of the marriage budget spent on the two daughters, and η_i is the distaste factor that converts this difference into disutility, measured in rupees. A parent who weights the two daughters equally, $\alpha_i = 1/2$, can still choose unequal transfers when the resulting improvement in outcomes is large enough to outweigh this equality cost. Equal spending and equal treatment are not the same object here. We suppress i and z where they are held fixed.

Family type and the elder match. The family privately knows its type $\theta \in \{0, 1\}$, where $\theta = 1$ is the type prospective in-laws value more highly. This includes social standing, networks, and other persistent characteristics that affect marriage opportunities but that an outside family cannot perfectly observe. Its prior probability is $p \in (0, 1)$.

The elder daughter obtains a low or high match, $m \in \{L, H\}$. Securing match m requires a transfer t_m and a non-pecuniary acquisition cost $c_\theta(m)$. We assume $t_H \geq t_L$, so the better elder match is weakly more expensive.[7] Prospective groom families for the younger daughter observe the elder match and the public family context, but not the transfer used to obtain it. Their posterior assessment of the family is $\mu(m) = \Pr(\theta = 1 \mid m)$, and we write μ_H and μ_L for the posteriors following a high and a low elder match.

The younger daughter's market. After the elder marriage, an opportunity state ω determines the outside options available to prospective grooms for the younger daughter. Let the ordered set of groom qualities be $\mathcal{G} = \{g_1 < \cdots < g_J\}$. For groom quality g, let $a_g(\omega)$ denote the groom family's outside option and $v_g \geq 0$ the value it places on a connection to a $\theta = 1$ family. Its expected surplus from accepting a transfer d is $d + v_g\mu - a_g(\omega)$. The minimum transfer it will accept is therefore

$$D_g(\mu, \omega) = \max\{\underline{d}_g,\ a_g(\omega) - v_g\mu\}, \tag{2}$$

where $\underline{d}_g > 0$ is a customary minimum transfer for groom quality g. A more favourable assessment substitutes for part of the transfer needed to secure the same groom, but only until this floor binds.

[7] For the elder match, type affects the non-pecuniary acquisition cost rather than the transfer schedule.

Choosing the younger match. After spending t on the elder marriage, the family has $B = W - t$ remaining. Writing $\rho \equiv \eta/W$, the equality penalty becomes $\rho\,|d_1 - d_2|$, so ρ is its cost per rupee of spending gap. We assume $\rho < 1$, or equivalently $\eta < W$. The effective cost of a transfer d to the younger daughter, inclusive of the penalty for departing from what was spent on the elder, is then

$$h_t(d) \equiv d + \rho\,|t - d|.$$

Away from $d = t$, we get $h_t'(d) = 1 + \rho\ \mathrm{sign}(d - t)$, so the restriction $\rho < 1$ makes h_t strictly increasing. A lower transfer is always preferred for an otherwise identical match, because the rupee saved outweighs the cost of any wider gap it opens. We assume at least one groom quality is affordable in every relevant state.

Given (t, B, μ, ω), let $g^*(t, B, \mu, \omega)$ denote a groom quality that maximises $(1-\alpha)u(g) + B - h_t(D_g(\mu, \omega))$ over the affordable set $\{g : D_g(\mu, \omega) \le B\}$. Before ω is realised, the continuation value of the younger marriage is

$$V(t, B, \mu) = \mathbb{E}_\omega\Big[\max_{g\,:\,D_g(\mu,\omega)\le B} \{(1-\alpha)u(g) + B - h_t\big(D_g(\mu, \omega)\big)\}\Big], \tag{3}$$

with induced match distribution $F^*(\bar{g} \mid t, B, \mu) = \Pr_\omega\big[g^*(t, B, \mu, \omega) \le \bar{g}\big]$. Because D_g is weakly decreasing in μ, a more favourable assessment weakly expands the younger daughter's feasible set and raises the continuation value. Appendix A.1 establishes these monotonicity results formally.

3.2 Why the elder match conveys information

The elder match is informative only if the favourable type is more willing to obtain it. Let,

$$k_\theta = c_\theta(H) - c_\theta(L) - \alpha[u(H) - u(L)]$$

denote the additional non-pecuniary cost of the high elder match, net of the direct value the parent places on it. With $B_m = W - t_m$, the family chooses H whenever

$$k_\theta \le V(t_H, B_H, \mu_H) - V(t_L, B_L, \mu_L).$$

The right-hand side is what the elder choice does for the younger daughter. It combines the resources left after the elder marriage with the assessment subsequently attached to the family.

Writing $\Gamma \equiv V(t_H, B_H, 1) - V(t_L, B_L, 0)$ and $\Gamma^p \equiv V(t_H, B_H, 1) - V(t_L, B_L, p)$ for the gain from a high elder match that reveals the favourable type, measured against a low match that reveals the unfavourable type and against one that leaves beliefs at the prior, a sufficient condition for separation is

$$0 < k_1 < \Gamma^p \leq \Gamma < k_0. \tag{4}$$

The favourable type finds the continuation benefit large enough to justify the high elder match; the unfavourable type finds imitation too costly. The inequality $\Gamma^p \leq \Gamma$ follows because V is increasing in μ.

Proposition 1 (Informative elder match)**.** *Under* (4) *and* $\rho < 1$*, conditional on the competitive continuation, the unique sender outcome satisfying the intuitive criterion has type* $\theta = 1$ *choose* H *and type* $\theta = 0$ *choose* L*. On-path beliefs are* $\mu(H) = 1$ *and* $\mu(L) = 0$*.*

Appendix A.2 proves Proposition 1. Full separation is sufficient rather than necessary. Everything below needs only that the favourable type is more likely to obtain H, so that $\mu(H) > \mu(L)$; Appendix A.3 gives the corresponding partial-sorting result.

The model also separates two objects that are easy to conflate. Under separation,

$$\Gamma = \underbrace{V(t_H, B_H, 1) - V(t_H, B_H, 0)}_{\text{information benefit}} - \underbrace{\left[V(t_L, B_L, 0) - V(t_H, B_H, 0)\right]}_{\text{resource cost}}. \tag{5}$$

The first term asks what happens to the younger daughter when the family is assessed more favourably, holding elder expenditure fixed. The second is the cost of the larger elder transfer at a fixed assessment, counting both the resources displaced and the expenditure gap it opens between the sisters. A better elder match can therefore help the younger daughter even when the spending required to obtain it has a negative net return. Appendix A.1 shows that the resource-cost term is nonnegative under $\rho < 1$ and $t_H \geq t_L$.

3.3 Predictions

The model delivers four predictions. Predictions 1–3 concern beliefs, transfers, and younger-match opportunities; Prediction 4 concerns allocation once match consequences are known.

Prediction 1: Family assessment. If favourable-type families are more likely to obtain a high elder match, then

$$\mu(H) > \mu(L).$$

The argument follows directly from Bayes' rule and does not require full separation; see Appendix A.3.

Thus, a better elder match should improve how prospective groom families assess the family, which we measure with a question on family-standing from both the bride and groom's perspectives as outlined in Sections 4.1,4.2.

Prediction 2: Transfers for a fixed groom. Holding younger-groom quality and the opportunity state fixed, let $\mu_H > \mu_L$. Then

$$D_g(\mu_H, \omega) - D_g(\mu_L, \omega) = -\min\{v_g(\mu_H - \mu_L),\ [a_g(\omega) - v_g\mu_L - \underline{d}_g]_+\} \leq 0,$$

where $[x]_+ \equiv \max\{x, 0\}$. A better elder match weakly lowers the transfer required for the same younger groom. When the floor does not bind, the reduction is proportional to the improvement in assessment. Once the floor binds, further improvements in assessment no longer lower the required transfer, so family assessment and the transfer response need not share the same shape. Appendix A.4 derives this condition.

We test whether a better elder match weakly lowers the dowry required for the younger's groom, using bride-side parents' anticipated dowry for a fixed high-quality groom and groom-side parents' minimum acceptable dowry across elder-match qualities, described in Sections 4.1,4.2.

Prediction 3: Younger-match prospects. For $\mu_H > \mu_L$, define the information concession for groom g as

$$C_g(\omega) \equiv D_g(\mu_L, \omega) - D_g(\mu_H, \omega).$$

Suppose required transfers are weakly increasing in groom quality at each posterior and $C_g(\omega)$ is weakly increasing in groom quality. Then, at fixed elder expenditure and remaining resources,

$$F^*(\bar{g} \mid t, B, \mu_H) \leq F^*(\bar{g} \mid t, B, \mu_L) \qquad \text{for every } \bar{g}.$$

A more favourable assessment shifts the younger daughter's match distribution weakly upward in the sense of first-order stochastic dominance. Better grooms remain more expensive, but they receive at least as large a concession when family assessment improves. When the transfer floor does not bind, v_g weakly increasing in groom quality is sufficient for this concession ordering. Appendix A.5 gives the proof.

A better elder match should shift the younger daughter's likely matches towards better grooms, which we measure using bride-side parents' elicited probabilities across low-, medium-, and high-quality grooms, described in Section 4.1.

Prediction 4: Allocation when outcomes are known. Compare equal spending $E = (W/2, W/2)$ with an elder-tilted allocation $T = (rW, (1-r)W)$, where $r > 1/2$. Writing $\Delta u_{ji} = u_i(m_j^T; z) - u_i(m_j^E; z)$ and $\Delta c_i = c_i^T - c_i^E$, and noting that both alternatives exhaust the budget so that x cancels,

$$A_i(z) \equiv U_i(T; z) - U_i(E; z) = \alpha_i(z)\Delta u_{1i} + [1-\alpha_i(z)]\Delta u_{2i} - \eta_i(z)(2r-1) - \Delta c_i.$$

Holding family context, elder outcomes, transfers, preferences, and Δc_i fixed, improving the younger daughter's match under tilted spending, or worsening it under equal spending, weakly raises willingness to tilt. The prediction survives a direct distaste for unequal expenditure: η_i shifts the level of A_i but not the sign of its response to Δu_{2i}. Appendix A.6 derives this result and its choice-probability counterpart.

Parents should become more willing to tilt resources towards the elder daughter when doing so improves the younger daughter's marriage outcome relative to equal spending. We test this using allocation choices in which the consequences of both spending alternatives are specified, comparing scenarios that improve the younger daughter's relative outcome while holding the elder daughter's outcomes fixed, as outlined in Section 4.3

Two limits are worth stating. First, parents in the model hold correct beliefs about the reservation schedule D_g. For the same younger groom, opportunity state, and public family context, the benchmark therefore contains no wedge between the concession parents expect and the concession groom families require. A difference between the two is a departure from the common-beliefs benchmark; interpreting that difference as misperception additionally requires comparable targets

and contexts across the two sides. Second, sibling spacing enters through z as a conditioning characteristic. We do not model the search process through which a closely spaced younger sister may constrain the elder daughter's attainable match.

4 Survey Instruments and Experimental Design

Our empirical design separately measures the three objects that realised marriage outcomes combine: the cross-sister return perceived by parents of daughters, the response reported by parents of sons, and parents' willingness to act on that return when allocating resources across daughters. We use experimentally varied hypothetical vignettes, following work that uses stated beliefs and choices to separate perceived returns from the preferences governing subsequent decisions (Giannola, 2024; Calvi et al., 2024; Berry et al., 2020; Andrew and Adams, 2025). This approach is particularly useful in marriage markets, where realised matches are an equilibrium outcome jointly determined by both families and observed choices need not reveal the beliefs under which they were made (Manski, 2004).

Respondents first complete either the bride- or groom-side marriage-market module and then an allocation module common to the full sample. The bride module varies the elder daughter's realised match while holding the resources used for her marriage fixed. The groom module elicits the corresponding response from parents of sons. The allocation module then supplies the marriage outcomes generated by alternative resource allocations directly, allowing us to study choices without requiring respondents to infer how spending maps into match outcomes. All modules concern hypothetical families whom respondents are asked to advise. Throughout, marriage transfers are described as "gifts," using familiar examples such as jewellery, household goods, vehicles, land, and cash.[8] The treatment dimensions and outcome measures in each module are described below.

4.1 Bride-side beliefs

The bride-side module asks whether parents of daughters perceive the cross-sister linkage described in Section 3. Each respondent evaluates eight scenarios concerning hypothetical families

[8] We use 'dowry' in the paper when referring to these gifts or transfers. Piloting indicated that respondents were more forthcoming about marriage payments when they were described as gifts rather than using the legally and socially sensitive term *dowry* or *dahej*.

with two daughters. The elder daughter is of marriageable age and receives a proposal, while the younger daughter's future marriage remains uncertain. The family has saved a fixed amount for both daughters' marriages, and the vignette states the transfer required for the elder daughter's proposal and the resources that would remain for the younger daughter.

We exogenously vary three features of the family across scenarios. First, the proposed elder groom is either low quality, a Class 8 pass factory labourer, or high quality, a college graduate working as a government officer. For context, among recently married men in our setting, roughly a quarter have less schooling than the low-quality groom and roughly a quarter have at least as much as the high-quality groom.[9] Second, the household either has a son or does not.[10] Third, the daughters are either two or five years apart. Across respondents, total marriage savings are either Rs. 5,00,000 or Rs. 10,00,000, and accepting the elder proposal requires either 60 or 70 percent of those savings. These quantities remain fixed within-respondent.

The transfer required for the elder marriage, and therefore the resources remaining for the younger daughter, is held fixed when the proposed groom changes from low to high quality. The High–Low comparison therefore varies the elder daughter's realised match while holding the resource cost of obtaining it fixed, separating the perceived effect of the elder match from the mechanical effect of leaving fewer resources for the younger sister. This corresponds to the information component of (5).

Each scenario elicits four outcomes. Respondents first advise whether the family should accept the elder daughter's proposal at the stated transfer. Then, *conditional on the elder daughter marrying the proposed groom*, they report three beliefs about the younger daughter's subsequent marriage market. First, they report the likelihood that she obtains a low-, medium-, or high-quality groom.[11] Second, they report how groom families would view the bride's family, on a five-point scale from "much worse" to "much better". Third, they state the value of gifts they think the younger daughter would need to marry a high-quality groom. Respondents who did not provide a numeric amount were asked whether this reflected a substantive response, such as no realistic gift

[9] Our calculations from the Individual Recode of NFHS-5, 2019–21. The sample comprises currently married women in Uttar Pradesh whose first cohabitation began within ten years of the interview ($N = 3{,}616$), and the outcome is the husband's completed years of schooling as reported by the wife. All figures use the women's sampling weights. The share of husbands with fewer years of schooling than each groom type: 23.5 percent have fewer than eight years, and 75.7 percent have fewer than fifteen, the threshold for a three-year undergraduate degree. The medium-quality groom, with twelve years of schooling, sits at the 59th percentile.

[10] We do not tell respondents that the son contributes to household liquidity or otherwise specify his role. Respondents are left to interpret what his presence implies for the family.

[11] The medium groom is defined by a Class 12 education and a private-sector job. The low- and high-quality groom continue to be defined as stated earlier.

being sufficient, or refusal or difficulty answering. The exact questions, descriptors, and sample vignettes are shown in Appendix B.1.

The acceptance question captures the broader decision of whether the elder match is worth committing a large share of available marriage resources and can reflect both its direct value and its expected consequences for the younger daughter. We therefore do not interpret it as a separate test of a single prediction. The remaining outcomes map directly into the model. The expected match distribution corresponds to Prediction 3, family standing to Prediction 1, and the transfer required for a fixed high-quality groom to Prediction 2. Because elder expenditure, remaining resources, and the younger groom are held fixed in the relevant High–Low comparisons, these outcomes isolate the perceived effect of the elder daughter's realised match conditional on family context (z).

For clarity, let $\widetilde{D}_g$ denote the bride-side perceived counterpart to the reservation-transfer schedule D_g in (2). The High–Low difference in the fixed-groom question measures the concession parents of daughters expect prospective groom families to grant,

$$\Delta \equiv \widetilde{D}_g(\mu_L) - \widetilde{D}_g(\mu_H).$$

Under the common-beliefs benchmark in Section 3, this perceived concession should coincide with the concession elicited directly from groom-side respondents for an otherwise comparable match. Section 4.2 examines this benchmark.

The age gap and brother treatments allow the perceived linkage to vary with family context. The age gap captures variation in the queueing pressure created by the younger daughter's proximity to marriageable age, while brother presence may alter how respondents interpret the elder match or the resources available across siblings. In the model, both enter through z.

Empirical specification. The eight scenarios form a complete $2 \times 2 \times 2$ factorial within respondent. For respondent i and scenario s, let H_{is} be a binary indicator equal to one when the elder daughter's proposed groom is of high rather than low quality, and define B_{is} and Q_{is} analogously

for the presence of a brother and a two-year rather than five-year age gap. We estimate

$$\begin{aligned} Y_{is} = \alpha &+ \beta_H H_{is} + \beta_B B_{is} + \beta_Q Q_{is} \\ &+ \beta_{HB} H_{is} B_{is} + \beta_{HQ} H_{is} Q_{is} + \beta_{BQ} B_{is} Q_{is} \\ &+ \beta_{HBQ} H_{is} B_{is} Q_{is} + \nu_i + \zeta_{p(is)} + \varepsilon_{is}. \end{aligned} \tag{6}$$

where ν_i denotes respondent fixed effects and $\zeta_{p(is)}$ vignette-position fixed effects. Standard errors are clustered by respondent.

