Marriage Markets and the Rise of Dowry in India*

Gaurav Chiplunkar
Yale University†

Jeffrey Weaver
UC San Diego ‡

October 2017

Please find latest draft of the paper here

Abstract

Dowry payments are an important part of household finances in India, typically exceeding one to two years of household earnings. Yet there is little empirical evidence on determinants of dowry payments, with existing work relying on small and non-representative samples. In the first part of the paper, we leverage data on over 76,000 marriages to document stylised facts about changes in Indian marriage markets between 1930-2000. We show that although many marriage practices remain static over this time period, there were large changes in dowry payment. Between 1930-1975, the proportion of marriages with any dowry paid increased from 35-40% to nearly 90%. Over the same period, median real dowry more than doubled, but decreased after 1975 in real terms as well as a fraction of household income. In the second half of the paper, we use this data to test major theories of dowry: (i) whether dowry serves as a bequest to female children or is a groom price; (ii) if the increase in dowry prevalence resulted from lower castes adopting high caste practices (Sanskritization); (iii) how changes in sex ratios on the marriage market affect dowry (Marriage squeeze hypothesis); and (iv) if changes in dowry can be explained by hypergamy and cross-caste competition for grooms. We find that the patterns in the data do not support these theories, but instead that the changes are explained by shifts in the quality (earnings/education) distribution of brides and grooms. This has important implications for designing anti-dowry policies.

*We thank seminar participants at Yale's Labor/Public Economics Prospectus Workshop for helpful comments. We are particularly grateful to Dan Keniston for his advice on the project and sharing his encyclopedic knowledge of Indian historical data. The research was conducted with support from a NSF Graduate Research Fellowship. All views expressed are those of the authors and do not necessarily reflect the opinions of any funding organizations.

†gaurav.chiplunkar@yale.edu

‡jbweaver@ucsd.edu

1 Introduction

One of the most significant financial transactions in the life of an Indian household occurs at the time of marriage. Dowry, transfers from the household of a bride to that of her groom, are nearly universal in contemporary India and typically exceed a year of household earnings. Due to the size of these transfers, there is much popular concern about the consequences of the institution, where households may become impoverished as a result of paying dowry or engage in sex-selective abortion to avoid payment of dowry for female children. In this paper, we leverage data on over 76,000 marriages between 1900 and 1999 to document new facts about nature and evolution of dowry and marriage practices over the past century. We then use the data to test prominent theories of dowry, finding that most do not match the patterns in the data.

The scope of payments at the time of a wedding in India is astonishing, especially given the poverty of many households and relative weakness of formal financial markets. In addition to the high cost of a wedding, dowry payments claim a significant fraction of households' income and savings. The Indian government considers the payment of dowry to be a major social ill and has passed multiple pieces of legislation to combat it,[1] but none have proven effective. Recent papers have found that dowry encourages sex-selective abortion (Bhalotra et al., 2016; Alfano, 2017), alters investment behavior of households (Anukriti et al., 2017), and can encourage violence against wives in hopes of extracting further payment (Bloch and Rao, 2002). Each of these has a variety of consequences for general welfare and development.

It is also critical to understand the underlying factors that drive the payment and size of dowry, since different explanations would merit different policy responses. One theory suggests imbalances in the number of men and women on the marriage market due to population growth is the cause of high dowries (e.g. Rao (1993b)), in which case population control policies will prove effective.[2] Another suggests that competition over grooms drives up dowry, and that dowry emerges due to the discrepancy in economic value of men and women. If correct, this suggests that improving female labor force participation and economic opportunities for women will prove more effective at

[1] This began with the Dowry Prohibition Act of 1961, which was followed by additional tightening, such as in the Dowry Prohibition Rules of 1985.

[2] If this theory were true, the government need not do anything since sex-selective abortion is likely to eliminate dowry in many parts of India in the near future. It also suggests that if regulation of sex selective abortion is successful and the population of India continues to grow, dowry may increase.

eliminating dowry. Understanding what drives dowry is the first step in constructing meaningful policy responses.

In this paper, we first examine how marriage markets have changed in India over the past century. As was the case 70 years ago, marriages are almost entirely arranged by the parents of the couple, and brides move to the household of their husband's family, typically in another village. Despite the rise in communications technology, marriage markets are almost entirely within a small geographical area (district), and caste remains the most important determinant of who one marries. 94% of marriages occur between individuals from the same *jati*, a specific identity group based on the ancestral occupation within a village economy. Strikingly, this rate does not decline at all between 1930 and 2012 in rural areas, with a decline of only 2 percentage points in urban areas.

While those marriage customs have not changed, there have been large shifts in dowry practices. Although only 35-40% of brides paid dowry between 1900 and 1940, dowry is rapidly adopted across all of India in the 1940-1975 period. After 1975, dowries are almost universal, with no sign of decline. Others have observed that dowry may be more prevalent than in previous eras Caldwell et al. (1983); Srinivas (1984), but to the best of our knowledge, this is the first quantitative evidence on the scale and geography of adoption. In 1945, the value of dowry began to rise, as has been noted in other papers (e.g. Rao (1993b)). However, the changes within the distribution of payments follow a suggestive pattern that has not been documented. In the pre-1975 period, the increase in dowry is initially driven by increases in the top half of the distribution of dowry payments, followed by increases throughout the distribution. There is then a substantial decline in top percentiles of dowry payments in the post-1975 period with almost no change in other parts of the distribution. This motivates the main question of the paper: why do we observe such huge changes in dowry amount, and what do these changes tell us about why dowry occurs?

In the second part of the paper, we provide rigorous tests of numerous prominent theories of dowry, many of which have never been empirically tested. One of the theoretical debates around dowry is whether it stems from a parental desire to bequeath resources to their daughters or as a market clearing payment for grooms in a competitive market. We show that the latter model more closely fits the data since theories of bequest have difficulty rationalizing the massive increase in dowry. Such an increase would either come from increases in family wealth, which are small over this time period, or the desire to provide daughters with a greater share of the inheritance, which

is inconsistent with other investment decisions of the family.

Another well-known explanation for changes in dowry prevalence is the Sanskritization hypothesis by Srinivas (1984), which proposes that dowry was always practiced among the upper castes and spread as lower castes emulated them. This emulation is attributed to low caste groups attempting to "*Sanskritize*", or increase their caste status by emulating practices of high caste groups (e.g. vegetarianism, dowry). We show that this theory cannot explain the rise of dowry. Both low and high caste groups began wide-scale adoption of dowry at around the same time and at similar rates.

A number of papers rationalise changes in dowry as the result of population growth (e.g. Rao 1993b; Billig 1991, 1992; Dalmia and Lawrence 2005; Sautmann 2011). Since men marry at older ages than women, there is a surplus of women on the marriage market in a growing population. This creates a "marriage squeeze", where competition over grooms leads to an increase in dowry. Previous empirical papers have documented associations between surpluses of women on the marriage market and increases in dowry using small data sets of fewer than 200 individuals (Rao, 1993b; Edlund, 2000), but there is significant debate over whether such associations exist. A related theory proposed by Anderson (2007b) argues that a surplus of women in the marriage market should instead result in changes in the age of marriage that relieve the population pressure. We find that sex ratio in the marriage market is not related to increases in the prevalence or size of dowry, but is strongly correlated with reducing the age gap.

Anderson (2003) proposes a theory linking dowry increases to increases in wealth dispersion within caste groups. Intuitively, the paper argues that low caste families prefer to marry their daughters to high caste men and as the number of wealthy low caste households increases, they attempt to attract high caste grooms by offering larger dowries. These high caste grooms can then use the threat of out-marriage to increase the dowry payments they receive from in-group women, which pushes up the size of dowry. We devise a number of empirically testable predictions of the theory and find that they are violated. For example, increases in dowry size are not found among the types of high caste men that the theory would predict. This can be explained by a violation of one of the assumptions on individuals' preferences over potential partners from other castes.

We provide evidence that increased differentiation in groom (and bride) quality can explain the rise of dowry during the process of modernization, which is consistent with the theory discussed in

Anderson and Bidner (2015). Starting in the 1930s and 1940s, a class of higher earning and salaried positions for Indian men began to emerge, and the education level of men began to rise. Those grooms command higher dowries since those characteristics are valued on the marriage market, and as their numbers increase, this leads to a rise in dowry. We first demonstrate that dowry increases are concentrated in areas and caste groups whose educational attainment is increasing. We then show that dowry increases are attributable to groom quality, using variation in dowry payments within families and caste.

Our paper makes a number of contributions to the literature. First and most important, we contribute an empirical element to an almost entirely theoretical literature on dowry (Anderson, 2003, 2007b; Anderson and Bidner, 2015; Becker, 1973; Bhaskar, 2016; Botticini and Siow, 2003; Choo and Siow, 2006). Due to a lack of comprehensive large scale data on dowry payments and marriage patterns, there is a wide disconnect between the theory and its empirical validations. Providing tests (and finding little empirical support for many theories) enables us to describe how marriage markets in India actually function, which is relevant for future theoretical frameworks as well as for designing anti-dowry policies.

Second, we provide a thorough characterization of Indian marriage markets in the post-1930 period and document novel stylised facts on the evolution of these institutions and practices. We show that even as India has economically modernized, many features of marriage markets have not changed. Our findings contrast with much of the existing literature, which only had data on a small samples of marriages or has not considered national trends (Rao, 1993b; Edlund, 2006; Arunachalam and Naidu, 2008; Sautmann, 2011). Anukriti et al. (2017) carry out a similar exercise, but focus on a later time period and different aspects of dowry. Our paper's results are useful for understanding overall patterns in marriage markets, which are relevant for a broader understanding of household finances in India.

The remainder of the paper is organized as follows. Section 2 discusses the data used in the paper, while section 3 uses that data to document stylised facts on the nature and evolution of marriage markets in India since 1930. Section 4 provides quantitative tests for the most prominent theories of dowry. Section 5 concludes.

2 Data

The main analysis on dowry is based on the 1999 round of the Additional Rural Incomes Survey/Rural Economic and Demographic Survey (REDS), a detailed survey of rural households conducted by the National Council of Applied Economic Research (NCAER) across 17 major states in India. The REDS data set, collected in four rounds between 1971-2008 collects detailed information on social, demographic and economic characteristics of households. Most importantly, the REDS collects detailed information on the marriages of the household head, their parents, all of their brothers and sisters, and all of their sons and daughters.[3] Critically, this includes a monetary valuation of transfers made at the time of each marriage from the household of the bride to that of the groom, and the household of the groom to that of the bride. We use this to construct a retrospective history of marriages across decades comprising over 76,000 observations. Our main analysis is based on the 1999 round, but we use the 2008 round in some sections as a check of data quality.[4]

Decade	Pre-1930s	1930s	1940s	1950s	1960s	1970s	1980s	1990s	Total
Marriages	2,253	3,198	5,586	8,353	11,683	15,141	16,559	11,573	74,346

Table 1: Number of Marriages in the 1999 ARIS-REDS data by decade of marriage

The REDS data is a substantial improvement over other sources. The most widely cited empirical articles about marriage in India have been based on a retrospective survey on marriages collected in 1983 by the International Crops Research Institute for Semi-arid Tropics (e.g. Behrman et al. 1999; Deolalikar and Rao 1995; Edlund 2000, 2006; Rao 1993b). These data only contain 127 observations on marriages between 1923 and 1978 from six villages in three districts of South India (see Edlund (2006) for descriptive statistics) and may not be representative of larger trends across the country, particularly given India's cultural and regional heterogeneity. Due to the small sample size, we do not use the ICRISAT data in this paper. Another source of data was collected by the National Council of Applied Economics Research in 1995 from two states in India (Anderson, 2007a; Dalmia, 2004; Dalmia and Lawrence, 2005; Sautmann, 2011). We also elect not to use this data since it is still relatively small in size, has poor coverage of the period prior to 1970 (when

[3] A particularly nice feature of the data is that it contains information even for family members who have died, so there is no data loss due to mortality.

[4] Note that this wave is sometimes referred to as the 2006 round of the ARIS/REDS survey. Although the first survey round was conducted in 2006, 84% of responses are from 2008. Thus we refer to it as the 2008 round in this paper.

most of the changes in dowry practices are occurring) and lacks all-India coverage.

We also utilize two additional sources of data for information on marriages: (i) the Indian Human Development Survey (IHDS)[5] and (ii) the National Family Health Surveys (NFHS)[6]. Both are large, multi-round, representative surveys conducted across India. The IHDS was coordinated by researchers at the University of Maryland, and has two waves of data on approximately 41,000 households (2005 and 2011). This includes detailed information on marriage practices over time that we will use for descriptions of changes in Indian matrimonial markets. However, it only collects perceptions of average dowry payments at the time of the survey, rather than actual payments made at the time of marriage. India's Ministry of Health and Family Welfare coordinates the NFHS, and has completed three waves of data collection: 1992-93, 1998-99, and 2005-06. In each round of the survey, around 90,000 married women were surveyed about their marriage and family. This is an excellent source of information on age of marriage and fertility, but is limited in scope since it contains no information on dowry payments. These data play a secondary role in this paper and are used only to show demographic patterns at particular points in time.

Lastly, for some of the analysis, we use the Survey of Women and Fertility (SWAF, Smith et al. (2000))[7]. Between 1993 and 1994, researchers at the University of Pennsylvania conducted the survey in 1,551 households in Tamil Nadu and 895 households in Uttar Pradesh, spread across four districts. It addresses similar topics to the NFHS, but contains additional retrospective information on bride/groom selection, marriage practices as well as dowry and gifts exchanged at the time of marriage. It does not contain information on the total value of dowry payments made. Nearly all of the SWAF data is from marriages occurring after 1975, so it is not useful for most historical patterns. We use it to confirm patterns observed in more recent times.

Given the lack of extant historical data on dowry, it is necessary to use retrospective data. One possible issue is that respondents may be unable to properly recall dowry transactions, particularly given that many occurred well before the time of the survey. However, there are a numerous reasons to think that recall will not be a problem– first, marriage is one of the main focuses of life in rural India and events around the time of marriage will be particularly salient. Second, the magnitude

[5]See `http://ihds.info/`

[6]See `http://rchiips.org/nfhs/`

[7]See `http://swaf.pop.upenn.edu/` for data and details. This data was also used by Logan and Arunachalam (2014) and Alfano (2017).

of dowry payments makes it likely that respondents will accurately recall them (it would be like asking a home owner in the developed world what they had paid for their house).

Appendix B.1 provides a detailed discussion of possible biases and two direct tests of the validity of retrospective data on dowry in the REDS. The first test takes advantage of the panel nature of the REDS survey. Since respondents were interviewed in 1999 and 2008 and asked similar questions about dowry payments, a systematic recall bias (or any temporal bias in general) would lead to differences in responses to the 1999 and 2008 waves of the survey. However, responses are very similar on average, particularly when asking about marriages in the post-1945 period, the main focus of our analysis. As a second check, we use data from the SWAF. It conducts separate interviews with husbands and wives, but asks a series of identical questions about whether particular items were given as part of the dowry in their marriage[8]. If recall is poor, then we would expect that the answers of the two parties would be poorly aligned. Instead, we observe identical responses in 87.8% of cases. It is particularly notable that there is nearly no decline in the match rate between couples whose marriages occurred well before the survey and whose marriages were just a year before the survey (see appendix figure A22). This increases our confidence in the validity of using retrospective dowry about dowry: while recall might be poor for less significant economic transactions, it does not appear to be so for one as large and important as dowry.

We combine two other data sources with the marriage data to construct measures of demographic factors that may be relevant to marriage markets, such as the number of men and women of a particular age group. The first of these is the Census of India, which is conducted every 10 years, and administered by the Government of India. This is a complete survey of all households in India. We use the rounds between 1911-1991, primarily its population counts of men and women at different age groups. The second data set is the National Sample Survey (NSS), also administered by the Government of India. We pool rounds 38 [1983], 43 [1987-88], 50 [1993-4], 55 [1999-2000], and 62 [2005-6] to get individual-level information on the location, broad caste grouping, gender, religion, education and birth year of all individuals in a surveyed household[9].

[8]The items included are land, jewelry, cash, a vehicle such as a car, a TV, furniture, a radio, utensils, a bicycle, livestock and clothing

[9]We re-weight according to the weighting file to produce representative figures of the population.

3 Marriage Markets in India: 1930-present

In this section we provide an in-depth look at the functioning of Indian marriage markets over the past century and document some stylised facts by linking multiple sources of data to examine aggregate trends as well as regional variation. This lays the groundwork for understanding the prevalence and evolution of marriage practices and dowry over the last century. While some of these facts have been discussed in other papers, it remains a useful background for understanding Indian marriage practices.

3.1 Marriage Practices

Marriages in India are almost entirely monogamous with extremely low rates of divorce. Data from the NFHS (2006) shows that fewer than 1% of women are divorced, separated or deserted by their husbands. Parents play an important role in choosing the bride/groom. According to the IHDS 2005 (which only asked the wife), the wife chose her husband in 5-10% of marriages between 1960 and 2005, with only a slight increase over that time period (see figure 1). Women seem to have increasingly had at least some input into the choice of groom: from 25% of marriages in 1960 to around 40% by 2005[10]. This pattern is true for men as well, though traditionally, men hve had more influence over the selection of their spouse (figure A2). Over 60% of couples between 1970 and 2000 met only on the day of the wedding, though this rate has declined rapidly in recent years (appendix figures A3 and A5).[11]

Lastly, over 90% of couples live with their husbands' family after marriage (patrilocal exogamy) and only 10-15% of women marry someone from their own village (IHDS, 2005).[12] [13] While brides move outside of their village, they don't move far. 78.3% of marriages are within same district (REDS, 1999), with an average travel time of 3 hours from the household of the bride to that of the

[10]There is some heterogeneity in levels of input by women by state (appendix figure A1), but aggregate trends are fairly similar across states.