We estimate (6) separately for acceptance, the younger daughter's expected match distribution, perceived family standing, and the stated transfer required for the fixed high-quality younger groom. The coefficient β_H gives the High–Low effect for the reference family with no brother and a five-year age gap, while the interaction terms show how that response varies with sibling composition and spacing. For the fixed-groom transfer outcome, $-\beta_H$ is the perceived High–Low concession, $\widetilde{\Delta}$, for the reference family.

4.2 Groom-side responses

The groom-side module elicits the corresponding response from parents of sons. Although the model is written from the bride household's allocation problem, Predictions 1 and 2 require prospective groom families to respond to the elder daughter's realised match. The bride module elicits what parents of daughters expect this response to be. The groom module instead asks parents of sons whether the elder match changes their assessment of the bride's family and the minimum transfer they would require for their son to marry the younger daughter. Under the common-beliefs benchmark in Section 3, the concessions elicited from the two sides should coincide for otherwise comparable matches.

Respondents advise the parents of a hypothetical son. Two features are assigned across respondents and remain fixed throughout the module. First, the son is either a medium-quality groom, a Class 12 pass private-sector worker, or a high-quality groom, a college graduate working as a government officer. Second, the marriage-wealth level used to construct the bride-family scenarios is either Rs. 5,00,000 or Rs. 10,00,000. Respondents are not told this total directly; it determines the rupee values of the gift offers they see. Within-respondent, we vary the bride family's composition and the realised match of its elder daughter. Respondents evaluate one family

with two daughters and no son and another with two daughters and a son.

The module proceeds in two stages for each bride-family type. In the baseline scenario, respondents are told that the elder daughter is already married but receive no information about whom she married. The younger daughter is proposed to the son with a gift offer equal to half of the bride household's marriage wealth. Respondents report the likelihood that they would advise accepting the proposal, on a scale from zero to ten, and the lowest gift offer at which they would advise acceptance.

Respondents are then given new information about the same family. The elder daughter is revealed to have married a low-, medium-, or high-quality groom.[12] Before each update, respondents are reminded of their answers to the corresponding baseline scenario. The gift offered for the younger daughter moves in the opposite direction to elder-match quality, mimicking the depletion of bride-household resources associated with securing a better elder match. The family offers 70 percent of its marriage wealth following a low-quality elder match, 40 percent following a medium-quality match, and 30 percent following a high-quality match. The acceptance question therefore presents a deliberate trade-off: a better elder match may make the proposing family more desirable, while the transfer offered becomes smaller.

After receiving the new information, respondents again report their acceptance likelihood and minimum acceptable gift. They also report whether the elder sister's realised match makes them view the proposing household as a better or worse family to marry into, using the same five-point scale as in the bride-side module. Respondents who did not provide a numeric minimum for their reservation dowry were also asked whether this reflected a substantive response, including that no gift amount would induce acceptance or that gifts were irrelevant, or refusal or difficulty answering. The exact questions and sample vignettes are shown in Appendix B.2.

The family-standing question provides the groom-side counterpart to Prediction 1, while the minimum acceptable gift corresponds to Prediction 2 and the reservation schedule $D_g(\mu, \omega)$ in (2). Acceptance instead combines these margins; a better elder match can raise the value of the family connection and lower the reservation transfer, while the gift actually offered simultaneously falls. The model therefore does not sign the change in acceptance across update states.

Unlike the bride module, the groom module includes a medium-quality elder match. This

[12] The low-quality groom is defined by a Class 8 education and employment as a factory labourer; the medium-quality groom by a Class 12 education and private-sector employment; and the high-quality groom by a college education and government employment.

intermediate state allows us to examine the shape of the response: the transfer floor in (2) permits family assessment to continue improving even after the minimum acceptable transfer has stopped falling. Brother presence provides an additional source of heterogeneity and, as in the bride module, enters the family context z without a signed prediction.[13]

Empirical specification. The two-stage structure of the groom-side vignette allows us to measure updating relative to the respondent's own baseline assessment of the same bride-family type. For outcomes observed in both stages, if $f(s)$ denotes the family corresponding to update scenario s, define

$$\Delta Y_{is} = Y_{is}^{\text{update}} - Y_{if(s)}^{\text{baseline}}. \tag{7}$$

Differencing removes the respondent's baseline assessment of that family. Contrasts across the low-, medium-, and high-match updates then measure how responses change with the elder sister's realised match.

Let M_{is} and H_{is} be binary indicators equal to one when the revealed elder match is of medium and high quality, respectively, so that the low-quality match is the omitted category, and define B_{is} analogously for the presence of a brother in the bride's family. We estimate

$$\Delta Y_{is} = \alpha + \beta_M M_{is} + \beta_H H_{is} + \beta_B B_{is} + \beta_{MB} M_{is} B_{is} + \beta_{HB} H_{is} B_{is} + \nu_i + \zeta_{p(is)} + \varepsilon_{is}. \tag{8}$$

where ν_i denotes respondent fixed effects and $\zeta_{p(is)}$ vignette-position fixed effects. Standard errors are clustered by respondent. For the family-standing outcome, which is elicited only after the elder match is revealed, we estimate the same specification in levels rather than changes from baseline.

For the reference family without a brother, β_H gives the High–Low response and β_M the Medium–Low response. The change in the minimum acceptable transfer,

$$-\beta_H = D_g(\mu_L, \omega) - D_g(\mu_H, \omega) \equiv \Delta$$

is the groom-side High–Low concession, while $-\beta_M$ is the corresponding Low–Medium conces-

[13] The groom module also provides a supplementary check on Prediction 3. Prediction 3 requires the concession generated by a more favourable family assessment to be weakly increasing in groom quality. Because son quality is assigned between respondents and held fixed within respondent, we can compare the High–Low reservation-transfer response for medium- and high-quality sons. A larger concession for high-quality sons is consistent with this ordering condition, although the bride-side match-distribution elicitation remains the direct test of Prediction 3. It is presented in Appendix E

sion. For comparisons that hold the younger-groom target and family context comparable across modules, the common-beliefs benchmark in Section 3 implies $\widetilde{\Delta} = \Delta$. We denote the empirical difference

$$\Omega \equiv \widetilde{\Delta} - \Delta,$$

with departures from zero measuring disagreement between the bride-side anticipated concession and the groom-side stated concession rather than a separate prediction of the model.

4.3 Allocation choices under certainty

The prior modules ask respondents how they believe an elder daughter's realised match changes the opportunities facing her younger sister. The allocation module instead takes those consequences as given. Respondents are shown exactly whom each daughter would marry under two alternative divisions of a fixed marriage budget and choose between the resulting bundles. Beliefs about how marriage expenditure maps into match outcomes therefore do not enter the choice. The module provides the direct empirical counterpart to Prediction 4: whether parents become more willing to accept unequal expenditure when doing so improves the younger daughter's relative marriage outcome.

Each scenario describes a family with two daughters and a fixed lifetime marriage budget, either Rs. 5,00,000 or Rs. 10,00,000. Respondents choose between an equal allocation, which divides the budget 50–50, and an allocation tilted towards the elder daughter. The tilted allocation assigns either 60 or 70 percent of the budget to the elder daughter, with the remainder going to the younger. Total marriage savings and the size of the tilted allocation remain fixed within-respondent. The resulting matches are shown on a five-point ladder ranging from very poor to very good, with each category defined by the groom's education and occupation.

The full design varies three features of the family environment, namely, whether the outcome mapping contains a cross-sister return, whether the daughters are two or five years apart, and whether the family has a brother who contributes to marriage resources. These dimensions determine the match bundle associated with each spending option. There are eight possible cells in the full design, of which each respondent evaluates four. A simple example illustrates the cross-sister variation. Consider a family with a wide age gap and no brother. Without a cross-sister return, the two alternatives are

	Elder daughter	Younger daughter
Tilted allocation	Good	Poor
Equal allocation	Medium	Medium

When the cross-sister return is introduced, the tilted allocation instead produces (Good, Medium), while the equal allocation remains (Medium, Medium). The spending alternatives and the elder daughter's outcomes are unchanged; only the younger daughter's outcome under the tilted allocation improves. Prediction 4 therefore implies greater willingness to choose the tilted allocation in the latter case. The complete set of eight outcome bundles is reported in Appendix B.3.

The remaining dimensions change the family context (z) in which the allocation is made. A tight age gap lowers the elder daughter's attainable match in the displayed bundles, capturing the queueing pressure associated with the younger daughter's marriage following closely behind (Vogl, 2013; Guha et al., 2026). The brother contributes resources and therefore improves the matches attainable from a given division of the parents' marriage budget for both daughters.[14] In the conceptual framework, both characteristics enter through z and allow the response to the same cross-sister return to vary with family context.

The module also lets us examine parents' preference for equal spending separately from the cross-sister return itself. Across some scenarios, tilting resources towards the elder daughter improves her match while leaving the younger daughter's match unchanged. Choosing equality in these cases means forgoing a better outcome for the elder without improving the younger daughter's displayed outcome. Such choices are consistent with a direct value placed on equal expenditure, η_i, in the equation stated in Prediction 4. We therefore use these scenarios to ask whether equal spending retains value even when it comes at a cost in displayed marriage outcomes. We also compare them with scenarios in which tilting improves the elder daughter's match but worsens the younger daughter's, providing a broader measure of how strongly parents respond to the sacrifice imposed on the younger daughter. The full set of outcome bundles is reported in Appendix B.3.

Empirical specification. The outcome Tilt_{is} equals one when respondent i chooses the elder-tilted allocation in scenario s and zero when they choose equal spending. Let S_{is} be a binary indicator equal to one when the displayed outcomes include a cross-sister spillover, and define B_{is}

[14] Liquidity is an important concern for parents in this setting (Anukriti et al., 2022), and brothers may alleviate such constraints (Ahsan and Jain, 2026). We examine the validity of these design choices in Section 7.

and Q_{is} analogously for the presence of a contributing brother and a two-year rather than five-year age gap. Our primary specification is

$$\text{Tilt}_{is} = \alpha + \beta_1 S_{is} + \beta_2 B_{is} + \beta_3 Q_{is} + \beta_4 S_{is} Q_{is} + \beta_5 S_{is} B_{is} + \nu_i + \zeta_{p(is)} + \varepsilon_{is}. \tag{9}$$

where ν_i denotes respondent fixed effects and $\zeta_{p(is)}$ vignette-position fixed effects. Standard errors are clustered by respondent. The coefficient β_1 summarises the change in tilt choice associated with the cross-sister spillover for the reference family, while the interactions allow this response to vary with sibling spacing and brother presence.

Prediction 4 is tested most directly by comparisons that hold the spending alternatives and the elder daughter's match outcomes fixed while improving the younger daughter's outcome under the tilted allocation relative to equality. These comparisons isolate the change in the younger daughter's relative consequence. We report these comparisons, together with cases in which tilting improves the elder daughter's match without worsening the younger daughter's, in Section 6.4, Table 4.

Finally, the earlier marriage-market module allows us to ask whether respondents who previously perceived a cross-sister linkage respond differently when that linkage is built directly into the allocation consequences. Among respondents in the bride-side module, let $\widehat{\Delta}_i$ denote the High–Low concession implied by the fixed-groom transfer question, and define the binary indicator

$$\textit{Believer}_i = \mathbf{1}\{\widehat{\Delta}_i > 0\}. \tag{10}$$

We estimate

$$\text{Tilt}_{is} = \alpha + \beta_1 S_{is} + \tau\,(S_{is} \times \textit{Believer}_i) + \beta_2 B_{is} + \beta_3 Q_{is} + \nu_i + \zeta_{p(is)} + \varepsilon_{is}. \tag{11}$$

The interaction τ measures heterogeneity in the spillover effect by previously elicited belief type. It asks whether respondents who had already revealed a perceived cross-sister concession respond more strongly when the allocation consequences incorporate that same linkage with certainty.

5 Data and Field Implementation

Setting and sampling. The study was conducted in Lucknow and Prayagraj districts of Uttar Pradesh between 19 July and 3 August 2026. Uttar Pradesh is a useful setting for studying sibling linkages in marriage because arranged marriage, dowry, and birth-order norms remain salient in the region (Jones, 2017; Karve, 1965; Kolenda, 1987; Srivastava et al., 2021). Fieldwork was implemented with Shailam Longitudinal Research Pvt. Ltd., a Lucknow-based survey firm. The sampling design began with 26 primary sampling units (PSUs), 13 in each district, selected with probability proportional to population size from the 2011 Census. Within selected PSUs, enumerators recruited eligible parents through *anganwadi*-assisted canvassing in rural areas and door-to-door screening in urban areas. Bride-side respondents were required to have at least two daughters, and groom-side respondents at least one son. Recruitment within PSUs was therefore based on eligibility and availability. The final sample contains 516 respondents, split evenly between the bride and groom arms. All experimental blocks and scenario orders were assigned to survey slots before respondent contact and loaded automatically in SurveyCTO. We used Latin-square counterbalancing to distribute scenarios across presentation positions and prevented enumerators from selecting treatment values or order. Appendix F documents PSU selection and replacement, recruitment, piloting, assignment, consent, sample certification, and field-quality procedures.

Respondent characteristics. Appendix C, Table 8 describes the final sample. Respondents are on average 49 years old and have 3.1 living children, placing most in an age range where fertility is likely complete and marriage decisions for their children are salient. The sample is predominantly rural, with 22 percent of respondents living in urban PSUs, broadly reflecting the state's level of urbanisation. Sixty-nine percent belong to SC, ST, or OBC households, 66 percent own land, and 97 percent own their home. Educational attainment differs sharply by gender (Appendix C, Figure 13). Almost 60 percent of female respondents have either no formal schooling or only primary education, compared with roughly 19 percent of male respondents. The relatively low levels of respondent education also motivated the attention paid during piloting to reducing cognitive burden and simplifying repeated tasks.

Marriage transfers are economically substantial in this population. Figure 14 in Appendix C, plots the distribution of gifts reported at respondents' own marriages, expressed in 2010 rupees.

The median transfer is approximately Rs. 59,000, roughly 58% of national annual income per capita,[15]with an interquartile range of roughly Rs. 26,000 to Rs. 108,000 and a 90th percentile above Rs. 208,000. The distributions are similar for male and female respondents, indicating that large marriage transfers are familiar to respondents on both sides of the marriage market. Appendix C, Table 8 reports the remaining demographic, household-asset, social-group, and family-composition characteristics.

6 Results

We organise the results into five sections. Section 6.1 asks whether parents perceive the elder daughter's marriage as consequential for her younger sister. We then turn to the other side of the market and ask whether groom families in fact respond to the elder match, in Section 6.2. Third, Section 6.3 compares the magnitude of the perceived and realised responses from the two sides of the market. Finally, Sections 6.4 and 6.5 examine how parents allocate resources when cross-sibling returns are made certain, and test whether the beliefs elicited earlier in the survey predict subsequent allocation choices. The main specifications exploit within-respondent experimental variation and include respondent and scenario-position fixed effects, with standard errors clustered at the respondent level.

6.1 Do parents perceive cross-sibling linkages?

We first ask whether parents are willing to commit disproportionately large resources to the elder daughter's marriage when the prospective groom is sufficiently valuable, while the younger daughter's eventual marriage outcome remains uncertain. Replacing a low-quality elder groom with a high-quality groom raises the probability of advising acceptance, thus committing a dowry exceeding half of the household's total marriage resources, by 83.0 percentage points on average ($p < 0.001$). In the reference scenario, acceptance rises from 13.2 percent under the low-quality match to 96.1 percent under the high-quality match. The response is remarkably stable across the four brother-by-age-gap environments, with implied effects ranging from 81.8 to 85.3 percentage points (Appendix D, Table 9). Parents are therefore willing to allocate substantially more marriage

[15]We benchmark against GDP per capita, which stood at US$2,702.5 in 2025 (World Bank, *World Development Indicators)*. Converting at the 2025 period-average exchange rate of Rs. 87.15 per US dollar and deflating by the consumer price index (2010 = 100, IMF *International Financial Statistics*, via *World Development Indicators*) gives approximately Rs. 101,000 in 2010 rupees.

resources to the elder daughter, despite uncertainty over whether the remaining resources will be sufficient for her younger sister, when the elder match is sufficiently attractive. This first-stage decision can reflect both the direct value of the elder match and its expected consequences for the younger daughter; the remaining outcomes separate these channels more directly.[16]

Table 1: Perceived cross-sister effects of the elder daughter's match

	Younger sister's expected match			Family	Required Dowry
	Low	Medium	High	standing	high groom
High elder match	-0.385***	-0.093***	0.478***	0.851***	-198.6***
	(0.017)	(0.011)	(0.017)	(0.026)	(31.8)
High × Brother	0.009	0.015	-0.024	-0.031	54.0**
	(0.013)	(0.013)	(0.017)	(0.026)	(24.0)
High × Tight gap	-0.017	0.030***	-0.012	-0.019	67.2**
	(0.012)	(0.011)	(0.014)	(0.026)	(27.1)
High × Brother × Tight gap	-0.007	-0.028*	0.035*	0.031	-54.4
	(0.017)	(0.015)	(0.021)	(0.037)	(35.9)
Reference Group Mean	0.439	0.424	0.137	0.081	887.9
Observations	2,064	2,064	2,064	2,064	2,001
Respondents	258	258	258	258	258

Notes: Each column reports a separate within-respondent regression. The reference group scenario is a low elder match, no brother, and a wide age gap. Dowry is reported in Rs '000 (2010 prices). All specifications include respondent and scenario-position fixed effects. Standard errors are clustered by respondent. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 1 presents estimates from (6), first reporting parents' beliefs about the distribution of potential grooms available to the younger daughter conditional on the elder daughter marrying the proposed groom. The results provide direct support for Prediction 3. In the reference scenario, a high-quality rather than low-quality elder match lowers the expected probability of a low-quality younger match by 38.5 percentage points and that of a medium-quality match by 9.3 percentage points, while raising the probability of a high-quality match by 47.8 percentage points (all $p < 0.001$). Relative to the reference-group mean of 13.7 percent, the probability assigned to a high-quality younger match therefore rises to 61.5 percent. The shift is also strikingly stable across family composition and sibling spacing. The implied increase in the probability of a high-quality younger match ranges from 45.4 to 47.8 percentage points across the four environments. Parents thus perceive the elder daughter's realised match as substantially changing the younger daughter's marriage prospects. Perceived marriage-market returns of this kind matter for investment.