[11]There is substantial regional heterogeneity, where in north India (e.g. Uttar Pradesh, Bihar, Punjab, Haryana) around 90% of couples only meet on the wedding day. In southern India, it is more typical for couples to have met prior to the wedding, but this is often due to marriages between cousins rather than social contact.In Tamil Nadu, where the SWAF data gives a more detailed account, the bride and groom were related in 90% of marriages where partners knew each other prior to the wedding.

[12]There is significant regional heterogeneity, with less than 10% of marriages in north India occurring between partners from the same village, but approximately 25% of marriages in South India occurring within the same village (appendix figure A7).

[13]Over the time period 1970 to 2005, rates of marrying within the same village (appendix figure A6) and co-residing with the husband's family after marriage are unchanged across all regions of India (appendix figure A10).

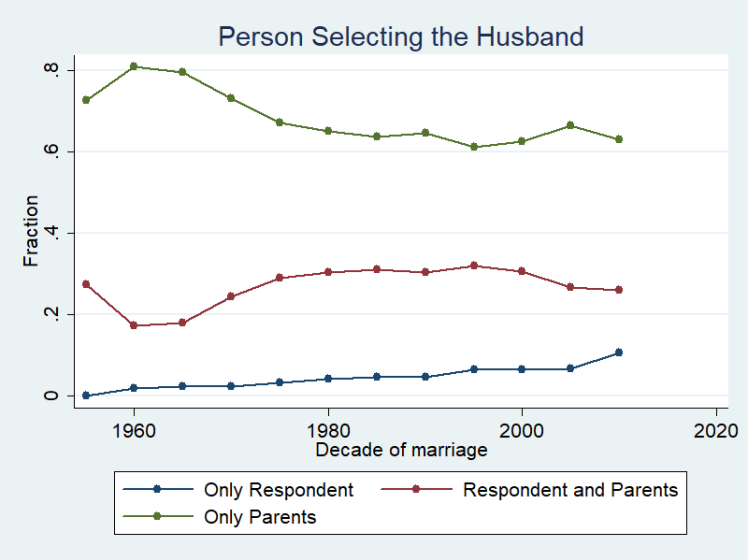

Figure 1: Decisions Over Identity of Husband (IHDS)

groom (IHDS, 2005). One might have suspected that the advent of information communications technology would expand markets and thus the distances over which marriages are made, but we find no changes in the average distance between 1930 and 2010 (see appendix figures A8 and A9). Consistent with earlier literature (Rosenzweig and Stark (1989)), this suggests the role of social connections in screening prospective partners.

In recent years, there has been a great deal of interest in assortative matching in developed country marriage markets, with numerous papers highlighting increased positive assortative matching as a reason for increases in income inequality (e.g. Eika et al. (2014); Greenwood et al. (2014)). In India, female labor force participation is low, so positive assortative matching on education need not have such substantial implications for income inequality as it does in the developed world. However, it is still useful to note whether the presence of transfer payments (dowry) reinforces or reduces the assortativity of matching. Appendix figure A12 uses the REDS data to show the education level of the spouse based on the education of their partner. Clearly, the matches are strongly positive assortative: a man with no education virtually always marries a woman with no education, whereas those with higher educational attainment tend to marry women with similar education. However, it is difficult to tell how assortativity changes over time, since the figure does not provide information on the share of individuals in each education bin by gender. Appendix D.1 follows the methodology in both Eika et al. (2014) as well as Greenwood et al. (2014) and shows

that assortativeness of matching has increased across education groups.

3.2 The Role of Caste in Marriage

One of the most significant features of the Indian marriage market is caste. Indian society has traditionally divided individuals in different sub-castes (*jatis*), based on the traditional occupation of their ancestors within a village economy (e.g. water carriers, leather workers). As has been documented in many studies, individuals have a strong preference for marrying within their own *jati*, or sub-caste group (Dugar et al., 2012; Banerjee et al., 2013). A recent all-India survey by CNN-IBN found that 74% of respondents were opposed to inter-caste marriage. Banerjee et al. (2013) find that the preference is so strong that a woman would be indifferent between a husband from the same *jati* with no education and a husband from a different *jati* with a master's degree.

In both the REDS and IHDS data, the prevalence of marriages across caste boundaries is incredibly low in rural areas, with only 6% of marriages occurring between individuals from different sub-castes.[14] [15] Most importantly, we find no change in inter-caste marriage in rural areas over the period 1920-2011, contrary to many notions of modernization (see figure A11). Even in urban areas, there is a scarcely perceptible decline, and rates of within jati marriage are well above 90%. Thus different sub-castes can be thought of as part of independent marriage markets.

3.3 Dowry Prevalence and Size

Transfers made at the time of marriage typically include both cash payments as well as gifts of physical items such as jewelry or kitchen utensils. The SWAF provides a breakdown of the composition of transfers (see appendix figure A14). Cash transfers (or dowry payments) account for a significant

[14] The concept of *varna*, which divides caste groupings into five levels (Brahmin, Kshatriya, Vaishya, Shudra, Untouchable), is often used as a method of categorization. Instead of *varna*, we define levels on the caste hierarchy along the ARIS-REDS categories of Brahmin, non-Brahmin upper caste, scheduled caste, scheduled tribe, other backwards class, and non-classified members of the major religious groups (Hindu, Muslim, Sikh), and check what percent of marriages involve couples from different groups. Only 2.4% of marriages occur between individuals from different broad categories, with no change over time.

[15] The rural REDS data has slightly lower rates of within caste marriage than the IHDS rural sample. This is likely because the IHDS explicitly asks individuals whether they married within their caste group (yes/no), whereas the REDS asks about the caste of the bride and groom and we analyze if that is the same. If REDS surveyors made a mistake in even 1% of cases when noting the *jati* of individuals, this would fully account for the difference between IHDS and REDS. Respondents may also be more likely to respond to the IHDS question framing in the affirmative if inter-caste marriage is socially undesirable.

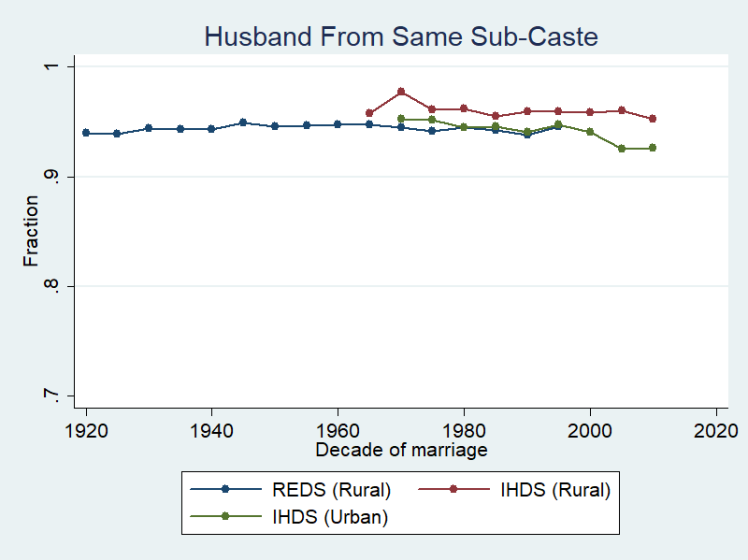

Figure 2: Prevalence of Inter-jati Marriage from 1930-1999 (REDS)

proportion of transfers made during marriage.[16] Since the seminal work of M.N. Srinivas in 1940's Maharashtra, sociologists have stated that dowry has become increasingly prevalent in marriages in India (e.g. Srinivas (1976, 1984)). Other researchers discuss transitions away from payment of bride price to a dowry/groom price system (e.g. Rajaraman 1983; Billig 1992). Srinivas (1984) states that dowry was prevalent among higher caste groups, while low caste groups paid brideprice, a claim that we will return to in section 4. However, these papers are entirely based on qualitative impressions or small studies of particular villages, typically in an ethnographic context. Likely due to lack of data availability, there is no quantitative evidence on the size and timing of shifts to dowry in India.[17] Using data from the 1999 wave of the REDS data, we document three stylised facts with regard to the prevalence and evolution of dowry payments in India over the past century:

Fact #1: On the extensive margin, there was a rapid increase in the prevalence of dowry payment

[16] Gift transfers rarely included transfers of land (1.2% of marriages), and at the time of the SWAF (1993-4), cars or motorcycles (1.2%), TV/VCRs (2.5%), or large kitchen equipment such as a refrigerator (2.1%). On the other hand, nearly all weddings featured gifts of jewelry (91%) , kitchen utensils (94.5%), and clothing (95%). Other items are relatively common, such as furniture (23% in Tamil Nadu, 50.3% in Uttar Pradesh), radios (3% in Tamil Nadu, 34% in Uttar Pradesh), bicycles (0.3% in Tamil Nadu, 37.2% in Uttar Pradesh) and livestock (4% in Tamil Nadu, 22% in Uttar Pradesh). Over the 1965-1993 period, the fraction of marriages in which each of these are given is roughly constant, with decreases in the frequency of livestock transfers and increases in frequency of consumer goods such as bicycles, furniture and radios.

[17] Arunachalam and Logan (2016) address a similar question on whether the motivation for dowry transfers has shifted from bequest to groom price in Bangladesh, and do find some shifts, but look entirely within among households paying dowry.

between 1940-1975. After 1975, dowry payments were nearly universal across all marriages.

Fact #2: On the intensive margin, dowry payments rose across all parts of the distribution between 1945-1975. This is initially lead to increase among the upper tail of the distribution, followed by a shift of the bottom half of the distribution. Post 1975, there was a decline in the average dowry payments, driven by declines in the upper percentiles of the dowry distribution. The lower percentiles and the median were mostly unchanged.

Fact #3: As a fraction of annual household income, median dowry payments were around twice the annual household income in the 1960s. They have been declining over time to around twice the annual household income by 1990.

We now examine the above facts in detail. Figure 3 shows the prevalance of dowry on the extensive margin (Fact #1). Before 1940, only 38% of households engaged in the payment of dowry, but this increased to 88.2% by 1970 and has remained steady since then.[18] [19] [20] The timing of adoption is consistent with existing ethnographic evidence, such as Srinivas (1976) noting the rapid adoption of dowry in Maharashtra in the 1940s. This raises one of the central questions of the paper: what drove the sudden flourishing of dowry in the post-1940 period?

The existing literature on dowry payments in India has focused on the intensive margin of dowry payments, i.e. the size of dowry payments made by those who are paying dowry. It is commonly thought that the size of dowry payments has increased over time (Billig 1991, 1992; Epstein 1973; Rajaraman 1983; Bhaskar 2016; see Anderson (2007a) for a review). Quantitative evidence for this claim comes from Rao (1993b), who finds an increase in dowry size in the ICRISAT data. Claims of dowry inflation have been challenged by Edlund (2006), who redoes the analysis of Rao (1993b)

[18] Appendix figure A13 shows the geographical heterogeneity in adoption of dowry. Some states, such as Himachal Pradesh and Madhya Pradesh had near universal payment of dowry over most of the entire sample period, while others, such as West Bengal had a consistently low rate of dowry payment. Bihar and Rajasthan had higher initial rates of dowry payment and quick adoption, while Uttar Pradesh had low initial rate of dowry payment and slower adoption.

[19] As a further test, we also examine payment of dowry in the SWAF. That data set is not ideally suited to this task, since nearly all of its marriages occur after 1965, at which point dowry adoption is nearly universal. These data also separate payments from the bribe to groom's household into gifts and dowry payments, while the ARIS-REDS data combine the two. Nonetheless, when we compare the ARIS-REDS dowry prevalence data from Uttar Pradesh and Tamil Nadu to the SWAF (see appendix figure A14), the patterns are qualitatively the same.

[20] One concern is that the pattern could be produced by poor recall, where individuals have systematically worse recall in earlier periods and hence state they did not pay dowry. However, the data separates non-payment and non-recall of dowry and the rate of non-recall is quite low and not skewed towards older marriages (see appendix figure A25). As a second test, we compare the 1999 and 2008 rounds of the REDS and the patterns are similar (appendix figure A24).

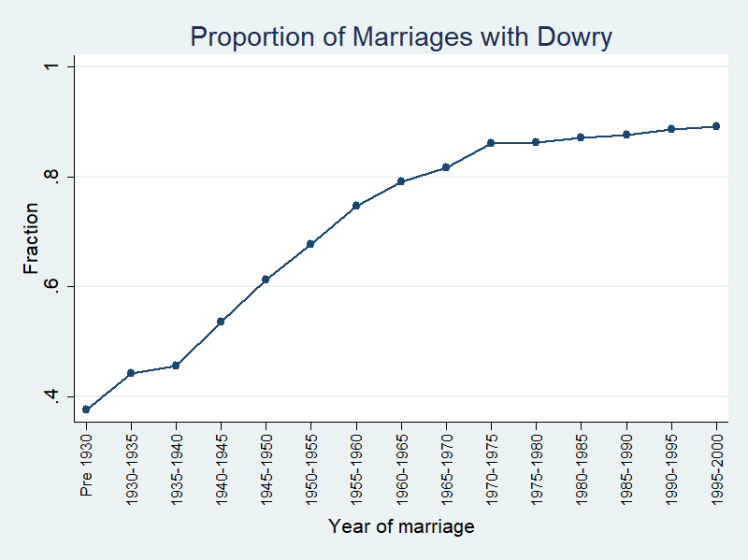

Figure 3: Prevalence of Dowry from 1930-1995

with the ICRISAT data and finds no inflation, as well as others (Logan and Arunachalam, 2014). However, the small sample size of the ICRISAT data makes it poorly equipped to find aggregate trends. Anukriti et al. (2017) use the 2008 round of the REDS data to investigate changes in dowry value after 1960 and similarly find little inflation.

We consider dowry size between 1930 and 1999 in the 1999 wave of the REDS data, which has three key advantages over the existing research. First, its coverage is much broader. Aside from Anukriti et al. (2017), all other quantitative evidence comes from extremely small sample sizes (Edlund, 2000; Rao, 1993b) or only have data from part of the country (Sautmann, 2011; Logan and Arunachalam, 2014). Second, all of the existing literature has considered the size of payments made by those paying dowry, but not the extensive margin of dowry payment. Due to the enormous changes on the extensive margin over time, such an analysis could produce misleading results. For example, if individuals who previously did not pay dowry shift into paying low dowries, this would produce the illusion of a decline in mean dowry payments, when in fact, total dowry payments have increased and always-payers of dowry did not change their dowry payments. Third, we consider how the overall distribution of dowry payments has changed, rather than just the mean, since focusing only on the central moment masks important patterns.

Figure 4 plots the 35^{th}, 50^{th} and 75^{th} percentile of real dowry payments[21] (inclusive of payments

[21]We will define the value of dowry as the net of payments made to the household of the groom at the time of

of zero) as a three year moving average.[22] Dowry size steadily rises between 1930-75, with a deviation from the upwards trend in the 1940s. That deviation is likely attributable to World War 2, during which India experienced high inflation and a famine in 1943. A large pool of marriage-age men also served in the British army during World War II (around 2.5 million), so the sample of those marrying during this time period may be unusual. After 1975, dowry payment is nearly universal, but the upper percentiles of the distribution falls in value while median and lower payments slightly increase. This is a stark contrast to journalistic accounts, which suggest large increases in dowry payments. That may be because individuals experience money illusion, where they think in nominal rather than real terms (Shafir et al., 1997), since there have been large increases in nominal dowry (appendix figure A17).

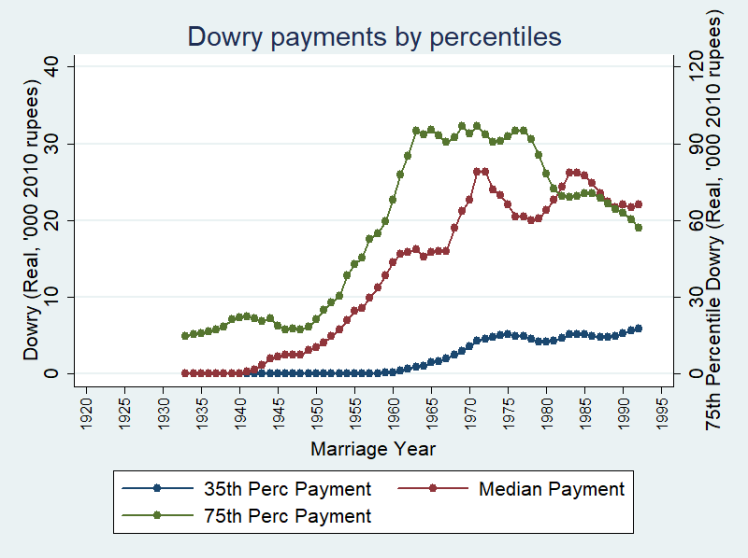

Figure 4: Dowry Payment By Year of Marriage

Figure 5 plots the full distribution of dowry payments by decade between 1940 and 1999. Between 1940 and 1950, the main shift in the distribution is around the median, along with some increase in the upper tail around the 80^{th} percentile. In the following decade, the increase was felt throughout the distribution, with a nearly uniform shift outwards. In the 1970s, the main

marriage minus those made to the household of the bride. In addition to dowry, there often gifts transferred from the family of the bride to that of the groom, and this is the standard definition in the literature (e.g. Rao (1993b); Edlund (2006)). As shown in the appendices, using either gross dowry or net dowry has little effect on our results since the two are highly correlated.