The next outcome speaks directly to the information channel in Prediction 1. Under a low-quality elder match, only 8.1 percent of respondents in the reference scenario expect prospective

[16]The bride module presents only low- and high-quality elder matches. An earlier design included a medium-quality groom, but piloting indicated that presenting all three qualities would also have raised the scenario count beyond what respondent cognitive load allowed. Field implementation and piloting procedures are outlined in Appendix F

groom families to view the family more favourably when arranging the younger daughter's match. A high-quality elder match raises this probability by 85.1 percentage points ($p < 0.001$). The High–Low effect remains between 82.0 and 85.1 percentage points across the four environments. Bride-side parents therefore attach substantial informational value to the elder daughter's realised match. Marrying well is expected to improve how the family itself is assessed in the subsequent marriage market.

Finally, we ask whether this perceived improvement in family standing translates into the terms on which the younger daughter can obtain a high-quality groom. This corresponds to Prediction 2. Respondents expect the required dowry to fall by Rs 198,600 when the elder daughter marries the high-quality rather than the low-quality groom ($p < 0.001$). This is a 22 percent reduction relative to the reference-group mean of Rs 887,900. The dowry response is more sensitive to the surrounding family environment than either the perceived match distribution or family standing. Brother presence attenuates the High–Low concession by Rs 54,000 ($p < 0.05$), while a tight age gap attenuates it by Rs 67,200 ($p < 0.05$). One interpretation, suggested by our qualitative interviews, is that a brother dampens the positive informational value of a strong elder match because prospective groom families expect daughters to receive a smaller share of the family inheritance. One possible explanation as to why tight age gap may similarly reduce the value of the signal is by compressing the search time available before the younger daughter is expected to marry. Even so, the implied concession remains large and negative in every brother-by-spacing cell, ranging from approximately Rs 131,000 to Rs 199,000. Nor does the concession vary detectably across the randomised household-resource or elder-expenditure blocks (Appendix E,Table 13).

The monetary estimates condition on respondents who provide a numeric amount. This selection is itself informative. Respondents are more likely to provide a numeric answer following a high-quality elder match, while responses that "no realistic gift amount" would secure the target groom are concentrated in the low elder-match state (Appendix D, Figure 15). Under the natural interpretation that "no realistic amount" corresponds to a required transfer above the range respondents are willing to state numerically, restricting the regression to numeric responses understates the High–Low concession, thus our estimates are conservative.

Taken together, the bride-side experiment provides a consistent account of the cross-sister linkage perceived by parents. A better elder match is expected to shift the younger daughter's

match distribution sharply upward, improve the family's standing among prospective groom families, and reduce the transfer required to obtain a fixed high-quality groom. These responses line up respectively with Predictions 3, 1, and 2. The next question is whether these beliefs describe an actual response on the other side of the market (Section 6.2).

6.2 Does the marriage market respond to the elder match?

We next ask whether the cross-sibling linkage perceived by parents of daughters is also present on the other side of the marriage market. Groom-side parents evaluate proposals from otherwise comparable bride families after learning where that family's elder daughter married. As elder-match quality improves, the dowry offered by the bride family is deliberately reduced, reflecting the greater resources required to finance the elder daughter's better match. Respondents therefore face a trade-off between a more favourable signal about the bride family and a smaller current transfer. The acceptance decision captures the net effect of these two forces, while the subsequent family-assessment and reservation-dowry questions isolate the response to the elder match more directly.

Table 2: Groom-family responses to the elder sister's realised match

	Acceptance update	Family standing	Reservation Dowry
Medium − Low	0.017	0.254***	-80.4***
	(0.014)	(0.027)	(22.1)
High − Medium	-0.028***	0.463***	3.9
	(0.011)	(0.030)	(11.7)
High − Low	-0.011	0.717***	-76.5***
	(0.018)	(0.026)	(25.2)
Brother differential in High − Low	0.012	-0.015	11.1
	(0.017)	(0.030)	(17.6)
Reference Group Mean	-0.052	0.066	83.0
Observations	1,548	1,548	1,488
Respondents	258	258	255

Notes: Medium–Low, High–Medium, and High–Low are equal-weighted over brother status. Reservation dowry are in Rs '000 (2010 prices). Reference group mean reports the mean of the dependent variable for scenario: Low elder match and no brother. All specifications include respondent and scenario-position fixed effects. Standard errors are clustered by respondent. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2 presents the results from estimating (8). Column 1 shows little movement in the resulting acceptance decision. Relative to a low-quality elder match, revealing a medium-quality match raises acceptance by 1.7 percentage points ($p = 0.226$), while revealing a high-quality match lowers it by 1.1 percentage points ($p = 0.542$). Acceptance does fall by 2.8 percentage

points between the medium- and high-quality states ($p = 0.012$), although the magnitude is small. This muted response is consistent with the structure of the task. A more favourable assessment of the bride's family need not translate into greater willingness to accept when it is accompanied by a lower dowry offer. One interpretation of the small but statistically significant High-Medium decline is that the additional improvement in the bride family's perceived quality does not compensate respondents for the accompanying reduction in the offered dowry. We therefore turn to outcomes that capture the reputational value and the corresponding dowry concession it induces, more directly.

The informational response to the elder match is large. Column 2 provides direct support for Prediction 1. Relative to a low-quality elder match, a medium-quality match raises the probability that the bride family is viewed more favourably by 25.4 percentage points ($p < 0.001$), while a high-quality match raises it by 71.7 percentage points ($p < 0.001$). The additional increase from Medium to High is 46.3 percentage points ($p < 0.001$). Groom families therefore treat the elder daughter's realised match as information about the household, and the assessment improves monotonically with elder-match quality. The High–Low response also changes little with brother presence; the estimated differential is -1.5 percentage points ($p = 0.618$). Appendix D, Table 10 reports the full coefficient vector.

This reputational update also changes the terms on which groom families are willing to match. Column 3 corresponds to Prediction 2. Moving from a low- to a medium-quality elder match lowers the groom family's reservation dowry by Rs 80,400 in 2010 prices ($p < 0.001$). Moving from Low to High lowers it by Rs 76,500 ($p = 0.003$). The High–Low concession is similar across brother conditions; the brother differential is Rs 11,100 and is statistically not significant ($p = 0.529$), but is directionally consistent with what the bride's parents anticipate. Groom families are therefore willing to accept a lower transfer from the younger daughter's family after observing a better elder match, as predicted by the model.

The shape of the response is also informative. Family standing continues to improve sharply between the medium- and high-quality elder matches, but reservation dowry does not; the High–Medium difference in reservation dowry is only Rs 3,900 ($p = 0.739$). This is consistent with the truncation in Prediction 2. Once the reservation-transfer floor binds, further improvements in the family's assessment need not generate an additional reduction in the transfer required by a fixed

groom. The data therefore displays precisely the distinction allowed by the model, namely, beliefs about family quality can continue to improve even when the price response has flattened.

The block-level estimates reveal two further patterns (Appendix E, Table 14). First, the reservation-dowry response scales with the monetary stakes of the vignette. Relative to the 500,000 block, the Medium–Low and High–Low concessions are larger by Rs 96,100 ($p = 0.036$) and Rs 118,300 ($p = 0.020$), respectively, in the 1,000,000 block. Second, acceptance responds more strongly to elder-match quality when respondents evaluate a high-quality son. Relative to the medium-quality-son block, the High-Low acceptance response is 14.1 percentage points more positive ($p < 0.001$), with a 7.9 percentage-point differential for Medium–Low ($p = 0.005$). Since the reduction in the offered transfer is fixed by design, this pattern is consistent with high-quality grooms placing greater weight on the favourable information conveyed by the elder sister's match. The reservation-dowry concession itself does not vary with son quality. Thus, we interpret the acceptance result as suggestive evidence of Prediction 3.

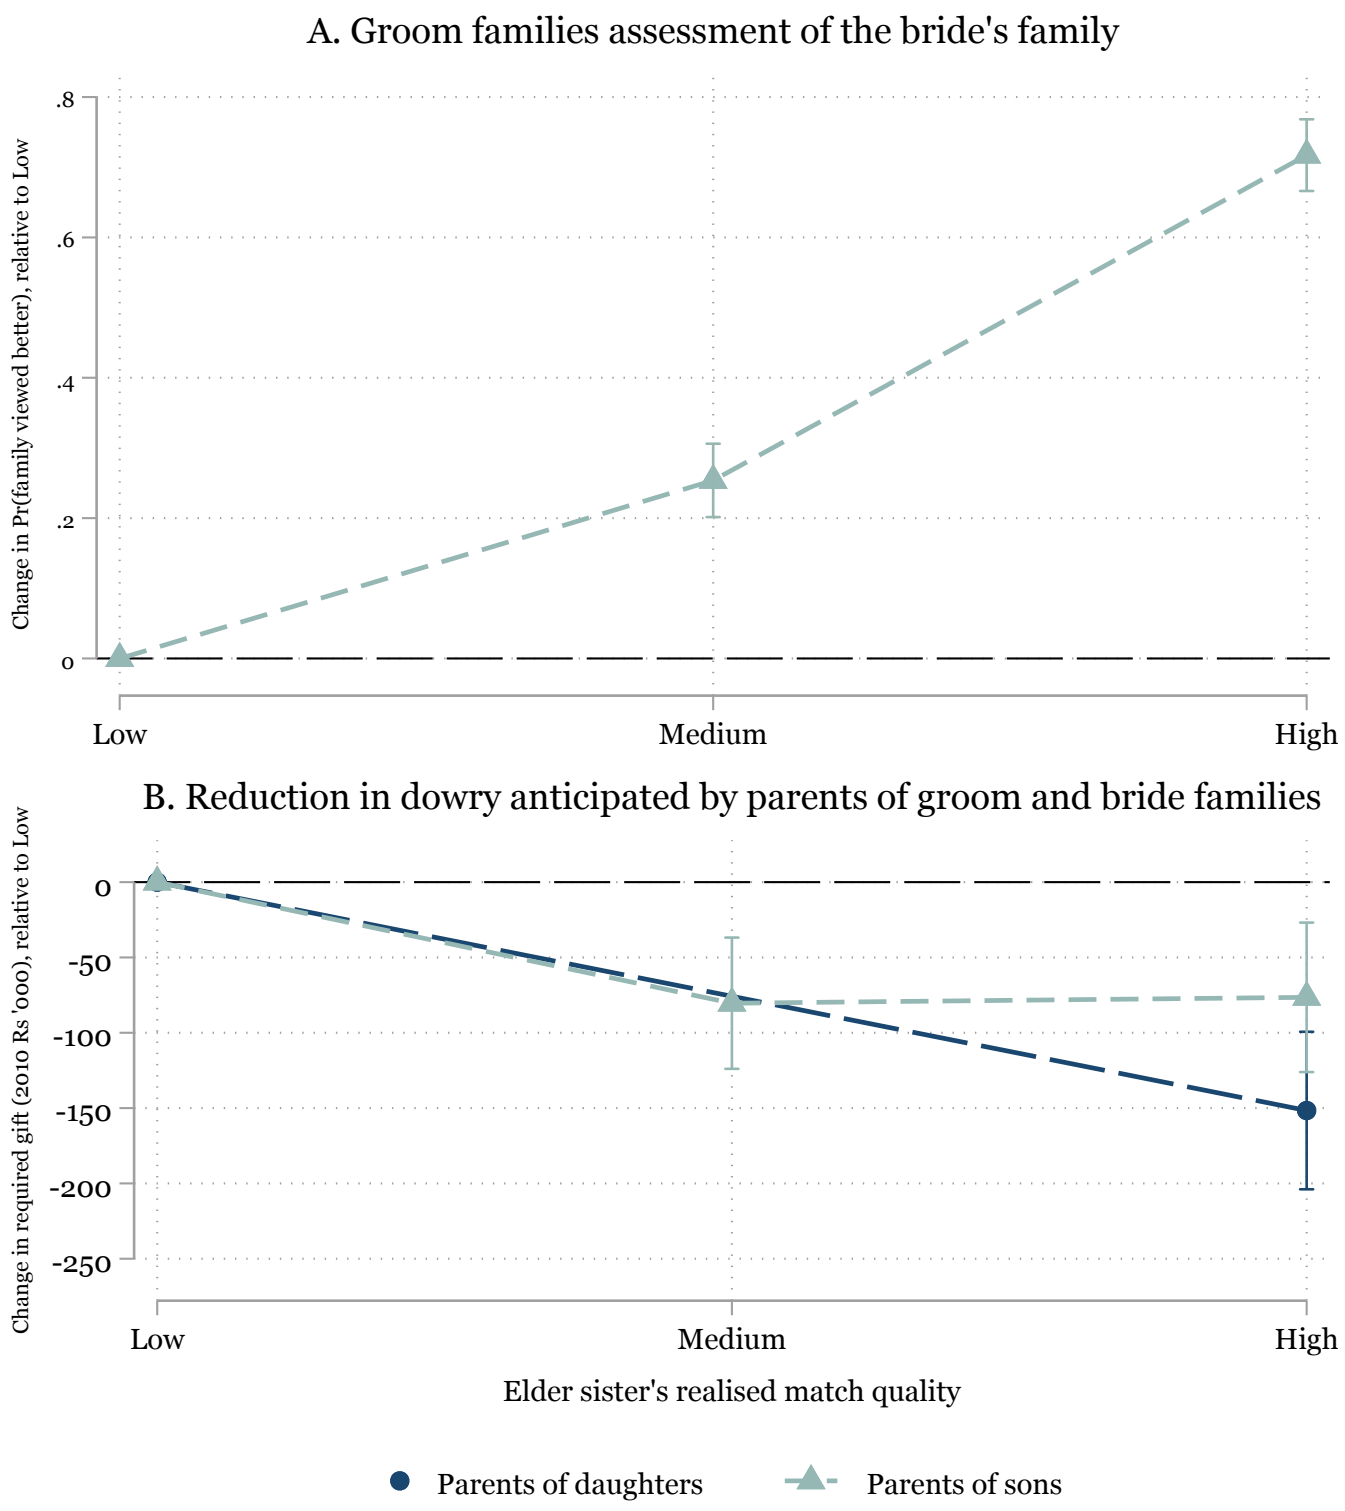
Taken together, the groom-side results show that the linkage perceived by parents of daughters is present in market behaviour as well. A better elder match raises the bride family's perceived standing and lowers the reservation transfer demanded for the younger daughter. The divergence between the smoothly improving family assessment and the flatter reservation-dowry response becomes central in the next subsection, where we compare the market response with the magnitude anticipated by parents of daughters.

6.3 How large is the market response relative to parental beliefs?

The preceding results establish agreement across the two sides of the market on the existence and direction of the cross-sibling linkage. We now ask whether they agree on its magnitude. The conceptual framework provides two useful benchmarks for this comparison. Predictions 1 and 2 allow family assessment to continue improving with elder-match quality even when the associated transfer response becomes flat at the reservation-gift floor. Under common beliefs, the concession anticipated by parents of daughters should also coincide with that granted by parents of sons when they evaluate a comparable groom and family context.

Panel A of Figure 2 first shows the groom-side assessment response across all three elder-match qualities. Family standing therefore continues to improve substantially between the medium-

Figure 2: Family assessment and transfer responses to elder-match quality



Notes: All changes are relative to a Low elder match. The bride instrument does not include a Medium elder-match treatment.

and high-quality matches (25.4 and 71.7 percentage points, respectively, both $p < 0.001$). Panel B shows that the reservation-dowry schedule behaves differently. A medium-quality elder match lowers the groom family's reservation dowry by approximately Rs 80,400 in 2010 prices, while a high-quality match lowers it by Rs 76,500. We cannot reject equality of the Medium–Low and High–Low concessions ($p = 0.738$). This pattern is consistent with the truncation allowed by Prediction 2, namely, family assessment can continue to improve even after further improvements cease to translate into lower transfer requirements.

Panel B also allows us to price the same High–Low change in elder-match quality from opposite sides of the market. Parents of daughters anticipate that moving from a low- to a high-quality elder match will lower the younger daughter's required dowry by approximately Rs 151,600 in 2010 prices (averaged across the age-gap environments not present in groom-module), but parents of sons lower their reservation dowry by only Rs 76,500. The difference is Rs 75,200 and is statistically significant ($p = 0.040$). Parents of daughters anticipate approximately 1.98 times the concession that groom-side parents are willing to grant.

This comparison provides the empirical counterpart of the model's common-beliefs benchmark. For a comparable groom and family context, agreement about how the elder match changes family type implies agreement about the resulting transfer concession. The substantial bride–groom gap therefore indicates that parents of daughters anticipate more pass-through from the elder match into subsequent marriage terms than groom-side parents report. More broadly, this result relates to evidence that individuals can systematically misperceive how others will respond, with such misperceptions affecting behaviour (Bursztyn et al., 2020). In our setting, parents of daughters correctly anticipate the direction of the market response, but substantially overstate its magnitude. The Rs 75,200 wedge is also conservative due to the substantive non-numeric responses. The bride- and groom-side estimate conditions on respondents who provide a numeric transfer. Responses that no realistic gift would secure the high-quality groom are concentrated in the Low elder-match state.

Finally, brother presence provides common variation in the family context z across the two modules. The High–Low transfer concession is smaller in the presence of a brother on both sides of the market, although the groom-side differential is imprecisely estimated. The common direction is suggestive of the mapping from elder-match information into subsequent marriage terms varying

with observable family circumstances. Since the model does not sign these interactions, we treat this pattern as contextual heterogeneity rather than a separate prediction.

6.4 Do parents act on cross-sibling returns?

The preceding modules establish that both sides of the marriage market recognise a cross-sibling return to the elder daughter's match. We next ask whether parents act on this return when allocating marriage resources across daughters. The allocation experiment removes uncertainty about the consequences of each choice. Respondents choose between an equal division of resources and an allocation tilted towards the elder daughter, and are shown the resulting match quality of each daughter under both options. Across scenarios, we vary whether a cross-sister spillover operates, whether the elder daughter suffers from a tight age gap, and whether the household has a contributing son who can help finance better matches for both daughters. Unlike the preceding modules, respondents do not need to infer how the elder daughter's marriage affects her younger sister.

Table 3: Allocation response to a cross-sister return

	Pr(tilt resources to elder)
Spillover	0.304***
	(0.031)
Spillover × Brother	-0.012
	(0.033)
Spillover × Tight gap	0.242***
	(0.033)
Reference Group Mean	0.142
Observations	2,064
Respondents	516

Notes: The outcome equals one when the respondent chooses the allocation tilted toward the elder daughter. The omitted scenario has no cross-sister spillover, no contributing brother, and a wide age gap. Reference Group Mean reports the mean of the dependent variable in that scenario within the estimation sample. The specification includes respondent and scenario-position fixed effects; standard errors are clustered by respondent. The three-way Spillover × Brother × Tight-gap interaction is not included because the complementary half-fraction design does not identify it jointly with respondent fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 3 first asks whether parents respond to the cross-sister return itself, using (9). The presence of a spillover raises the probability of tilting resources towards the elder daughter by 30.4 percentage points ($p < 0.001$). This provides direct support for Prediction 4. Once a better outcome for the elder daughter is known to alter the consequences faced by her younger sister, parents become substantially more willing to depart from equal spending. The response is considerably larger when the sisters are closely spaced. The presence of a tight age-gap penalty in the elder's

match outcome increases the spillover effect by a further 24.2 percentage points ($p < 0.001$). Brother presence changes the spillover response by only −1.2 percentage points ($p = 0.716$). Thus, the strategic response is particularly strong in the tight-gap environment but is essentially unchanged by the presence of a contributing son.

The full estimates are present in Appendix D, Table 11, which also shows that, in the absence of a spillover, moving from a wide to a tight age gap raises the probability of tilting by 15.4 percentage points ($p < 0.001$). Without a brother or spillover, tightening the age gap lowers the elder daughter's match by one category under both equal and tilted allocation while leaving the younger daughter's outcomes unchanged. Parents respond with compensatory allocation once the effect of birth-order pressure, of the type highlighted by Vogl (2013), is made explicit. In the model, sibling spacing forms part of the family context z; the model does not separately sign this direct effect. The spillover response does not significantly vary across the randomised between-respondent blocks and respondent type (whether they were screened into the bride arm or groom arm). Appendix E, Table 15 reports these estimates.

We next use the supplied marriage outcomes themselves to understand what changes parents' willingness to tilt. Table 4 compares all scenarios in which one daughter's outcomes improve while the other's are held fixed. These are cross-cell comparisons using a within-respondent regression. The comparisons are easiest to see with an example. If we move from Cell 1 to Cell 2 (Appendix B.3, Table 7), the elder daughter's outcomes are unchanged as she obtains a "Medium" quality under equal allocation and "Good" quality under the tilted allocation in both cells. The younger daughter receives a "Medium" match quality under equality in both cells, while her outcome under the tilted allocation rises from "Poor" to "Medium". Cells 7 and 8 generate the same comparison one category higher. In each case, the cost imposed on the younger daughter by tilting disappears while the elder daughter's gain is held fixed.

Column 1 shows a very large response. A one-category improvement in the younger daughter's outcome under the tilted allocation, holding her equal-allocation outcome and both elder outcomes fixed, raises the probability of tilting by 56.9 percentage points ($p < 0.001$). This is the most direct empirical counterpart of Prediction 4. Parents are substantially more willing to favour the elder daughter when doing so becomes less costly for her younger sister.

Column 2 considers the converse comparison. We check how willingness to tilt changes when

Table 4: Allocation choices and supplied marriage outcomes

	Elder fixed, Younger better	Elder even better, Younger fixed
Change in Pr(tilt), one-category improvement	0.569***	-0.212***
	(0.023)	(0.020)
Observations	2,064	2,064
Respondents	516	516

Notes: Both columns come from the same within-respondent regression of the tilt choice on the supplied outcome scores. Column 1 reports the change in the probability of tilting from a one-category improvement in the younger daughter's outcome under the tilted allocation, holding her equal-allocation outcome and both elder outcomes fixed. Column 2 reports the change from a one-category improvement in both elder outcomes, holding both younger outcomes fixed. By design, the elder daughter's outcome under the tilted allocation is always one category above her outcome under equal allocation, so a common elder-outcome level is used in Column 2. The specification is linear in the ordered match-quality scores used to construct the survey's total-outcome scores and includes respondent and scenario-position fixed effects. Standard errors are clustered by respondent. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

the elder daughter's overall marriage prospects improve while the younger daughter's outcomes remain unchanged. For intuition, compare Cells 3 and 1 (Appendix B.3, Table 7). In Cell 3, the younger daughter obtains a "Medium" match under equality and a "Poor" match under tilting, while the elder obtains "Poor" and "Medium", respectively. In Cell 1, the younger daughter's outcomes remain "Medium" and "Poor", while the elder's outcomes rise to "Medium" and "Good". The gain to the elder from tilting therefore remains one match-quality category in both cells; what changes is her overall level of marriage prospects. Cells 7 and 5 provide an analogous comparison at a higher outcome level. A one-category improvement in the elder daughter's outcomes under both allocations is associated with a 21.2 percentage-point decline in the probability of tilting ($p < 0.001$). Put differently, parents are substantially less willing to tilt when the choice moves the elder from "Medium" to "Good" rather than from "Poor" to "Medium", even though the gain to the elder remains one category and the younger daughter bears exactly the same loss. This pattern is consistent with compensatory allocation, wherein, parents are more willing to direct additional resources towards the elder daughter when her underlying marriage prospects are worse. Like parents in Giannola (2024), who direct more resources to the child with the higher perceived return, our respondents also allocate strategically to capture the cross-sister return, tilting towards the elder when doing so also benefits the younger. However unlike those parents, our respondents also care about equality, becoming less willing to tilt when equal spending already gives both daughters the same match and tilting would open a big gap between them, in line with the strong value parents place on treating their children equally (Berry et al., 2020).

Taken together, the allocation results show that parents are willing to act on the cross-sibling return once its consequences are made explicit. They are much more willing to tilt resources

towards the elder daughter when that choice ceases to harm the younger sister, as predicted by Prediction 4. Unequal expenditure can therefore arise from parents responding to unequal returns and unequal marriage prospects across daughters, even when they place value on both daughters' outcomes.

6.5 Do elicited beliefs predict strategic allocation?

The final exercise asks whether the beliefs elicited earlier in the survey predict how parents respond when the cross-sibling return is subsequently made explicit. We focus on bride-side respondents as they are the population facing the household allocation problem described in the vignette, namely how to divide scarce marriage resources across daughters. We classify a respondent as a believer when their earlier answers imply that a high-quality elder match lowers the dowry they expect the younger daughter to require. We then ask whether the experimentally supplied spillover in the allocation module has a larger effect among these respondents.

Table 5: Elicited beliefs predict the response to cross-sister returns

	Pr(tilt resources to elder)
Spillover effect, non-believers	0.327***
	(0.040)
Believer differential	0.117**
	(0.053)
Spillover effect, believers	0.444***
	(0.035)
Reference Group Mean	0.270
Observations	1,028
Respondents	257

Notes: Bride-side respondents only. A believer is a respondent whose earlier bride-side elicitation implies a positive High–Low expected concession in the younger daughter's required gift. The Believer differential is the additional effect of the experimentally supplied spillover among believers relative to non-believers; the final row is the implied spillover effect among believers. Reference Group Mean is the mean probability of tilting among non-believers when the spillover is absent, averaged over brother and age-gap environments. The regression controls for brother presence and age gap and includes respondent and scenario-position fixed effects. Standard errors are clustered by respondent. Belief status is elicited before the allocation task and is not randomly assigned. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 5 presents estimates from (11), to show a clear difference. Among respondents whose earlier answers do not imply a positive cross-sister concession, introducing the spillover raises the probability of choosing the elder-tilted allocation by 32.7 percentage points ($p < 0.001$). Among believers, the corresponding effect is 44.4 percentage points ($p < 0.001$). The believer differential is 11.7 percentage points and is statistically significant ($p = 0.027$). This cross-module relationship provides behavioural support for the mechanism linking Predictions 2 and 4. Respondents who previously expected a favourable elder match to reduce the younger daughter's required transfer

subsequently respond more strongly when a return to the younger daughter is built directly into the allocation choice. The belief elicited therefore predicts behaviour in our survey, in the direction implied by the framework. This is consistent with evidence that parents' beliefs about returns govern how they invest in their children (Attanasio et al., 2022; Andrew and Adams, 2025; Calvi et al., 2024).

Appendix D, Table 12 asks whether this relationship also varies with the *magnitude* of the expected concession among believers. It does not appear to. The spillover effect remains large in every quin-tile, ranging from 38.2 to 47.7 percentage points, but there is no systematic gradient with the size of the anticipated concession. The predictive relationship therefore appears primarily at the extensive margin. The block-level estimates in Appendix E, Table 16 show that this believer differential is present more strongly in higher wealth and higher inequality allocation environments.

This result also sharpens the interpretation of an elder-daughter penalty in observed marriage outcomes (Vogl, 2013; Guha et al., 2026). Among parents who already perceive a cross-sister return, a worse elder match cannot be attributed simply to a failure to recognise the value of improving her marriage outcome. These respondents understand the linkage and react more strongly when its consequences are made certain. Any remaining elder-daughter disadvantage in realised marriages may therefore reflect the uncertainty and constraints present in actual marriage decisions, rather than an absence of perceived cross-sibling returns.

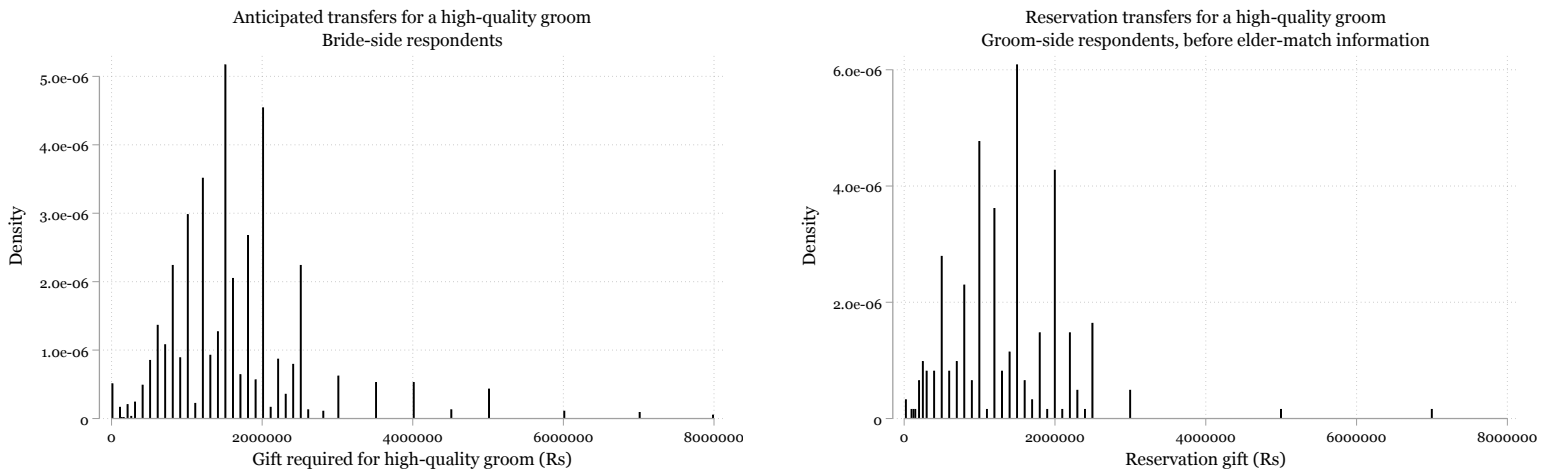
7 Design Validation

We next assess whether the institutions, prices, and constraints embedded in the experimental design correspond to features respondents recognise in their own marriage market. We administered most of these questions after completion of all the experimental modules, so they did not affect their vignette responses. We use them to validate three features of the design. First, that marriage transfers are a familiar and meaningful price-like object; second, that parents are the most relevant decision-makers in our context; and third, that the sibling-composition and birth-order margins varied in the experiment reflect constraints respondents recognise in their own marriage market.

Data collected prior to administration of vignettes show marriage transfers are pervasive in the study population. 99.42 percent of respondents report that some amount of gifts were given

or received at their own marriage, and Section 5 notes that these transfers were economically substantial. Respondents also appear to share a common monetary frame for evaluating current marriage transfers. Figure 3 compares two independently elicited distributions during the vignette exercises. The left panel reports the amount bride-side respondents state that a younger daughter would need to marry a fixed high-quality groom. The right panel reports the minimum transfer groom-side respondents assigned to advise the family of a high-quality son require before receiving any information about the elder sister's match. Despite being elicited from opposite sides of the market and through different questions, the responses occupy a similar range and distribution. We interpret this, along with respondents' own participation in the practice, as evidence that they understand the monetary elicitations as economically meaningful marriage-market prices and draw on a broadly shared understanding of the scale on which such transfers are negotiated.

Figure 3: Marriage-transfer valuations on both sides of the market



Notes: Left panel: shows bride-side responses to the transfer the younger daughter would require to marry a fixed high-quality groom. Right Panel: shows baseline reservation transfers among groom-side respondents assigned to advise the family of a high-quality son, before the elder daughter's match is revealed. All monetary values are in nominal rupees.