[22] Appendix figure A16 shows the mean of dowry payments over time. The patterns are qualitatively the same, but more overstated, likely due to the presence of outliers in the upper tails of the distribution. The figure also shows the mean without and without incorporation of the extensive margin, and finds that ignoring changes on the extensive margin overstates the rise in dowry payments.

increase occurred around the 60^{th} to 70^{th} percentile of the distribution, with stagnation elsewhere. Following the 1970s, there is a clear and concentrated decline in dowry, driven entirely by the upper half of the distribution. Among the bottom 35% of payments, there is no shift at all, while the upper tail of payments shifts inwards substantially. This pattern will be helpful when evaluating what theory of dowry match these changes.

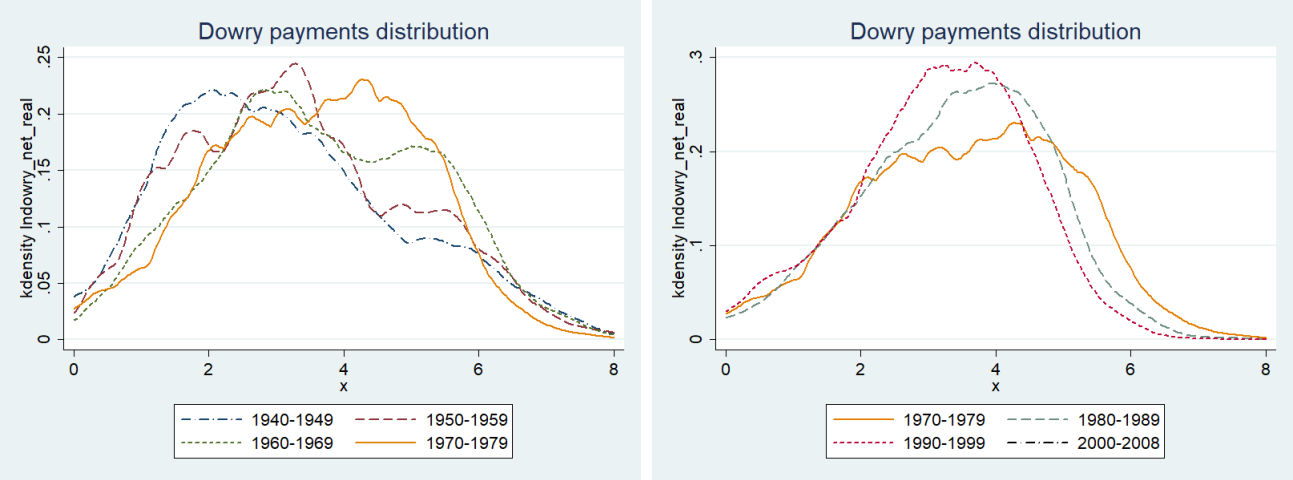

Figure 5: Dowry Payment Distribution

Finally, we document trends in dowry payment as a fraction of household earnings, using wage and expenditure data from the National Sample Surveys between 1960-1995 to compute average annual household earnings for each state in our sample (see appendix section B.4 for details). Using a three-year moving average, figure 4 shows that both mean and median dowry payment decline as a fraction of average annual household earnings. Nonetheless, payments are still substantial, with the median payment at between one to two times annual household income. Unfortunately, the NSS data does not contain percentiles of the income distribution, and we do not observe income at the time of marriage for households in the REDS data, so we cannot capture how this changed by the economic status of households: it may be that poorer households pay a larger or smaller fraction of household income, or experienced changes that differ from the average household[23] .

[23]Our analysis on dowry is similar to and consistent with that of Anukriti et al. (2017), who look at national trends in dowry using the 2008 wave of the REDS data. While Anukriti et al. (2017) focus on geographical, case and religious variation in the 1960-2008 period, we look between 1930-1999 and document changes on the extensive margin as well as distributional changes in the dowry size. See appendix section B.2 for a discussion of the trade-offs of different data sets on dowry, such as the 1999 and 2008 waves of the REDS data.

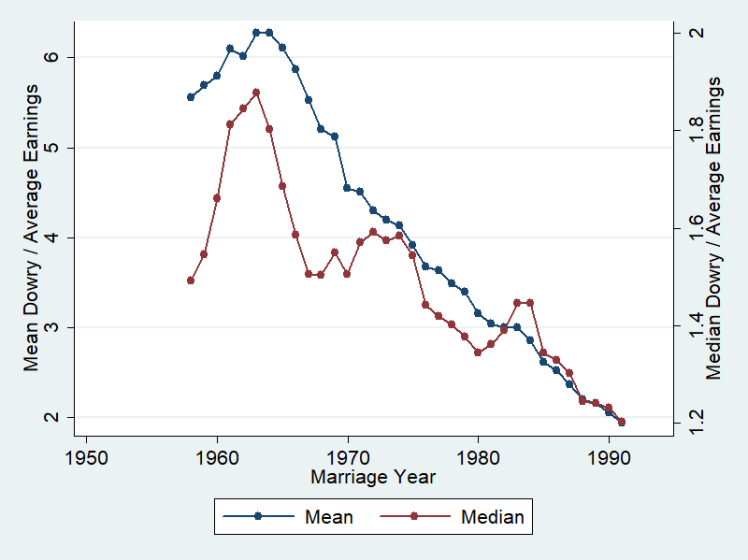

Figure 6: Income and Dowry Payments

4 Testing Theories of Dowry

The previous section documents that most features of Indian marriage markets have remained static over the past century. Despite that, there is a massive increase in prevalence and size of dowry between 1930-75, and a potential increase in the positive assortative matching on education.[24] This section of the paper examines why those changes occurred and what this implies for understanding payment of dowry more broadly. The literature makes many theoretical arguments on the evolution and prevalence of dowry practices in India and yet, there has been virtually no rigorous testing of these theories, primarily due to the lack of comprehensive data.

Within the economics literature, there are two broad explanations for payments at the time of marriage (dowry/brideprice). The first considers dowry to be a pre-mortem bequest, where parents use dowry to give their daughters their inheritance. This idea is extended by Botticini and Siow (2003), who note that in developing countries with patrilocal endogamy, male children remain part of household production, but daughters do not. If parents passed on part of earnings to their daughters after death, this would disincentive effort by male children, and thus it is efficient to provide the inheritance as dowry at the time of separation from the family. Under the second set

[24] In a traditional matching model, the increase in positive assortative matching implies that there has been an increase in the complementarity of male and female education. This may have occurred because children's education became increasingly important, and so families that placed a high value on education became less likely to want to select uneducated wives (who are an input into a child's education).

of explanations, dowry is a price in a matching market so that the markets clear. This could take the form of bride price, where transfers are made from the groom's family to the bride's, or dowry, where brides pay higher prices to match with higher quality grooms (Corno et al., 2017). Some papers have fused these explanations, noting that both motivations may play a role, there may be transitions to and from dowry over time (Caldwell et al., 1983; Kapadia, 1993), or that there may be heterogeneity in dowry motives across different families (Logan and Arunachalam, 2015).[25]

We do not provide direct quantitative tests for bequest theories since it is immediately clear that they cannot explain the sudden and massive increase in dowry prevalence and size documented in the previous section. If dowry is a bequest, then a given daughter's dowry will be equal to the product of a family's expectation of its future earnings and the share of earnings that it wishes to bequeath to that daughter. The rise in dowry after 1945 and 1970 would have to be caused by either an increase in expectation of future earnings or the desire to pass a much greater share of earnings to ones daughters. The former is inconsistent with the limited income growth over this period. The latter explanation would imply a massive increase in parental weight over the utility of their daughters relative to sons, so at reasonable discount rates, there should be huge shifts in the allocation of resources to daughters relative to sons prior to the daughter's marriage. In those outcomes that we can measure, such as education, there is no evidence for such an effect. Data also indicate that dowry transfers are not direct gifts to a daughter: the SWAF data finds that only 9% of wives have majority control over the cash component of dowry, with a split between whether it is controlled by their husbands or their husband's family. It may be that the bride's family is bequesting her with a husband rather than the cash value of dowry transfers, but while bequest motives may play a role in some portion of dowry payments, this role cannot be large enough to explain the major shifts in dowry observed in the data.[26]

[25] Related concepts are used to categorize the components of dowry in India: *stridhana*, gifts from parents to the bribe that remain her property (often jewelry or gold) and pure groom price, cash payments that are the property of the groom and his family. Over time, the latter category is speculated to have become increasingly dominant (Anderson and Bidner, 2015). While *stridhana* payments are clearly bequests, payments that appear as groom price may also be bequests, where the family purchases a higher quality groom as their daughter's inheritance.

[26] Qualitative data from the ICRISAT survey is also suggestive of non-bequest motivations. Household heads were asked about the dowries they had paid at the time of marriage of their daughters. They were then asked the principal reason why the dowries of particular daughters deviated from the average dowry payment. As reported by Behrman et al. (1999), the primary reason was schooling differences across the daughters in 34.1% of cases, and the wealth of the groom as the reason in 32.7%. No other reason is nearly as common, with the third most common reason, the physical characteristics of the daughter, only being cited in 8% of cases. They do not report responses consistent with a bequest explanation for differences in dowry payments, i.e. favoring one child over another or changes in household finances between the time of each marriage.

4.1 Theory of Sanskritization

One of the most important and widely cited theories in Indian sociology is that of Sanskritization, first theorized by the prominent sociologist M.N. Srinivas (Srinivas, 1956). The theory states that Brahmins, the priestly caste, traditionally carried out a number of practices that reinforced their high caste status, such as payment of dowry, vegetarianism and particular forms of dress (e.g. the sacred thread). He argues that lower castes began to emulate these practices, including dowry, in order to increase their perceived rating in the caste hierarchy. As a result, this emulation and status seeking explains the rise of dowry.

Other authors have disputed this explanation. Rao (1993a) argues that the increase in status conferred by dowry could not possibly justify the enormous dowry payments made by households. Caldwell et al. (1983) points out that in Indian popular conception, demands for dowry are viewed in a negative light, and thus are unlikely to confer higher caste status. There are two additional limitations of the Sanskritization theory. First, it does not explain why lower castes would begin to emulate higher castes in this regard at this point in Indian history rather than before or after. Second, as we show, it is inconsistent with the changes in the size of dowry payments over time. However, no previous research has attempted to test this theory quantitatively.

Sanskritization theory predicts that empirically, one should observe a high prevalence of dowry in Brahmin marriages, with an increase in the prevalence of dowry in non-Brahmin marriages as they start to emulate Brahmin practices. Figure 7 reports the proportion of marriages with dowry practices across four groupings of caste over the last century, as reported in the ARIS-REDS data. As seen in the figure, Brahmins were slightly more likely to pay dowry in the pre-1940 period, and began increasing dowry payments slightly before other groups. However, there is little differential prevalence and change in dowry practices across caste groups over time, certainly different from the implied patterns of Sanskritization.

Although it is obvious from figure 7, we formally test for and reject differential trends in the prevalence of dowries across other caste groups relative to the Brahmins (appendix figure A20). This does not invalidate Sanskritization generally, since it may be that lower caste individuals began to emulate upper caste practices other than dowry, such as vegetarianism. But it clearly

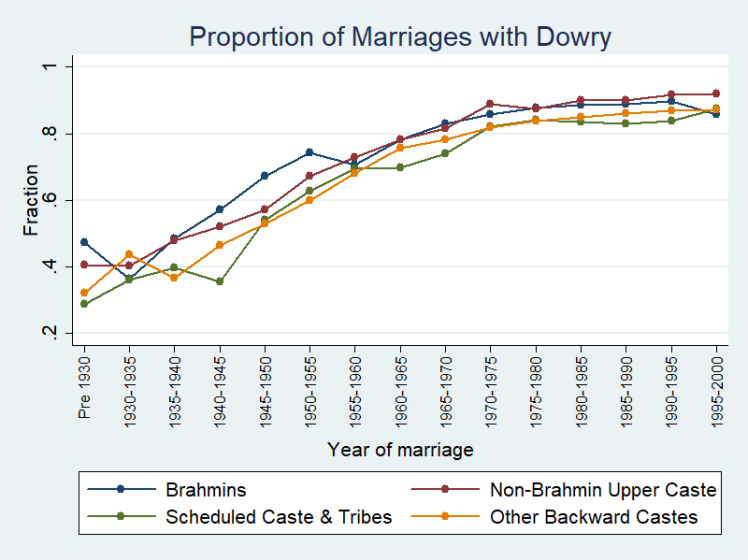

Figure 7: Prevalence of Dowry by Decade Across Caste Groups

does not explain changes in dowry practices over the past century.[27]

4.2 Modernization in Caste-Based Societies

Anderson (2003) constructs a novel theoretical framework to explain why in some societies, dowry declines in response to modernization, while in others, dowry size increases. In this model, the key difference is the society's status marker: in societies with status based on wealth (e.g Europe), dowry declines in response to modernization, whereas in caste-based societies (e.g. India), modernization prompts an increase in dowry size. In caste-based societies, caste is an inherited, hierarchical characteristic that is independent of wealth, with a vertical ranking of caste groups. In wealth-based societies, wealth is the determinant of status. On the marriage market, women are assumed to have vertical preferences over the status marker of potential grooms, where they always prefer to marry a higher status man. In wealth-based societies, this is a wealthier groom, whereas in caste based societies, this is a groom from a higher caste (as well as wealthier grooms). On the other hand, men only care about possible dowry payments and are indifferent to the status of potential brides.

[27] One concern is that this finding is due to recall bias. Recall bias would not produce these patterns unless Brahmins were systematically more forgetful than other caste groups; if anything, higher literacy rates among Brahmins should point in the opposite direction. It is also inconsistent with the slightly lower prevalence of missing responses to questions about dowry among Brahmins relative to other caste groups (appendix figure A25). It may be the case that Brahmins were more likely to practice dowry in the pre-1920 period, but even if so, that would still have been a minority of Brahmin marriages.

Modernization is assumed to involve two components: increasing average wealth and increasing income disparities within status (or caste) groups. The key result of the paper is that in a caste-based society, the increase in within-caste wealth dispersion along with an increase in average wealth leads to an increase in dowry payments. Intuitively, this is because there is an increasing stock of wealthy low-caste women, who attempt to lure the lower quality high-caste men to marry them by offering higher dowries.[28] This need not lead to inter-caste marriages, since high-caste women match the low-caste offers, but this causes inflation in the dowries that the high-caste men receive. As the income distributions of different caste groups begin to overlap and equalize, caste endogamy starts to break down.

We examine if this theory can rationalize the trends in dowry and inter-caste marriage in India by checking whether its predictions match the observed patterns in our data. First, the theory has strong testable predictions for how dowry payments evolve within caste groups. Within a particular caste grouping, dowry inflation for the highest quality grooms should be lower than that of the lowest quality caste grooms among castes ranked above them. It also should not be the case that lower caste grooms recieve higher dowries than those ranked higher on the caste hierarchy. The theory would predict that in such a case, the bride paying a high dowry to her low caste groom could profitably deviate to someone of the higher caste status. Neither of these patterns hold in the data. Figure 8 plots the percentiles of the dowry distribution within a given year. Among high quality, low caste grooms (here plotted as the 75^{th} percentile of the distribution), dowry inflation is substantial throughout the time period and dowries are reasonably large. For low quality, higher caste grooms (plotted as the 35^{th} percentile of the dowry payment distribution), the inflation is slower and they receive dowries that are lower in absolute magnitude. Thus it does not seem that the data match the predictions of the model.

Second, the model provides an additional testable prediction on dowry payments in inter-caste marriages. When a low-caste woman marries a high-caste man (i.e. she "marries up"), the theory would predict that she pays a larger dowry to outcompete women from the higher castes. On the other hand, if a high-caste woman marries a low-caste man (i.e. she "marries down"), she should be compensated by paying a smaller dowry. Taking the sample of women in the 1999 REDS data who marry outside their caste grouping, we test whether women who marry up or down are differentially likely to pay dowry and/or pay different amounts of dowry in regression 1:

[28] Although they would prefer high quality high-caste men, they would not be able to compete over them.