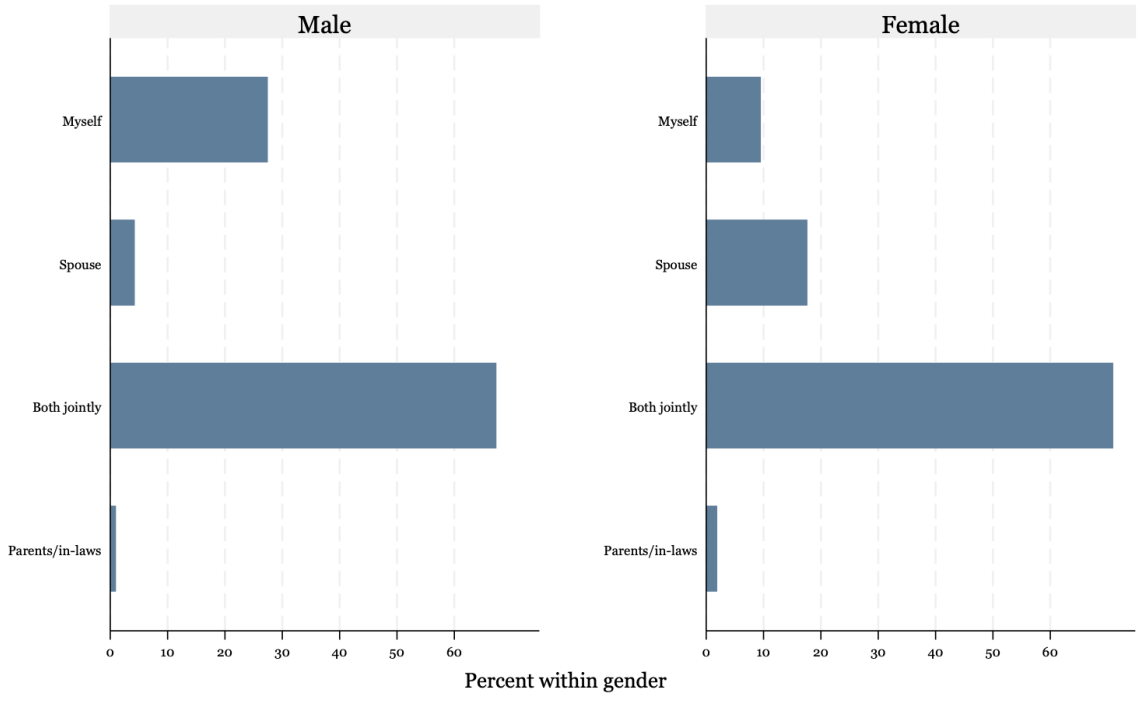
The respondents we survey also appear to be the relevant agents for the decisions studied in the experiment. Figure 4 reports whom respondents say has final authority over children's marriage decisions, separately by respondent gender. For both men and women, the dominant arrangement is joint decision-making of children's marriage by the husband and wife. Most of the remaining responses assign final authority to one of the two spouses, while almost none assign it outside the parental couple. The survey therefore elicits beliefs and choices from individuals who are directly involved in the household decisions represented by the vignettes.

Finally, we collected respondents' opinions on prevailing norms, but we use these as suggestive descriptive evidence, as they were collected after the vignettes were administered and thus may be subject to priming. Figure 5 helps validate that the sibling-related margins varied across the

experimental modules. Birth-order marriage is a near-universal norm in the sample, with roughly 90 percent of respondents reporting that an elder daughter must marry before her younger sister. Closely spaced daughters are also widely perceived to create a binding queue. More than 80 percent of respondents report that a one- to two-year age gap creates pressure to marry the elder daughter quickly, and roughly 70 percent report that this pressure can lower the quality of the elder daughter's eventual match. These responses closely mirror the queueing mechanism represented by the age-gap treatment and motivate the lower elder-match outcomes built into the tight-spacing cells of the allocation experiment.

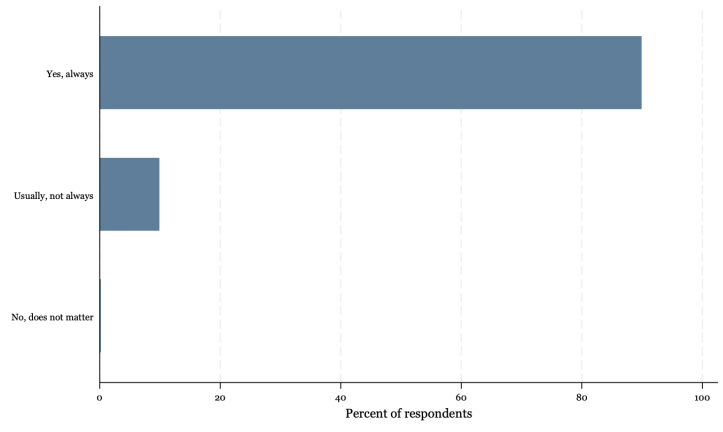
The brother treatment likewise reflects a salient household resource channel. Nearly 90 percent of respondents report that a brother is always expected to contribute to his sister's marriage expenses. This supports the post-pilot decision, outlined in Appendix F.2, to represent a brother in the allocation module as relaxing the household's marriage-resource constraint. Taken together, these descriptive responses show that the experimental variation is built around institutional features the respondents themselves recognise. We therefore interpret the experimental treatments in our design as controlled changes in familiar features of the local marriage environment rather than fully abstract hypothetical attributes.

Figure 4: Parents arrange children's marriages

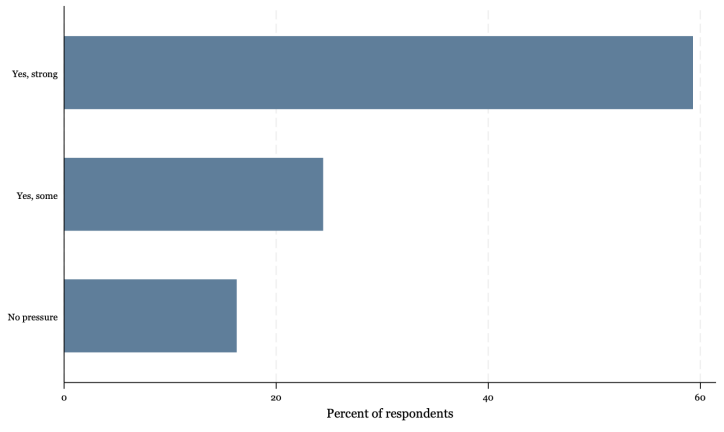


Notes: Bars report the percentage of respondents within each gender reporting who has the final say over marriage decisions for their children.

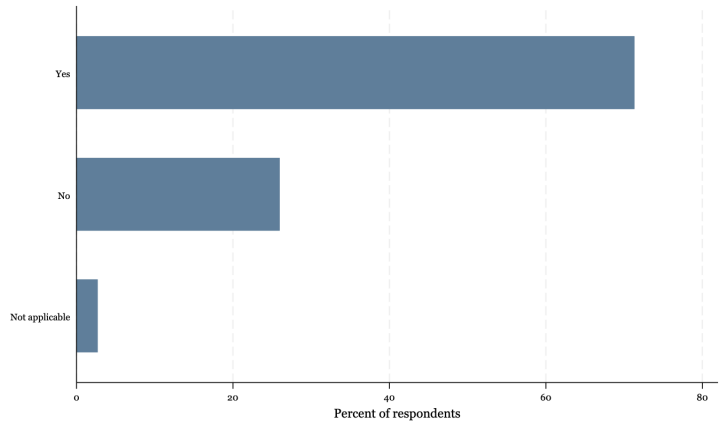
Figure 5: Sibling norms underlying the experimental design



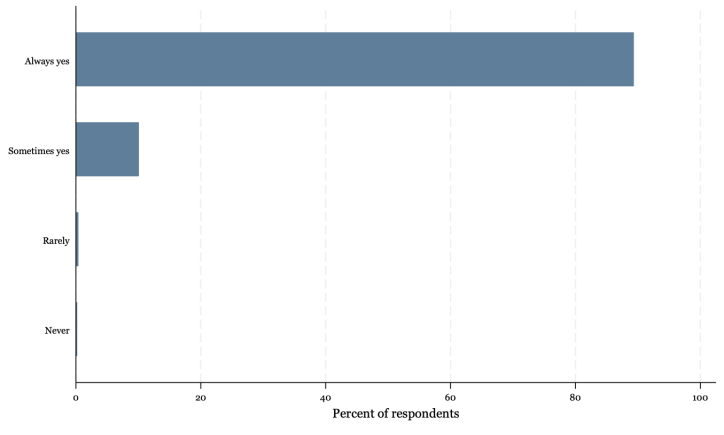
(a) Birth-order norm



(b) Queueing pressure



(c) Consequence of queueing



(d) Brother contribution

8 Conclusion

This paper studies how parents allocate resources across children when an investment in one child changes another child's opportunities. In the marriages of sisters in India, parents perceive a large positive cross-sister return. A better elder match improves the younger daughter's expected marriage prospects and lowers the dowry parents believe is required for a fixed high-quality groom. Parents of potential grooms respond in the same direction, although the concession they report is smaller than the one bride-side parents anticipate. When the consequences of alternative allocations are specified, parents are substantially more willing to tilt resources towards the elder daughter when doing so improves the younger sister's prospects. Together, these results show that parents recognise cross-child returns and are willing to incorporate them into allocation decisions. The findings also change how inequality within families should be interpreted. Unequal spending need not imply unequal parental concern when investment in one child also benefits another. The mechanism extends beyond marriage to any sequential allocations in which an earlier outcome is publicly observed. If others interpret that outcome as information about the allocator, strategically

investing more in the first recipient can improve the opportunities of those who follow.

Yet recognising such a return does not imply that parents can fully realise it. In actual marriage decisions, queueing pressure and liquidity constraints may outweigh the positive cross-sister return, while uncertainty about its magnitude may further limit strategic reallocation. The elder-daughter penalty documented in observational data can therefore coexist with parents who recognise the downstream value of a favourable elder match. Which constraints dominate, and whether relaxing them would allow families to capture this return, remain open questions. The answers matter for policy. A transfer to a household facing scarce marriage resources or closely spaced daughters may do more than simply improve the elder daughter's marriage prospects. If the transfer enables the family to secure a better elder match, it may also improve the terms subsequently available to her younger sister. Such transfers may therefore generate benefits across children that would be missed by evaluating their effects only on the targeted child. More generally, the returns to child-targeted policies depend on how investments, constraints, and opportunities propagate across siblings.

References

AHSAN, M. N. AND U. JAIN (2026): "Brothers, dowry, and spousal quality," *Journal of Population Economics*, 39, 24. [5], [21]

ALMÅS, I., A. W. CAPPELEN, E. Ø. SØRENSEN, AND B. TUNGODDEN (2026): "Fairness across the world," *The Quarterly Journal of Economics*, qjag044. [1]

ALTMEJD, A., A. BARRIOS-FERNÁNDEZ, M. DRLJE, J. GOODMAN, M. HURWITZ, D. KOVAC, C. MULHERN, C. NEILSON, AND J. SMITH (2021): "O Brother, Where Start Thou? Sibling Spillovers on College and Major Choice in Four Countries," *The Quarterly Journal of Economics*, 136, 1831–1886. [5]

ANDERSON, S. (2003): "Why dowry payments declined with modernization in Europe but are rising in India," *Journal of Political Economy*, 111, 269–310. [5]

——— (2007): "The economics of dowry and brideprice," *Journal of economic perspectives*, 21, 151–174. [1], [5], [6]

ANDERSON, S. AND C. BIDNER (2015): "Property Rights Over Marital Transfers," *The Quarterly Journal of Economics*, 130, 1421–1484. [5]

ANDREW, A. AND A. ADAMS (2025): "Revealed Beliefs and the Marriage Market Return to Education," *The Quarterly Journal of Economics*, 140, 2107–2162. [5], [14], [36]

ANUKRITI, S. AND S. DASGUPTA (2017): "Marriage markets in developing countries," Discussion paper, IZA Institute of Labor Economics. [6]

ANUKRITI, S., S. KWON, AND N. PRAKASH (2022): "Saving for dowry: Evidence from rural India," *Journal of Development Economics*, 154, 102750. [1], [5], [21]

ATTANASIO, O., T. BONEVA, AND C. RAUH (2022): "Parental beliefs about returns to different types of investments in school children," *Journal of Human Resources*, 57, 1789–1825. [36]

ATTANASIO, O. P. AND K. M. KAUFMANN (2017): "Education Choices and Returns on the Labor and Marriage Markets: Evidence from Data on Subjective Expectations," *Journal of Economic Behavior & Organization*, 140, 35–55. [5]

BANERJEE, A., E. DUFLO, M. GHATAK, AND J. LAFORTUNE (2013): "Marry for what? Caste and mate selection in modern India," *American Economic Journal: Microeconomics*, 5, 33–72. [5]

BEAUCHAMP, A., R. CALVI, AND S. FULFORD (2023): "Terms of Engagement: Migration, Dowry, and Love in Indian Marriages," Working paper, september. [5]

BECKER, G. (1981): "Treatise on the Family Harvard University Press Cambridge Mass," . [5]

BECKER, G. S. AND N. TOMES (1976): "Child endowments and the quantity and quality of children," *Journal of political Economy*, 84, S143–S162. [1]

BEHRMAN, J. R., R. A. POLLAK, AND P. TAUBMAN (1982): "Parental preferences and provision for progeny," *Journal of political economy*, 90, 52–73. [1]

BEHRMAN, J. R., M. R. ROSENZWEIG, AND P. TAUBMAN (1994): "Endowments and the allocation of schooling in the family and in the marriage market: the twins experiment," *Journal of political economy*, 102, 1131–1174. [4]

BERRY, J., R. DIZON-ROSS, AND M. JAGNANI (2020): "Not playing favorites: An experiment on parental fairness preferences," Tech. rep., National Bureau of Economic Research. [1], [4], [8], [14], [34]

BHARADWAJ, P., J. P. EBERHARD, AND C. A. NEILSON (2018): "Health at Birth, Parental Investments, and Academic Outcomes," *Journal of Labor Economics*, 36. [4]

BHATTARAI, J., R. CALVI, AND A. KESKAR (2026): "Gender Gap in Child Poverty: Evidence from India," Working paper, september. [4]

BLACK, S. E., P. J. DEVEREUX, AND K. G. SALVANES (2005): "The more the merrier? The effect of family size and birth order on children's education," *The Quarterly Journal of Economics*, 120, 669–700. [4]

BLOCH, F. AND V. RAO (2002): "Terror as a bargaining instrument: A case study of dowry violence in rural India," *American Economic Review*, 92, 1029–1043. [2], [6], [8]

BLOCH, F., V. RAO, AND S. DESAI (2004): "Wedding celebrations as conspicuous consumption: signaling social status in rural India," *Journal of Human Resources*, 39, 675–695. [2], [8]

BORKER, G., J. EECKHOUT, N. LUKE, S. MINZ, K. MUNSHI, AND S. SWAMINATHAN (2026): "Wealth, marriage, and sex selection," *American Economic Journal: Applied Economics*, 18, 107–146. [5]

BOTTICINI, M. AND A. SIOW (2003): "Why dowries?" *American Economic Review*, 93, 1385–1398. [5]

BROWN, C., R. CALVI, AND J. PENGLASE (2021): "Sharing the Pie: An Analysis of Undernutrition and Individual Consumption in Bangladesh," *Journal of Public Economics*, 200, 104460. [4]

BURSZTYN, L., A. L. GONZÁLEZ, AND D. YANAGIZAWA-DROTT (2020): "Misperceived social norms: Women working outside the home in Saudi Arabia," *American Economic Review*, 110, 2997–3029. [31]

CALVI, R., H. FAROOQI, AND E. KANDPAL (2024): "The Perceived Marital Returns to Education and the Demand for Girls' Schooling," Tech. rep., JSTOR. [5], [14], [36]

CALVI, R. AND A. KESKAR (2023): "'Til Dowry Do Us Part: Bargaining and Violence in Indian Families," *Review of Economics and Statistics*, 1–45. [6]

CHIPLUNKAR, G. AND J. WEAVER (2023): "Marriage markets and the rise of dowry in India," *Journal of Development Economics*, 164, 103115. [1], [5], [6]

CHO, I.-K. AND D. M. KREPS (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102, 179–221. [2]

GIANNOLA, M. (2024): "Parental investments and intra-household inequality in child human capital: Evidence from a survey experiment," *The Economic Journal*, 134, 671–727. [4], [14], [34]

GUHA, A., A. KESKAR, AND S. MOOKERJEE (2026): "Waiting Your Turn: Sibling Rivalry, Birth Order, and Dowry in Rural India," *Working Paper, R&R at Labour Economics*. [1], [4], [5], [7], [21], [36]

JACOB, S. AND S. CHATTOPADHYAY (2016): "Marriage dissolution in India: Evidence from census 2011," *Economic and Political Weekly*, 51, 25–27. [6]

JAKIELA, P., O. OZIER, L. C. H. FERNALD, AND H. A. KNAUER (2025): "Big Sisters," *Journal of Human Resources*, published online. [5]

JAYACHANDRAN, S. AND I. KUZIEMKO (2011): "Why Do Mothers Breastfeed Girls Less Than Boys? Evidence and Implications for Child Health in India," *The Quarterly Journal of Economics*, 126, 1485–1538. [4]

JAYACHANDRAN, S. AND R. PANDE (2017): "Why are Indian children so short? The role of birth order and son preference," *American Economic Review*, 107, 2600–2629. [4]

JONES, G. W. (2017): "Changing marriage patterns in Asia," in *Routledge handbook of Asian demography*, Routledge, 351–369. [23]

KARVE, I. K. (1965): "Kinship organization in India," *(No Title)*. [23]

KHALIL, U. AND S. MOOKERJEE (2019): "Patrilocal residence and women's social status: evidence from South Asia," *Economic Development and Cultural Change*, 67, 401–438. [6]

KOLENDA, P. (1987): *Regional differences in family structure in India*, Rawat Publications. [23]

MANSKI, C. F. (2004): "Measuring expectations," *Econometrica*, 72, 1329–1376. [7], [14]

NCAER NATIONAL DATA INNOVATION CENTRE (n.d.): "How Much Agency Do Indian Women Have in Marriage Choices?" India Human Development Survey research summary, reports evidence from IHDS 2011–12. [1]

NICOLETTI, C. AND B. RABE (2019): "Sibling Spillover Effects in School Achievement," *Journal of Applied Econometrics*, 34, 482–501. [5]

PENGLASE, J. (2021): "Consumption Inequality Among Children: Evidence from Child Fostering in Malawi," *The Economic Journal*, 131, 1000–1025. [4]

QURESHI, J. A. (2018): "Additional Returns to Investing in Girls' Education: Impact on Younger Sibling Human Capital," *The Economic Journal*, 128, 3285–3319. [5]

RAO, V. (1993): "The rising price of husbands: A hedonic analysis of dowry increases in rural India," *Journal of political Economy*, 101, 666–677. [6]

ROSENZWEIG, M. R. AND T. P. SCHULTZ (1982): "Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural India," *The American Economic Review*, 72, 803–815. [4]

RUBIO, G. (2014): "How love conquered marriage: Theory and evidence on the disappearance of arranged marriages," *Unpublished manuscript*. [1], [6]

SRINIVASAN, S. AND A. S. BEDI (2007): "Domestic violence and dowry: Evidence from a South Indian village," *World Development*, 35, 857–880. [6]

SRIVASTAVA, S., S. CHAUHAN, R. PATEL, S. P. MARBANIANG, P. KUMAR, R. PAUL, AND P. DHILLON (2021): "Banned by the law, practiced by the society: The study of factors associated with dowry payments among adolescent girls in Uttar Pradesh and Bihar, India," *PLoS ONE*, 16, e0258656. [23]

VOGL, T. S. (2013): "Marriage institutions and sibling competition: Evidence from South Asia," *The Quarterly journal of economics*, 128, 1017–1072. [1], [4], [5], [7], [21], [33], [36]

WISWALL, M. AND B. ZAFAR (2018): "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, 133, 457–507. [7]

——— (2021): "Human Capital Investments and Expectations about Career and Family," *Journal of Political Economy*, 129, 1361–1424. [5]

A Proofs for the Conceptual Framework

This appendix establishes the results used in Section 3. Throughout, parent and public-context subscripts are suppressed when they are held fixed. The ordered groom set is finite, $\mathcal{G} = \{g_1 < \cdots < g_J\}$, and at least one groom quality is affordable in every relevant state.

A.1 Competitive continuation

For fixed elder expenditure t, recall

$$h_t(d) = d + \rho|t - d|, \qquad 0 \le \rho < 1.$$

The function is convex. Away from its kink at $d = t$,

$$h_t'(d) = \begin{cases} 1-\rho, & d < t, \\ 1+\rho, & d > t. \end{cases}$$

Both slopes are strictly positive, so h_t is strictly increasing.

Lemma 1. For fixed (t, B), the continuation value $V(t, B, \mu)$ is weakly increasing in μ.

Proof. From (2),

$$D_g(\mu, \omega) = \max\{\underline{d}_g, a_g(\omega) - v_g\mu\}$$

is weakly decreasing in μ because $v_g \ge 0$. Hence any groom quality feasible at μ remains feasible at any $\mu' \ge \mu$. For every previously feasible quality, the required transfer weakly falls; because h_t is increasing, its effective cost weakly falls. The maximised objective in (3) therefore weakly rises state by state. Taking expectations over ω proves

$$V(t, B, \mu') \ge V(t, B, \mu) \qquad \text{for } \mu' \ge \mu.$$

□

Lemma 2. Suppose $t_H \geq t_L$ and $B_m = W - t_m$. For any fixed posterior μ,

$$V(t_L, B_L, \mu) \geq V(t_H, B_H, \mu). \tag{12}$$

Proof. Since $t_H \geq t_L$, we have $B_L \geq B_H$. Every groom feasible after spending t_H on the elder is therefore also feasible after spending t_L. For any such groom requiring transfer d,

$$\begin{aligned}
[B_L - h_{t_L}(d)] - [B_H - h_{t_H}(d)] &= (t_H - t_L) - \rho\,(|t_L - d| - |t_H - d|) \\
&\geq (t_H - t_L) - \rho |t_H - t_L| \\
&= (1-\rho)(t_H - t_L) \geq 0,
\end{aligned}$$

where the first inequality uses the reverse triangle inequality. Thus every quality feasible after the larger elder transfer yields weakly greater continuation utility after the smaller one, and additional qualities may also become feasible. Maximising state by state and taking expectations gives (12). □

Lemma 2 implies that

$$V(t_L, B_L, 0) - V(t_H, B_H, 0) \geq 0,$$

so the second term in (5) is a genuine resource cost.

A.2 Proof of Proposition 1

The gain to type θ from choosing H rather than L is

$$V(t_H, B_H, \mu_H) - V(t_L, B_L, \mu_L) - k_\theta. \tag{13}$$

By Lemma 1, the largest possible continuation gain from H occurs at $\mu_H = 1$ and $\mu_L = 0$, and equals Γ. Since $k_0 > \Gamma$ under (4), type $\theta = 0$ strictly prefers L for every possible pair of posterior beliefs.

Suppose type $\theta = 1$ chooses H with positive probability. Because type 0 never chooses H, Bayes' rule then implies $\mu_H = 1$. The posterior following L is no greater than the prior, so $\mu_L \leq p$.

Lemma 1 gives

$$V(t_H, B_H, 1) - V(t_L, B_L, \mu_L) \geq V(t_H, B_H, 1) - V(t_L, B_L, p) = \Gamma^p > k_1.$$

Type 1 therefore strictly prefers H and cannot mix. Any equilibrium in which H occurs consequently has the separating sender outcome.

Pooling on H cannot be an equilibrium because type 0 strictly prefers L regardless of off-path beliefs. Consider pooling on L, for which the on-path posterior is p. If a deviation to H is attributed to type 1, that type obtains the continuation gain

$$V(t_H, B_H, 1) - V(t_L, B_L, p) = \Gamma^p > k_1$$

and therefore prefers to deviate. Type 0 cannot profit from the same deviation under any posterior because $k_0 > \Gamma$. The intuitive criterion therefore attributes the deviation to type 1 and rules out pooling on L. The separating sender outcome is the unique outcome satisfying the criterion, with $\mu(H) = 1$ and $\mu(L) = 0$. □

A.3 Partial sorting and Prediction 1

Full separation is stronger than is needed for Prediction 1. Let

$$q_1 = \Pr(H \mid \theta = 1), \qquad q_0 = \Pr(H \mid \theta = 0),$$

and suppose $q_1 > q_0$. When both signals occur with positive probability, Bayes' rule gives

$$\Pr(\theta = 1 \mid H) = \frac{pq_1}{pq_1 + (1-p)q_0}$$

and

$$\Pr(\theta = 1 \mid L) = \frac{p(1-q_1)}{p(1-q_1) + (1-p)(1-q_0)}.$$

Cross-multiplication shows

$$\Pr(\theta = 1 \mid H) > \Pr(\theta = 1 \mid L) \quad \Longleftrightarrow \quad q_1 > q_0.$$

Thus any strict positive sorting of the favourable type into the high elder match is sufficient for $\mu(H) > \mu(L)$.

A.4 Prediction 2: transfers for a fixed groom

Fix groom quality g and opportunity state ω, and let $\mu_H > \mu_L$. Define

$$x_L = a_g(\omega) - v_g \mu_L, \qquad x_H = a_g(\omega) - v_g \mu_H = x_L - v_g(\mu_H - \mu_L).$$

Then

$$D_g(\mu_L, \omega) = \max\{\underline{d}_g, x_L\}, \qquad D_g(\mu_H, \omega) = \max\{\underline{d}_g, x_H\}.$$

There are three cases.

If $x_L \leq \underline{d}_g$, the floor binds already at the low posterior and both required transfers equal $\underline{d}_g$. The difference is zero.

If $x_H \geq \underline{d}_g$, the floor binds at neither posterior and

$$D_g(\mu_H, \omega) - D_g(\mu_L, \omega) = x_H - x_L = -v_g(\mu_H - \mu_L).$$

Finally, if $x_H < \underline{d}_g < x_L$, the high posterior reaches the floor while the low posterior does not, giving

$$D_g(\mu_H, \omega) - D_g(\mu_L, \omega) = \underline{d}_g - x_L = -[x_L - \underline{d}_g].$$

Combining the three cases yields

$$D_g(\mu_H, \omega) - D_g(\mu_L, \omega) = -\min\left\{v_g(\mu_H - \mu_L), [a_g(\omega) - v_g \mu_L - \underline{d}_g]_+\right\} \leq 0,$$

which is what we had in Prediction 2. Once the floor binds, the local derivative of the required transfer with respect to μ is zero even though the posterior itself may continue to improve. □

A.5 Prediction 3: younger-match prospects

Fix (t, B, ω) and two posteriors $\mu_H > \mu_L$. Write

$$D_g^L = D_g(\mu_L, \omega), \qquad D_g^H = D_g(\mu_H, \omega),$$

and define the information concession

$$C_g = D_g^L - D_g^H \geq 0.$$

Assume that D_g^L and D_g^H are weakly increasing in groom quality and that C_g is also weakly increasing in groom quality.

The payoff gain generated by moving from μ_L to μ_H while keeping groom quality fixed is

$$R_g = h_t(D_g^L) - h_t(D_g^H).$$

We first show that R_g is weakly increasing in g. Consider $g' > g$. By assumption,

$$D_{g'}^H \geq D_g^H, \qquad D_{g'}^L \geq D_g^L, \qquad D_{g'}^L - D_{g'}^H \geq D_g^L - D_g^H.$$

Because h_t is increasing and convex, the increment

$$h_t(a + \ell) - h_t(a)$$

is weakly increasing in both the starting point a and the interval length $\ell \geq 0$. Setting $a = D_g^H$ and $\ell = C_g$ therefore gives

$$R_{g'} \geq R_g.$$

Now let

$$Q_g^L = (1 - \alpha)u(g) + B - h_t(D_g^L)$$

denote the payoff from groom g under the low posterior. Under the high posterior,

$$Q_g^H = Q_g^L + R_g.$$

Let $g_L = g^*(t, B, \mu_L, \omega)$. Since required transfers weakly fall when the posterior improves, g_L remains feasible under μ_H. For any $g < g_L$,

$$Q^H_{g_L} - Q^H_g = (Q^L_{g_L} - Q^L_g) + (R_{g_L} - R_g) \geq 0,$$

because g_L maximises the low-posterior problem and R_g is weakly increasing in quality. Hence no lower-quality groom can strictly overtake g_L after the posterior improves. With ties resolved in favour of the highest maximiser,

$$g^*(t, B, \mu_H, \omega) \geq g^*(t, B, \mu_L, \omega) \qquad \text{for every } \omega.$$

Taking probabilities over ω gives

$$F^*(\bar{g} \mid t, B, \mu_H) \leq F^*(\bar{g} \mid t, B, \mu_L)$$

for every cutoff $\bar{g}$, proving Prediction 3. □

If the transfer floor does not bind for the relevant groom qualities, then

$$C_g = \nu_g(\mu_H - \mu_L),$$

so ν_g weakly increasing in groom quality is sufficient for the concession-ordering assumption. When floors bind, the ordering of the realised concessions is the relevant restriction because truncation can otherwise differ across groom qualities.

With three ordered groom categories, first-order stochastic dominance implies that the probability of the lowest category weakly falls and the cumulative probability below the highest category weakly falls. It does not separately sign the probability assigned to the middle category.

A.6 Prediction 4: allocation with known outcomes

Consider equal spending

$$E = \left(\frac{W}{2}, \frac{W}{2}\right)$$

and elder-tilted spending

$$T = (rW, (1-r)W), \qquad r > \frac{1}{2}.$$

Both allocations exhaust the marriage budget, so retained resources cancel in their utility difference. Under equal spending, the expenditure-equality term in (1) is zero. Under tilted spending,

$$\frac{|rW-(1-r)W|}{W} = 2r-1.$$

Let

$$\Delta u_{ji} = u_i(m_j^T; z) - u_i(m_j^E; z), \qquad \Delta c_i = c_i^T - c_i^E.$$

Subtracting $U_i(E; z)$ from $U_i(T; z)$ gives

$$A_i(z) = \alpha_i(z)\Delta u_{1i} + [1-\alpha_i(z)]\Delta u_{2i} - \eta_i(z)(2r-1) - \Delta c_i,$$

which is what we had outlined in Prediction 4.

Holding family context, elder outcomes, transfers, preferences, and Δc_i fixed, an improvement in the younger daughter's match under T weakly raises Δu_{2i} because $u_i(\cdot; z)$ is increasing. Since $1-\alpha_i(z) \geq 0$, this weakly raises $A_i(z)$. Worsening the younger daughter's match under E has the same effect. The direct expenditure-equality penalty is unchanged because the spending alternatives do not change.

If observed choice is generated by

$$T \text{ chosen} \quad \Longleftrightarrow \quad A_i(z) + \varepsilon_i \geq 0,$$

with a choice-error distribution that is stable across the compared scenarios, then the probability of choosing tilt is weakly increasing in $A_i(z)$. The same ordering therefore carries over to choice probabilities. □

If the younger daughter's outcomes improve under both E and T, monotonicity of u_i alone does not order the two utility gains. Prediction 4 therefore concerns changes in the younger daughter's *relative* consequence under tilted rather than equal spending.

B Vignettes

B.1 Bride Module Vignettes

The bride-side module varies two dimensions across respondents and three dimensions within respondent. Household marriage resources are either Rs.5,00,000 or Rs.10,00,000, and the share required for the elder daughter's marriage is either 60 or 70 percent; both are fixed throughout a respondent's module. Within respondent, we vary the elder daughter's realised match between a high-quality groom and a low-quality groom, the age gap between the sisters between two and five years, and whether the household has a son. Each respondent therefore evaluates eight scenarios, corresponding to the $2 \times 2 \times 2$ within-respondent design.

The age-gap and sibling-composition treatments define four family types. Respondents see both elder-match treatments for a given family before moving to the next family. The order of the four family types is assigned using a Williams Latin square, which balances the position of each family type and first-order carryover across family types. Within each family type, the order of the high- and low-quality elder matches is alternated. Crossing the four family orders with the two elder-match orders gives eight presentation sequences, assigned approximately evenly across respondents.

Survey wording. The following reproduces the wording of the module. Quantities in square brackets were populated according to the respondent's experimental assignment.

> *"I will now describe 4 hypothetical families from this area. These are not real families. The family has lifetime savings **Rs. [W]** for their daughters'marriages. In each scenario, we will describe a marriage proposal for the elder daughter of that family. Accepting the match requires giving gifts of a certain value to the groom's family. Examples of gifts may include a car, land, clothes, household items, jewellery etc. You will be advising the parents of the daughter in each scenario – you will answer 4 questions. There is no requirement to accept any match unless you think it is the right decision."*

Scenarios.

> *"A family in your area has [only two daughters and no son / two daughters and one son]. The daughters are [2 / 5] years apart."*

Figure 6: Family-composition image, bride module



Notes: The specific image would change to represent the described family scenario.

*'The family has lifetime savings **Rs. [W]** for their daughters' marriages. The elder daughter is of marriage-able age. The younger daughter is not yet married, and her future marriage outcome is uncertain.*"

*'A groom's family has proposed to the elder daughter. Accepting requires giving gifts of value **Rs. [D]** (**[60 / 70]%** of savings), leaving **Rs. [W-D] or [40 / 30]%** for the younger. The groom is:"*

Figure 7: Proposed elder groom, bride module



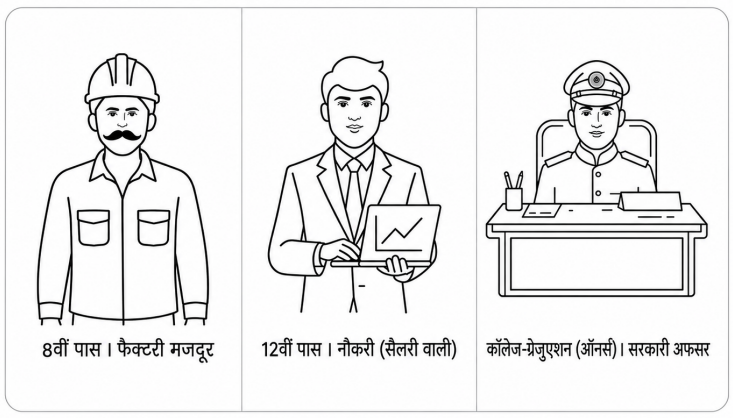
कॉलेज-ग्रेजुएशन (ऑनर्स)। सरकारी अफसर

Notes: This describes the high-quality groom who is a college graduate and working as a Government officer. This changes to the low-quality groom visual depending on the scenario.