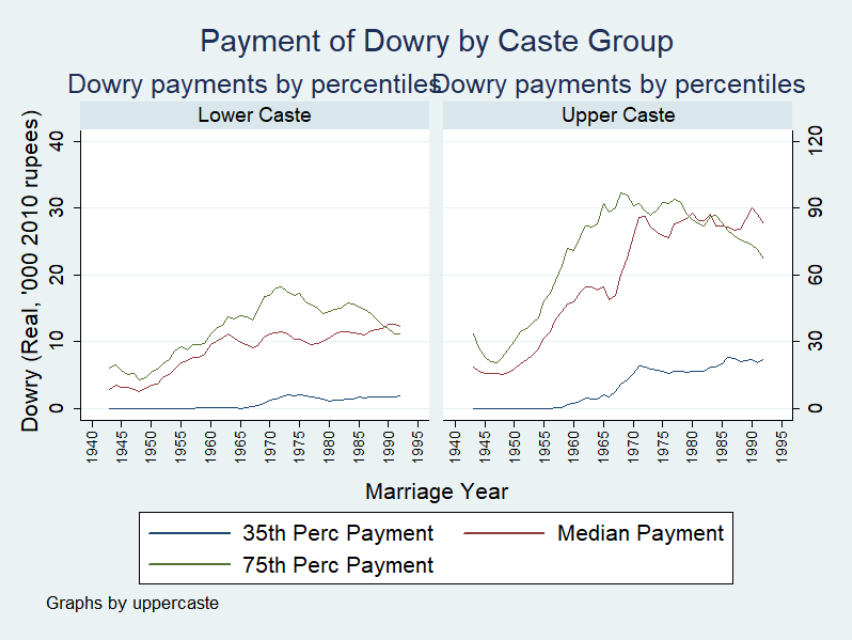

Figure 8: Dowry Payment By Caste and Year of Marriage

$$y_{mct} = \alpha + \gamma_{ct} 1^U_{mc} + X_{mct}\theta + \delta_t + \delta_c + \delta_d + \epsilon_{mct} \tag{1}$$

where y_{mct} is either a dummy for female payment of dowry or log of real dowry paid. 1^U is a dummy that takes the value 1 if the female married a male of a higher caste group ("marries up") and is 0 if the female marries down. δ_t,δ_d and δ_c are marriage year, district and jati (of the groom) fixed effects to take into account the underlying variation in dowry payments across time, districts and castes. As seen in table 2, the coefficient γ_{ct} is not distinguishable from 0, indicating that there is no statistical difference in either the probability of paying dowry or the dowry paid if the woman marries a man of a higher caste as opposed to marrying a man from the lower caste, contradicting the predictions of the modernization theory.

While this is a consistent theory of how dowry practices might evolve over time, the predictions of the theory do not match the Indian context. This is likely because the model assumes that female preferences over the caste of their groom are strictly vertical, i.e. they prefer to marry higher caste men (and avoid marrying lower caste men because that would degrade their caste status). In the Indian context, any vertical preferences have been shown to be dwarfed by the much stronger preferences for marrying within one's caste group. Banerjee et al. (2013) use data from responses

	Dowry Paid	Dowry Paid	ln(Dowry)	ln(Dowry)
Higher Caste Groom	0.00136	-0.0108	-0.0308	-0.0542
	(0.00970)	(0.0109)	(0.0579)	(0.0647)
Observations	2932	2778	1516	1401
Controls	Yes	Yes	Yes	Yes
District FE	Yes	No	Yes	No
Year FE	Yes	No	Yes	No
Caste FE	Yes	No	Yes	No
Household FE	No	Yes	No	Yes

Table 2: Dowry in Inter-caste Marriages

to matrimonial advertisements to estimate the underlying preferences over prospective partners. Their estimation allows for both vertical and within-caste preferences over caste, but they only find evidence of within-caste preferences. If within-caste preferences dominate, then even as low caste families become wealthier, they will continue to prefer to marry within caste. High caste brides will not need to increase their dowry payments to compete and so dowry inflation will be limited.

4.3 Marriage Squeeze Hypothesis

A prominent strand of the literature attributes changes in dowry payments to a *"marriage squeeze"*. This theory, initially proposed by Caldwell et al. (1983), notes that in India, as in many other developing countries, women typically marry at younger ages than men. As a result, if the population is growing, the cohort of men who are on the marriage market will be smaller than the cohort of women. This squeeze causes increased competition over the limited men and could lead to an increase in dowry size. Figure A18 plots the aggregate population sex ratio (number of women divided by number of men) and marriage market sex ratio (number of marriage age women (age 15-25) relative to number of marriage age men (age 20-30)). Although the aggregate sex ratio does not change, there is an increasing squeeze in number of women on the marriage market thorough 1970, followed by a decline. This pattern is consistent with the timing of increases and decreases in mean dowry.

Rao (1993b) first tested the theory empirically using the ICRISAT data on marriages between 1921 and 1983 and population data from the Indian census. He finds a strong positive relationship between the marriage market sex ratio and dowry size, as predicted by the theory. However, Edlund (2000) takes the same data and finds no relationship between sex ratio and dowry size. The limited

sample size of the ICRISAT data (127 marriages) makes it difficult to draw definitive inferences from either analysis. In addition to the possible empirical refutation by Edlund (2000), there are theoretical reasons why dowry size may not respond to population growth and sex ratio imbalances. Anderson (2007b) argues theoretically that sex ratio imbalance may not lead to changes in dowry size, but could instead shift the age at which individuals marry. For example, if there is a relative scarcity of men at typical ages of marriage, women may marry later and men may marry earlier to restore balance. Bhaskar (2016) extends her model in a dynamic setting to examine the impact of permanent and transitory shocks to cohort sizes on the equilibrium age gap and size of dowry. Under a standard transferable utility framework, he shows that the age gap will be immune to systematic growth in cohort sizes, while there will be an increase (decrease) in dowry paid with positive (negative) growth of cohort size. On the other hand, he argues that a transitory shock to cohort size could affect both age gaps and dowry sizes of the nearby cohorts and the dynamic adjustment will depend on how sensitive cohorts are to age considerations. It is unclear which theory is empirically relevant in this context, since individuals may not know whether a shock is transitory or permanent, which is crucial for their reaction.

We begin by replicating and extending the specifications from Rao (1993b) and related papers using the much larger sample of the REDS data. To determine sex ratio, we use Indian census data beginning in 1911. The marriage market relevant sex ratio is defined as the number of women aged 10-25 in a district divided by the number of men aged 20-35 in the same district. These ranges were selected to cover the range of ages at which marriages occur in our data (and are similar to that of Rao (1993a)).[29] In appendix section C, we measure sex ratio using two other data sets (1991 census data containing the full age distribution at the district level, and pooled data from the National Sample Survey that constructs sex ratio at the state-identity group level) and discuss the merits of each approach. Our findings across all three data sets are consistent, increasing confidence in the findings.

We include marriage-level characteristics such as education levels and *jati* of the bride and groom to control for compositional differences in the marriage market over time and add time

[29]The census data is only available each decade, but contains information about the number of individuals in five year age bins. This permits a more fine grained analysis, where observed marriages are placed in five year bins (e.g. 1971-1975, 1976-1980, etc.). We use the census to determine the number of marriage aged men and women in those bins.

(δ_t) and district (δ_d) fixed effects to account for unobserved heterogeneity across time and space. Identification comes from differential changes across districts over time in their (marriage-relevant) sex ratios. We therefore estimate the following specification:

$$y_{mdt} = X_m\beta + \gamma SexRatio_{dt} + \delta_d + \delta_t + \epsilon_{mdt} \tag{2}$$

where y_{mdt} is the outcome variable. We report the results in table 3. Column (1) examines the relationship between sex ratio and whether any dowry was paid. Therefore, the dependent variable is a dummy that takes the value 1 if dowry is paid and 0 otherwise. Column (2) examines the relationship between sex ratio and the size of dowry payment.

	(1) Dowry Dummy	(2) Log Dowry
Sex Ratio	0.0781	-0.276
	(0.116)	(0.327)
Observations	45414	36546
Controls	Yes	Yes
District FE	Yes	Yes
Year FE	Yes	Yes

Standard errors are clustered at the district level.
* p<0.10, ** p<0.05, *** p<0.01

Table 3: Sex Ratios and Dowry

The data offers little support for the marriage squeeze hypothesis. Neither the extensive nor intensive margin of dowry payment seem to be related to the marriage market sex ratio in the district, as the theory would predict. Given these results, there must be some other valve through which the pressure in the marriage market is relieved. In line with the theories of Anderson (2007b) and Bhaskar (2016), we turn our attention to age of marriage. We run the same specification as in (2), but now using the marriage ages of males and females and the marriage gap (defined as the difference between the male-female marriage age) as dependent variables. The results are reported in table 4. With the approximately 5000 marriages in the REDS for which the age at marriage for brides and grooms is observed,[30] we find results consistent with the predictions of marriage age shifts. Column (1) shows that as there is a relative increase in the number of women at traditional marriageable ages relative to men (i.e. an increase in the marriage-adjusted sex ratio), the difference

[30]In the REDS survey instrument, the question on age of marriage was asked on the roster administered only to the household head and hence we have the age of marriage for the household head and his/her spouse only, reducing the sample size.

in ages between brides and grooms decreases. The coefficient estimate implies a doubling of the sex ratio would cause a 3.6 year decrease in the age gap between men and women getting married that year. Interestingly, the smaller gap is entirely accounted for by a decrease in men's age of marriage. Families may be unwilling to postpone their daughter's marriage due to the high premium placed on female virginity and fears of loss of sexual purity (Jensen and Thornton, 2003). Anukriti (2013) also finds that sex ratio imbalance leads to shifts in the age of marriage using three rounds of the National Family Health Survey, although with a different definition of sex ratio.