Q1. *"Would you advise this family to accept this proposal, at the given cost, for the elder daughter?"*

Q2. *"Imagine the elder daughter ends up marrying the proposed groom. Purely based on where the elder daughter has married, where do you think the parents will be able to arrange a match for the younger daughter? Imagine each of these beans is a proposal for the younger; please divide them over each of the following types of groom based on likelihood."*

Figure 8: Bean-allocation card



Q3. *"Imagine the elder daughter ends up marrying the proposed groom. Purely based on where the elder daughter has married, how do you think families with sons looking for brides will view this family when the parents go to look for a groom for the younger daughter?"*

Much worse Worse No different Better Much better

Q4. *"Imagine the elder daughter ends up marrying the proposed groom. Purely based on where the elder daughter has married, what value of gifts (in rupees) would the younger need to marry a groom of* ***high quality – a college graduate (Hons.) and government officer****?"*

B.2 Groom Module Vignettes

The groom-side module varies two dimensions across respondents and two dimensions within respondent. Respondents are assigned either a high-quality son (college graduate (Hons.), government officer) or a medium-quality son (Class 12 pass, private-sector worker), and household marriage wealth of the bride's family (from which they make offers to the groom) is fixed at either Rs. 5,00,000 or Rs. 10,00,000 throughout the module. The initial gift offer is one-half of household marriage wealth. Within respondent, we vary whether the proposing family has a son and, after an initial baseline elicitation, whether the elder daughter has married a high-, medium-, or low-quality groom. Each respondent therefore evaluates two baseline families and six updated scenarios, for eight scenarios in total.

The scenarios are presented in family blocks. Respondents first evaluate one family before the elder daughter's match is revealed and then see all three possible elder-match updates for that same family. They subsequently repeat this sequence for the second family. The order of

the brother and no-brother families is counterbalanced across respondents. Within each family, the High, Medium, and Low elder-match updates are presented according to one of six balanced orders, with the second family using the reverse sequence. The resulting twelve presentation orders balance the position of each family type and elder-match treatment while avoiding repeated switching between family contexts.

In the updated scenarios, the gift offered by the bride's family also varies with the elder match. Relative to the baseline offer of $0.5W$, the family offers $0.3W$ following a high-quality elder match, $0.4W$ following a medium-quality elder match, and $0.7W$ following a low-quality elder match. Before each updated scenario, respondents are shown their own acceptance and minimum-gift answers from the corresponding baseline family.

Survey wording. The following reproduces the wording of the module. Quantities and descriptions in square brackets were populated according to the respondent's experimental assignment.

> *"You will be advising the parents of a son who is **[son description]**. I will present 2 hypothetical families from this area. Each of the families is coming with a marriage proposal. Every marriage proposal comes with the offer of gifts for the groom and his family. Examples of gifts may include a car, land, clothes, household items, and jewellery etc. We will ask how likely you are to recommend accepting or rejecting the marriage proposal. Later, I will share some new information about the family and ask you the same questions again. I will remind you of your earlier answers so you can see if and how your views have changed. You are always free to recommend rejection in every scenario and none of these are real families – there are no right or wrong answers."*

Baseline scenario. Respondents are first told:

> *"You are advising the parents of a son who is **[son description]**."*
>
> *'A family in another village with **[only two daughters and no sons / two daughters and one son]** has proposed their younger daughter for marriage with this son. The elder daughter has already married. The value of gifts being offered for the younger is **Rs.** [$0.5W$]."*

Figure 9: Son being advised, groom module



Notes: The image shows the Medium-quality son being advised, who has a Class 12 education and works in the private sector. This image changes across respondents (by block assignment).

They then answer:

Q1. *"On a scale of 0 to 10, how likely are you to advise this son's family to accept this match, based on what you know so far about this family? (0 = definitely reject, 10 = definitely accept)"*

Q2. *"What is the lowest value of gifts at which you would advise this son's family to say yes to this match?"*

Updated scenario. Before each update, respondents are reminded of their own answers for the corresponding baseline family:

Your earlier answer for this family:

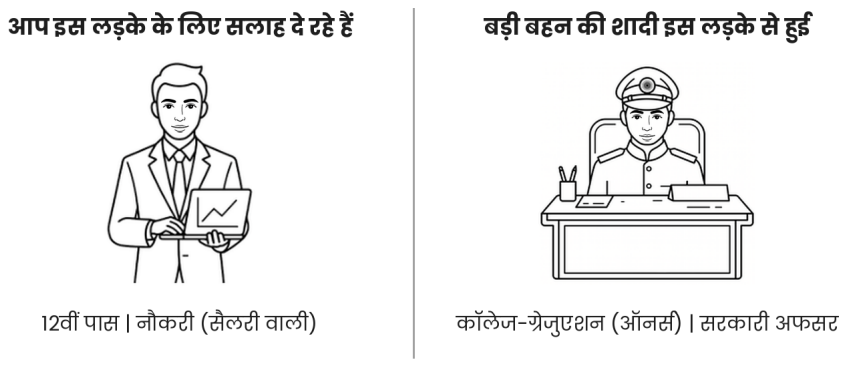
Family with [family composition], value of gifts offered Rs. [0.5W].

You said: Likelihood = [baseline answer]/10 Lowest acceptable gift offer = Rs. [baseline answer].

The elder daughter's realised match is then revealed:

*"Now imagine you found out, the elder daughter of the same family had married a groom who is a **[High / Medium / Low groom description]**. But now the family is offering gifts worth **Rs.** [0.3W / 0.4W / 0.7W] for marrying the younger daughter."*

Figure 10: Elder-match update screen, groom module



Notes: The image shows the Medium-quality son being advised, who the elder has married (here high quality) on the right - which changes according to the scenario.

Respondents then answer:

Q1. *"Now that you know* ***[the elder sister married well / the elder sister married someone of average quality / the elder sister married poorly]****, on a scale of 0–10, how likely are you to advise this son's family to accept this match,* ***given the new offer****?"*

Q2. *"What is the lowest value of gifts at which you would advise this son's family to say yes to this match?"*

Q3. *"Because the elder sister married this groom, do you now view this household as a better family, a worse family, or no different to marry this son into?"*

Much worse *Worse* *No different* *Better* *Much better*

B.3 Allocation Module Vignettes

The allocation module contains a full $2 \times 2 \times 2$ design varying the presence of a cross-sister return, the age gap between the daughters, and the presence of a contributing brother. Total household marriage wealth and the size of the elder-tilted split are fixed within respondent and vary across respondents. Household marriage wealth is either Rs. 5,00,000 or Rs. 10,00,000, and the tilted allocation assigns either 60 or 70 percent of these resources to the elder daughter, compared with an equal 50–50 allocation.

The three experimental dimensions generate the eight outcome bundles reported in Appendix Table 7. Respondents do not evaluate all eight cells. Each respondent evaluates four scenarios, one from each spillover-by-age-gap combination. Brother presence is assigned using complementary patterns such that each respondent sees two scenarios with a contributing brother and two without, once with and once without a spillover and once under each age-gap condition. Across respondents, this produces balanced coverage of all eight cells.

The order of the four assigned scenarios is counterbalanced using a Williams Latin square, balancing scenario position and first-order carryover. The labels "Option A" and "Option B" are also counterbalanced within respondent: each respondent sees the equal allocation as Option A twice and the tilted allocation as Option A twice.

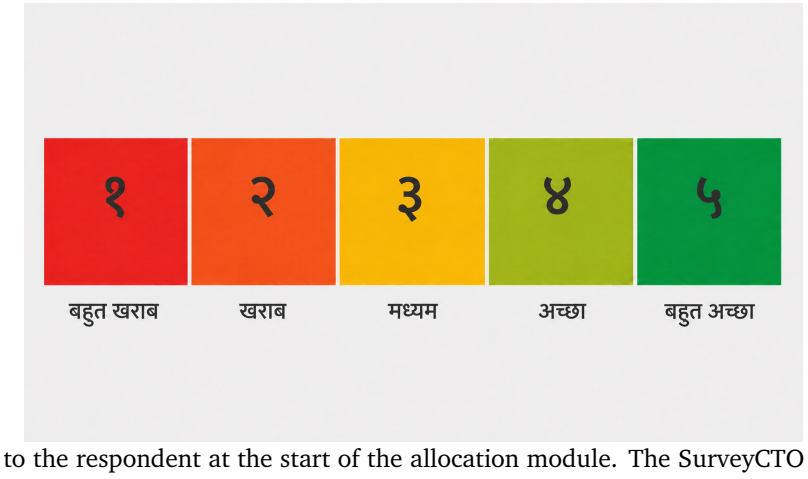
Respondents are shown the five-point match-quality scale in Appendix Figure 11. The match-quality level is accompanied by the corresponding education and occupation description outlined in Table 6.

Table 6: Match-quality scale used in the allocation task

Level	Label (Score)	Education	Job
+2	V. GOOD (5)	College graduate	Government officer
+1	GOOD (4)	Diploma	Government clerk
0	MEDIUM (3)	Class 12 pass	Private-sector worker
−1	POOR (2)	Class 8 pass	Labourer
−2	V. POOR (1)	Class 5 pass	Unemployed

Notes: Respondents see labels for a five-point match-quality scale alongside the corresponding education and occupation descriptions.

Figure 11: Match-quality card, allocation module



Notes: This card was handed to the respondent at the start of the allocation module. The SurveyCTO screen would show the quality level in a coloured box next to the description for easy comprehension.

The full eight-cell allocation design is shown in Table 7. Queue = 0 denotes a wider age gap between daughters. Queue = 1 denotes a tight age gap, where the younger daughter's marriage is approaching soon after the elder's. Spill = 1 represents the presence of a spillover between match outcomes, and Brother = 1 provides additional liquidity that raises both daughters' outcome.

Table 7: Eight-cell design of the allocation task

Cell	Spill	Queue	Brother	Tilted option	Equal option	$W(T)$	$W(E)$	$T-E$
1	0	0	0	(G, P)	(M, M)	0	0	0
2	1	0	0	(G, M)	(M, M)	+1	0	+1
3	0	1	0	(M, P)	(P, M)	−1	−1	0
4	1	1	0	(M, P)	(P, P)	−1	−2	+1
5	0	0	1	(VG, M)	(G, G)	+2	+2	0
6	1	0	1	(VG, G)	(G, VG)	+3	+3	0
7	0	1	1	(G, M)	(M, G)	+1	+1	0
8	1	1	1	(G, G)	(M, G)	+2	+1	+1

Notes: Each ordered pair reports the elder daughter's match quality first and the younger daughter's match quality second. VG, G, M, and P denote Very Good, Good, Medium, and Poor match quality, respectively, using the scale in Appendix Table 6. "Spill" indicates that the elder daughter's marriage generates a cross-sister return. $W(T)$ and $W(E)$ denote the combined match-quality scores under the elder-tilted and equal allocations, respectively; $T-E$ is their difference.

Survey wording. Before the allocation scenarios, respondents are told:

> *"You will now advising families in your area who have two daughters. Over their lifetime, the family has saved **Rs. [W]** for their daughters' marriages. The parents must decide how to split this budget between the two daughters. We will present 4 scenarios, and for each scenario, I will describe the family's situation and show you two ways they could split the budget. We will tell you exactly who each daughter will marry depending on how parent's split allocate the budget. You will tell us which option do you think the parents should choose."*

> *"In this family, the daughters are **[2 / 5] years apart**. There are **[two daughters and no son / two daughters and one son]**. The family has **Rs. [W]** available for the daughters' marriages."*

Figure 12: Allocation choice screen

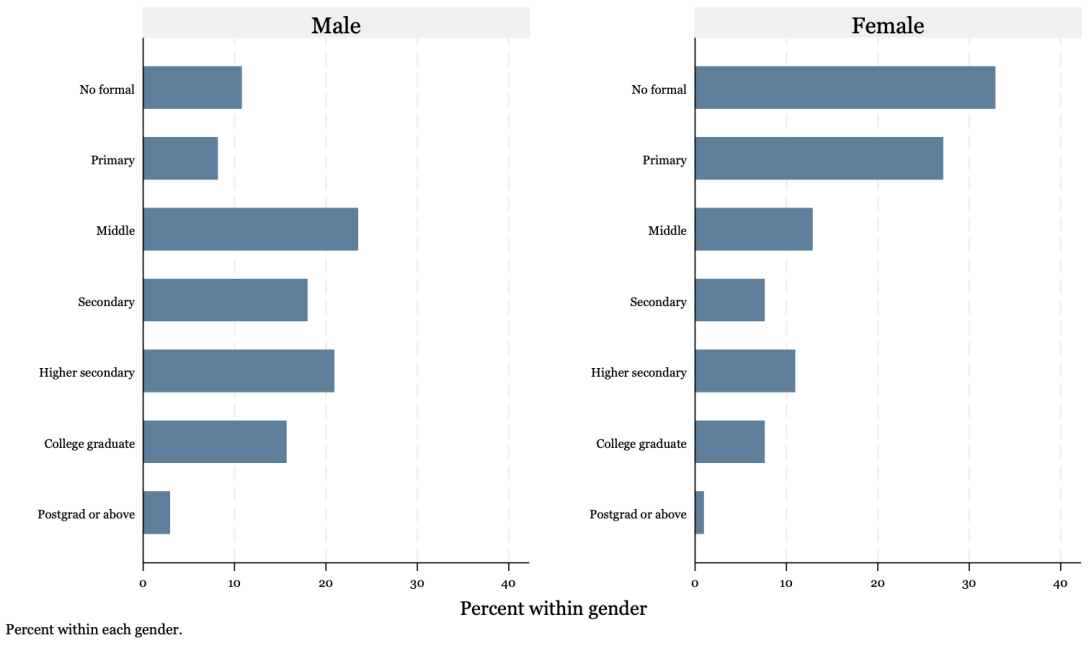
Tilted allocation (60%–40%)	**Equal allocation (50%–50%)**
Rs. 3,00,000 to elder daughter **Rs. 2,00,000 to younger daughter**	**Rs. 2,50,000 to elder daughter** **Rs. 2,50,000 to younger daughter**
Daughters' resulting matches:	
The elder daughter gets a **GOOD (4)** match: Diploma; government clerk.	**The elder daughter** gets a **MEDIUM (3)** match: Class 12 pass; private-sector worker.
The younger daughter gets a **GOOD (4)** match: Diploma; government clerk.	**The younger daughter** gets a **GOOD (4)** match: Diploma; government clerk.

Notes:This changes depending on the scenario, here we are showing Cell 8 in Table 7

"Now that you know what happens under each allocation, which allocation do you think is better for the parents?"

C Summary Statistics

Figure 13: Respondent education, by gender



Notes: Bars report the share of respondents within each gender in each education category.

Table 8: Respondent characteristics

	Mean	SD	N
Respondent and spouse			
Respondent age	49.05	5.91	516
Respondent female	0.41	0.49	516
Spouse age	48.61	7.03	507
Age at marriage	20.12	2.86	515
Location and residence			
Urban PSU	0.22	0.42	516
Resident in this area since birth	0.56	0.50	516
Resident in this area since marriage	0.34	0.47	516
Assets and housing			
Owns land	0.66	0.47	516
Land owned (bigha/acre)	2.97	5.20	331
Owns home	0.97	0.17	516
Owns a television	0.82	0.39	515
Owns a cooler or AC	0.15	0.36	515
Owns a two-wheeler	0.79	0.41	515
Owns a car	0.13	0.34	515
Social group			
Social category: SC/ST/OBC	0.69	0.46	516
Religion: Hindu	0.92	0.27	516
Family and sibship			
Living children	3.13	1.33	516
Daughters	1.69	1.19	516
Sons	1.44	0.97	516
Older brothers	1.20	1.08	510
Younger brothers	1.00	0.98	510
Older sisters	0.80	0.87	510
Younger sisters	0.99	0.99	510

Notes: Binary rows report shares. Land area is reported only for landowning households. Sibling counts describe the respondent's own sibship.

Figure 14: Marriage transfers in respondents' own marriages



Notes: Monetary values are expressed in 2010 rupees.

D Further Results

Table 9: Bride-side responses: full factorial coefficients

	Accept elder	Younger sister's expected match			Family	Required gift
	proposal	Low	Medium	High	standing	high groom
High elder match	0.829***	-0.385***	-0.093***	0.478***	0.851***	-198.6***
	(0.027)	(0.017)	(0.011)	(0.017)	(0.026)	(31.8)
Brother	0.008	-0.006	0.002	0.005	-0.004	27.8
	(0.017)	(0.010)	(0.008)	(0.011)	(0.019)	(19.1)
Tight gap	-0.016	0.023**	-0.016**	-0.006	0.001	-38.5*
	(0.015)	(0.009)	(0.008)	(0.009)	(0.019)	(20.5)
High × Brother	-0.008	0.009	0.015	-0.024	-0.031	54.0**
	(0.022)	(0.013)	(0.013)	(0.017)	(0.026)	(24.0)
High × Tight gap	0.024	-0.017	0.030***	-0.012	-0.019	67.2**
	(0.020)	(0.012)	(0.011)	(0.014)	(0.026)	(27.1)
Brother × Tight gap	0.020	0.003	0.008	-0.011	-0.004	32.4
	(0.025)	(0.013)	(0.011)	(0.013)	(0.027)	(25.0)
High × Brother × Tight gap	-0.027	-0.007	-0.028*	0.035*	0.031	-54.4
	(0.030)	(0.017)	(0.015)	(0.021)	(0.037)	(35.9)
Reference Group Mean	0.132	0.439	0.424	0.137	0.081	887.9
Observations	2,064	2,064	2,064	2,064	2,064	2,001
Respondents	258	258	258	258	258	258

Notes: Full coefficient vector from the bride-side factorial specification. The omitted scenario is a Low elder match, no brother, and a wide age gap. The acceptance outcome appears only in this appendix table. Reference Group Mean reports the mean of the dependent variable for the omitted scenario: Low elder match, no brother, and a wide age gap, within the estimation sample. Respondent and scenario-position fixed effects are included; standard errors are clustered by respondent. The required-gift outcome is in Rs '000 (2010 prices).

Table 10: Groom-family responses: full coefficient vector

	Acceptance update	Family standing	Reservation gift
Medium elder match	0.017	0.256***	-77.6***
	(0.015)	(0.031)	(22.8)
High elder match	-0.017	0.725***	-82.0***
	(0.019)	(0.030)	(25.0)
Brother	0.022	-0.004	-9.8
	(0.016)	(0.016)	(15.5)
Medium × Brother	0.002	-0.004	-5.5
	(0.014)	(0.030)	(15.0)
High × Brother	0.012	-0.015	11.1
	(0.017)	(0.030)	(17.6)
Reference Group Mean	-0.052	0.066	83.0
Observations	1,548	1,548	1,488
Respondents	258	258	255

Notes: Full coefficient vector from the groom-side specification. Low elder match and no brother are the omitted categories. Reference Group Mean reports the mean of the dependent variable for that omitted scenario within the estimation sample. Respondent and scenario-position fixed effects are included; standard errors are clustered by respondent. Reservation gifts are in Rs '000 (2010 prices).

Table 11: Allocation choices: full coefficient vector

	Pr(tilt resources to elder)
Spillover	0.304***
	(0.031)
Brother	-0.036*
	(0.020)
Tight gap	0.154***
	(0.020)
Spillover × Brother	-0.012
	(0.033)
Spillover × Tight gap	0.242***
	(0.033)
Reference Group Mean	0.142
Observations	2,064
Respondents	516

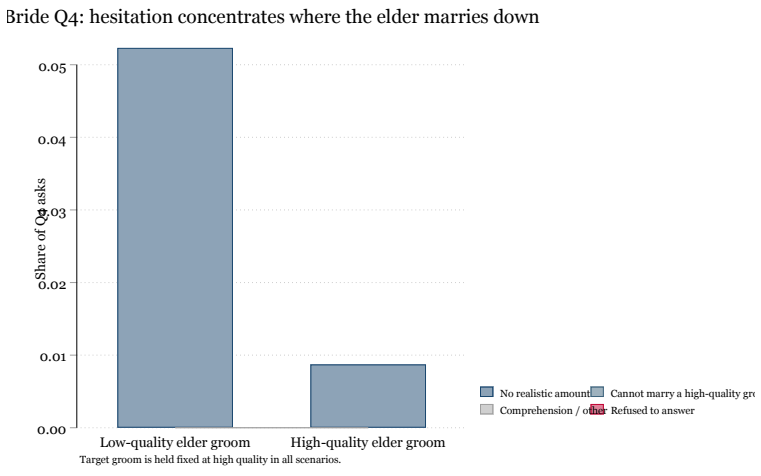
Notes: Full coefficient vector from the respondent fixed-effects allocation specification. The omitted scenario has no spillover, no contributing brother, and a wide age gap. Reference Group Mean reports the mean of the dependent variable in that scenario within the estimation sample. The specification includes respondent and scenario-position fixed effects; standard errors are clustered by respondent. The three-way interaction is not included because it is not jointly identified with respondent fixed effects under the complementary half-fraction assignment.

Table 12: Allocation response across expected-concession quintiles

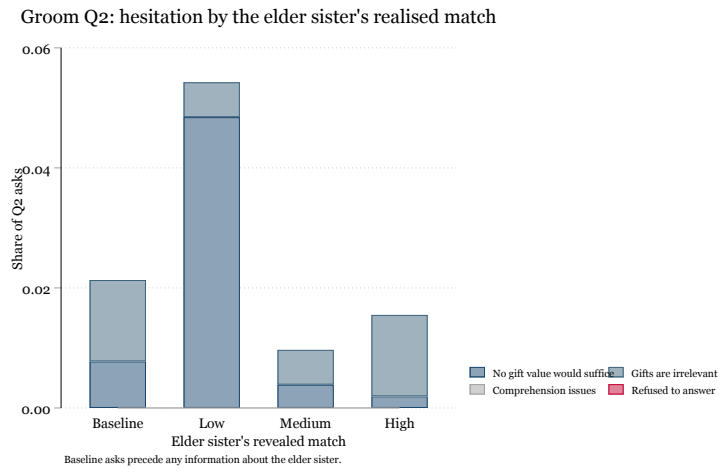
Expected concession	Spillover effect	Difference from Q1
Q1 (lowest positive)	0.477***	–
	(0.084)	
Q2	0.465***	-0.013
	(0.090)	(0.122)
Q3	0.462***	-0.015
	(0.071)	(0.111)
Q4	0.382***	-0.095
	(0.068)	(0.109)
Q5 (highest)	0.437***	-0.040
	(0.075)	(0.113)
Joint p-value: Q2–Q5 differences = 0	0.898	
Reference Group Mean (Q1, no spillover)	0.224	
Observations	576	
Respondents	144	

Notes: Bride-side believers only. Expected concession is the respondent-level difference between the numeric gift required for the younger daughter following a Low rather than High elder match; larger positive values indicate a larger anticipated concession. Quintiles are formed at the respondent level among respondents with a positive measurable concession. The first column reports the implied spillover effect within each quintile. The second reports the difference from Q1 from a common regression interacting Spillover with quintile indicators. The joint p-value tests equality of the spillover response across all five quintiles. The regression controls for brother presence and age gap and includes respondent and scenario-position fixed effects; standard errors are clustered by respondent. Respondents for whom the expected concession cannot be constructed from numeric responses are excluded. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure 15: Refusal to report dowry transfers by elder-daughter's realised matches



(a) Bride Module



(b) Groom Module

E Result heterogeneity by Block

Table 13: Bride-side responses: heterogeneity across randomised blocks

	Accept elder	Younger sister's expected match			Family	Required gift
	proposal	Low	Medium	High	standing	high groom
High × W=10 lakh	0.079	0.058*	-0.039**	-0.019	0.044	-28.7
	(0.048)	(0.031)	(0.016)	(0.029)	(0.044)	(53.2)
High × 70–30 split	-0.027	0.013	0.035**	-0.047*	-0.073*	-30.7
	(0.049)	(0.031)	(0.016)	(0.029)	(0.044)	(55.0)
Reference Group H–L Effect	0.805	-0.427	-0.076	0.503	0.850	-121.8

Notes: Entries are interactions of the High elder-match treatment with the between-respondent marriage-resource block (=1,000,000 rather than 500,000) and expenditure split (70–30 rather than 60–40). The regression also includes the full High × Brother × Tight-gap factorial. Reference Group H–L Effect reports the High–Low effect in the omitted between-respondent block (=500,000 and a 60–40 split), averaged over brother presence and age-gap environments; adding either reported interaction gives the corresponding High–Low effect in that treatment block. The required-gift outcome is in Rs '000 (2010 prices). Respondent and scenario-position fixed effects are included; standard errors are clustered by respondent.

Table 14: Groom-family responses: heterogeneity across randomised blocks

	Acceptance update	Family standing	Reservation gift
Medium × W=10 lakh	-0.020	-0.015	-96.1**
	(0.028)	(0.053)	(45.9)
High × W=10 lakh	-0.004	0.088*	-118.3**
	(0.034)	(0.052)	(50.9)
Medium × High-quality son	0.079***	-0.118**	-43.1
	(0.028)	(0.052)	(44.5)
High × High-quality son	0.141***	0.077	1.5
	(0.035)	(0.052)	(50.6)
Reference Group M–L Effect	-0.011	0.318	-12.6
Reference Group H–L Effect	-0.077	0.637	-19.4

Notes: Entries show whether the Medium–Low or High–Low response changes in the 1,000,000 gift-offer block or the high-quality-son block. The groom-side block scales the displayed gift environment and is not stated as bride-family wealth. The regression also includes the full elder-match × brother specification. Reference Group M–L Effect and Reference Group H–L Effect report the corresponding elder-match effects in the omitted between-respondent block (=500,000 and a medium-quality son), averaged over brother status. Adding a reported interaction to the relevant reference effect gives the corresponding elder-match effect in the 1,000,000 or high-quality-son block. The reservation-gift outcome is in Rs '000 (2010 prices). Respondent and scenario-position fixed effects are included; standard errors are clustered by respondent.

Table 15: Allocation response: heterogeneity across randomised blocks

	Pr(tilt resources to elder)
Spillover × W=10 lakh	0.008
	(0.038)
Spillover × 70–30 split	-0.014
	(0.038)
Spillover × Groom-side respondent	0.054
	(0.038)
Reference Group Spillover Effect	0.395
Observations	2,064
Respondents	516

Notes: Entries report how the spillover response changes in the =1,000,000 block, the 70–30 expenditure-split block, and among respondents previously assigned to the groom-side module. The regression includes Spillover, Brother, Tight gap, Spillover × Brother, and Spillover × Tight gap, together with respondent and scenario-position fixed effects. Reference Group Spillover Effect reports the equal-weighted spillover effect across brother and age-gap environments in the omitted between-respondent block (=500,000, 60–40 split, bride-side respondent). Adding a reported block interaction to this reference effect gives the corresponding average spillover effect in that block. Standard errors are clustered by respondent.

Table 16: Belief heterogeneity across randomised allocation blocks

	Believer differential in spillover response
Panel A: Resource block	
5-lakh block	0.052
	(0.078)
10-lakh block	0.184**
	(0.072)
Difference: 10 lakh − 5 lakh	0.132
	(0.106)
Reference Group Mean	0.264
Panel B: Elder-resource split	
60–40 split	0.025
	(0.073)
70–30 split	0.204***
	(0.075)
Difference: 70–30 − 60–40	0.178*
	(0.105)
Reference Group Mean	0.274
Observations	1,028
Respondents	257

Notes: Bride-side respondents only. Each panel reports the believer-versus-non-believer difference in the spillover treatment effect within the indicated randomized between-respondent block. The third row of each panel is the change in that differential across blocks. Reference Group Mean is the mean probability of tilting among non-believers with no supplied spillover in the lower block of the corresponding panel, averaged over the other randomized block and the brother and age-gap environments. Each regression controls for brother presence and age gap and includes respondent and scenario-position fixed effects; standard errors are clustered by respondent. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

F Field Implementation, Sample Construction, and Randomisation

This appendix documents the field procedures underlying the analysis sample. It separates three parts of implementation that are compressed in the main text: how respondents entered the study, how the instrument was refined before launch, and how experimental configurations were assigned and certified.

F.1 PSU selection and respondent recruitment

Fieldwork was conducted in Lucknow and Prayagraj districts between 19 July and 3 August 2026 with Shailam Longitudinal Research Pvt. Ltd., a Lucknow-based survey firm. The sampling design began with 26 PSUs, 13 in each district, selected with probability proportional to population size from the 2011 Census. Two initially selected Lucknow PSUs were replaced before interviewing began. In one, local community leaders did not permit fieldwork. The second had become urban since the 2011 Census and was replaced to preserve the intended sampling stratum (rural). Both were replaced with an alternative using the same PPS method.

Recruitment within selected PSUs was organised around the family-composition requirements of the two survey arms. Bride-side respondents were required to have at least two daughters; groom-side respondents were required to have at least one son. In rural PSUs, local *anganwadi* workers helped the field team canvass the area and identify households likely to satisfy these criteria. Enumerators then verified eligibility and approached households from these pools until the pre-assigned interview slots for the PSU were filled. In urban PSUs, where comparable *anganwadi* canvassing was unavailable, enumerators screened households door to door within the selected PSU.

This procedure gives probability-based geographic selection at the PSU stage but not probability sampling of households within PSUs. The field forms record 1,088 consent attempts or consent-form submissions, suggesting a refusal rate of $\sim 48.9\%$.

F.2 Piloting and instrument refinement

The final questionnaire was preceded by piloting and short open-ended conversations during field supervision. These exercises were used to refine wording, task structure, and respondent burden; they are not analysed as study outcomes. The pilot led to four changes that materially shaped the final instrument.

First, interviews became difficult to sustain once repeated scenario tasks pushed duration much beyond roughly 30–40 minutes, so we reduced the number of scenarios and simplified repeated instructions. Second, respondents found it easier to evaluate several proposals concerning the same hypothetical family-type before moving to another family. The final bride- and groom-side modules therefore group scenarios by hypothetical family and counterbalance the order in which family blocks appear. Third, pilot respondents consistently described brothers as contributing to sisters' marriage expenses. We therefore dropped a planned distinction between contributing and non-contributing brothers and use brother presence as the observable liquidity shifter in the final design. Fourth, respondents found two consecutive probability-elicitation tasks in the bride module difficult to distinguish. We replaced the first with a binary accept–reject decision and retained the probability elicitation for the younger daughter's expected match distribution.

The qualitative conversations also informed the language used for marriage transfers. Because dowry is illegal and can be socially sensitive, the instrument refers to these transfers as "gifts." The examples and wording were chosen to be familiar to respondents without changing the economic object elicited by the questions.

F.3 Pre-assigned treatment configurations and scenario order

Experimental assignment was fixed at the survey-slot level before an enumerator contacted a respondent. Each slot was associated with a unique identifier containing district, PSU, and respondent serial, for example `LK0-12-01`. For each selected PSU, the field team received a pre-generated roster of core interview IDs for each arm together with buffer IDs that could replace unusable submissions. Every ID was linked ex ante to the relevant between-respondent treatment block, scenario ordering, and allocation configuration.

The SurveyCTO form read these assignments from a lookup table after the enumerator entered the respondent ID. Enumerators therefore did not choose treatment values, scenario order,

or allocation configuration. Once an eligible household consented, the enumerator used one of the remaining IDs for the appropriate survey arm, and the form loaded that ID's pre-assigned configuration automatically. Buffer IDs were generated and assigned in the same way as core IDs, so replacing an unusable interview did not require the field team to construct a new treatment configuration.

Scenario order was counterbalanced using Latin-square designs. In the bride module, the order of hypothetical-family blocks and the order of elder-match scenarios within the relevant blocks were pre-assigned. In the groom module, family-type order and elder-match order were pre-assigned. In the allocation module, respondents were assigned one of the complementary four-scenario sets and the order in which those scenarios appeared; the position of the equal and elder-tilted allocations as Option A or Option B was also counterbalanced. These procedures distribute treatment conditions across presentation positions and limit systematic confounding by fatigue, anchoring, or carry-over. The substantive treatment dimensions are described in Section 4; the point here is that their delivery and order were determined before respondent contact rather than by enumerator discretion.

F.4 Consent, certification, and field quality

All respondents provided informed consent before the interview and were told that responses would be anonymous and used for research. No survey responses were retained for households that declined consent. Of the 1,088 recorded consent attempts, 555 produced consented interviews.

The analysis sample was defined using a common certification procedure. A submission had to satisfy the arm-specific family-composition screen, complete both the belief and allocation modules, and, where active interview timing was available, contain at least 17 active minutes. Two submissions identified contemporaneously during field supervision as unusable were also excluded. Applying these rules to the consented interviews yields 516 certified respondents, 258 in each survey arm. Lucknow contributes 263 respondents and Prayagraj 253. Active timing was introduced after fieldwork began and is available for 284 certified interviews, so the duration criterion is applied only when that measure exists. No additional outlier-based or discretionary cleaning rule determines inclusion in the headline analyses.

Table 17: Fieldwork and analysis-sample summary

	Count
Recorded consent attempts / form submissions	1,088
Consented interviews	555
Certified analysis sample	516
Bride-side respondents	258
Groom-side respondents	258
Lucknow district	263
Prayagraj district	253
Certified interviews with active timing	284

Notes: Recorded consent attempts need not correspond one-for-one to distinct approached households. Active timing was introduced during fieldwork and is therefore unavailable for part of the sample.

Eight enumerators conducted the interviews under daily field supervision. Incoming submissions were reviewed during data collection for eligibility, completion, timing where available, and consistency with the assigned survey configuration. Enumerators received feedback and retraining when concerns arose. Two enumerators were ultimately stood down after repeated protocol concerns. Interviews are not excluded solely on the basis of enumerator identity: their certified interviews remain in the analysis sample if they satisfy the same rules applied to all submissions.

The combination of pre-assigned survey configurations and common certification rules serves two distinct purposes. Pre-assignment protects the experimental comparisons from enumerator choice over treatment or order; certification defines the analysis sample using survey-completion and quality criteria that do not depend on respondents' experimental answers. Neither procedure changes the non-probability nature of household recruitment within PSUs.