	(1) Age Gap	(2) Age of Marriage (Male)	(3) Age of Marriage (Female)
sexratio	-3.558***	-3.197*	0.298
	(1.179)	(1.791)	(1.318)
Observations	5073	5155	5075
Controls	Yes	Yes	Yes
District FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Standard errors are clustered at the district level.
* p<0.10, ** p<0.05, *** p<0.01

Table 4: Sex Ratios and Marriage Age

Although both of these regression results suggest that sex ratio does not drive dowry, sex ratio is endogenously determined by factors that may also affect dowry size and/or the age of marriage. In particular, it will be a function of population growth and differential mortality across genders/sex-selection. In order to control for sex selection and differential mortality, we add controls for the contemporaneous sex ratio, i.e. the ratio of women aged 15-20 to men aged 15-20. This removes variation due to this differential mortality, and doing so has no effect on the results. After removing this variation, the key source of variation comes from population growth. The determinants of such growth occur at least 15 years prior to the marriage market of interest, and as a result, confounds that affected population growth of current participants, such as an increase in wealth, will already have been priced into the previous set of marriages. Since identification comes from differences in dowry from year to year, this should not create a problem for our estimates. This will only be a problem if the factor affecting population growth has a lagged effect on marriage markets aside from the distribution of individuals within that market.

4.4 Search and Matching Models with Heterogeneous Partner Quality

One hypothesis that has received surprisingly little attention in the economics literature is that changes in dowry practices are the result of changes in the distribution of groom quality. A number of sociologists have stated this informally, noting that India experienced an increase in salaried, non-agricultural jobs in the 1930s and 1940s, such as in the public sector (Srinivas, 1984; Caplan, 1984). These grooms tend to earn higher and more stable wages, and thus are more desirable for households to match with their daughters. If dowry acts as a market clearing price, an influx in higher quality grooms could cause dowry to rise. This is a feature of Anderson (2003) and Anderson and Bidner (2015), where marriage result from a two-sided matching market. In Anderson and Bidner (2015), both sides of the market have a single dimensional measure of quality and transfers may be made at the time of marriage. Households can invest in the quality of their offspring prior to entry on the marriage market, and as the returns to investment/education for men increase, this leads to larger dowry payments to compete over these men. This theory is intuitively plausible, particularly since the timing of initial increases in dowry payments lines up with improvements in the pool of salaried and educated grooms. We find that increases in the educational attainment of males of a caste group in a state are related to large increases in dowry size (appendix table A1).

As a test of this theory, we first examine whether higher quality grooms actually receive larger dowries. Since we do not directly observe earnings in the REDS data, we use groom's education as a proxy measure of their quality. For a marriage m of a groom in household h of district d at time t within the five-year band f, we regress:

$$y_{mhdt} = \beta_1 e_{mhdt} + \delta_{hf} + \delta_t + \epsilon_{mhdt} \tag{3}$$

where $y_{mhdt}$is the dowry payment (in real terms), e^i_{mhdt} is the years of education of the groom, and δ_t and δ_{hf} are marriage year and household-five year fixed effects respectively. Marriage year fixed effects account for changes in aggregate dowry payments over time across all of India. The household portion of the fixed effects removes differences in dowry payment related to household wealth or tastes over dowry that could be related to groom quality. However, household wealth and attitudes are not static, and might change over time in a manner related to groom quality. For example, a household may be relatively poor and have poorly educated grooms in the 1940s, but

become wealthier and educate its grooms in the 1960s. To account for this, we include time-varying household fixed effects: e.g. one for marriages in that household between 1940-1944, another for marriages between 1945-1949, and so forth. This is an extremely strict set of fixed effects, where identification comes from whether differences in dowry between brothers married within the same five-year window are related to differences in their education. β_1 is not a causal estimate of the effect of education, since there may be omitted qualities of a groom related to education that allow them to command a higher dowry. But since those omitted characteristics are additional measure of quality that are presumably positively related to education, that is still exactly what we seek to measure.

The identifying assumption is that there are not other within-household changes over that five year span that are simultaneously related to dowry and quality of the groom. That is particularly plausible given that education is completed well prior to a male entering the marriage market[31], so any wealth shocks that affected the education of one brother will either have dissipated or, if persistent, will affect both when they later are on the marriage market.

Table 5 indicates a substantial premium for education among grooms, where each additional year of groom's education increases the size of their dowry by Rs. 1,219. Given that the median dowry size is Rs. 25,062, that is substantial.

	(1) Dowry	(2) Dowry	(3) Dowry
Groom Education (Years)	1.219**		0.949***
	(0.430)		(0.264)
Groom Education Percentile		16.30**	2.914
		(6.758)	(4.478)
Observations	39865	31904	31707
Year FE	Yes	Yes	Yes
Village FE	No	No	No
Household FE	No	No	No
Household-5 year FE	Yes	Yes	Yes

Standard errors clustered at the state level

Table 5: Education and Dowry

There is a large premium to being a higher quality groom, but this does not necessarily mean that increases in educational attainment cause dowry inflation. That will depend on whether the returns to groom quality are due to the groom's relative position in the distribution of grooms or the absolute quality of the groom. It may be the case that a bride/bride's family only value a groom's

[31] Education is completed prior to marriage for 97.4% of males in our sample

ranking in the distribution of grooms, such as if they see grooms as conspicuous consumption, and getting a more educated groom matters only for signaling their status. In this case, a shift in the distribution of groom quality need not have an aggregate effect on dowry size, even if that shift is non-uniform. For example, if the distribution changes so that an 80^{th} percentile groom has 10 years instead of 6 years of education, the groom at that percentile will still get the same amount.[32] On the other hand, it is conceptually plausible that there are absolute returns to a bride marrying a more educated groom. In a matching model, the premium commanded by a groom depends on the returns to marrying them relative to outside options. If a groom is able to earn more due to his education, then the bride will gain utility relative to her outside option, and potentially be willing to pay more. If there remains a mass of grooms at the lowest quality rung (landless, illiterate laborers), increasing mass to the right of the distribution will lead to overall inflation in dowry amount.

By taking advantage of the segmented nature of Indian marriage markets, it is possible to separately identify whether dowry is affected by either the absolute or relative value of education. We estimate:

$$y_{mhdt} = \beta_1 e_{mhdt} + \beta_2 \tilde{e}_{mhdt} + \delta_{hf} + \delta_t + \epsilon_{mhdt} \tag{4}$$

where $\tilde{e}_{mhdt}$ is the groom's percentile rank in the educational distribution of grooms on the marriage market in the year of his marriage. This is calculated by combining multiple rounds of the National Sample Survey, which contains individual-level data on the age, gender, education, state, and broad identity grouping (caste and religion) of all individuals in the respondents household. Using the NSS data on the education distribution of grooms on the same market as this groom in his year of marriage, we calculate his relative position (e.g. for a scheduled caste Hindu groom who was married in 1984 in Rajasthan, this will be his position in the distribution of education for scheduled caste Hindu males in Rajasthan aged 18-25 in 1984). As before, we include household-five year fixed effects δ_{hf}, which remove time-varying characteristics of the family that might be related to groom quality, and identification comes from the differences in the absolute and relative level of education of brothers married within 5 years of one another.

[32] This additionally assumes that the distribution of groom quality is continuous rather than lumpy.

We are able to separately identify β_1 and β_2 due to the segmentation of marriage markets within identity groups. Suppose that there are two sets of brothers from identity groups A and B. In both sets of brothers, the first brother has 8 years of education and the second has 10 years of education. If their relative ranking were defined with respect to the aggregate educational distribution, then years of education would be almost entirely collinear with percentile ranking. However, since they are on the marriage markets of different identity groups, they face different distributions of educational attainment. It might be the case that the brothers from group A are in the 30^{th} and 50^{th} percentile of their distribution, while those from group B are in the 70^{th} and 80^{th} percentile of their distribution. Intuitively, β_1 and β_2 are derived from taking the difference in dowry between the two brothers and seeing if that is related to either the difference in their years of education (2 years) or their relative ranking in the educational distribution (20 percent and 10 percent).

Table 5 shows that the absolute return to education is the driver of dowry receipt. Relative rank is individually a strong predictor of dowry size, but once the two are put into a horse-race, relative rank drops to insignificance. This confirms the role of education in driving dowry inflation. While this is completely consistent with a model of dowry as groom-price, the question becomes why there was a decline in dowry payments at the upper percentiles of the distribution in the post-1975 period, with stability among the bottom 40% of the distribution of payments (see figure 5). Under the model of Anderson and Bidner (2015), such a decline will occur as the economic value of investing in women increases, such as due to increases in labor market opportunities. In future drafts of the paper, we will test whether this can be responsible for the decline.

Another possible explanation is that marriage markets are better described with a search model rather than a matching model. In this model, potential grooms are matched with a potential bride in every time period, and both sides bargain over the match. If they come to an agreement over dowry, the match is made; if not, then they rematch with another possible partner in the next period. As the distribution of groom quality takes on increasing mass in the right tail, this initially leads to an increase in dowry in high end marriages, with little effect on the lower end of the market. That is because high quality grooms are more desirable than the expected outcome of re-matching in the market, and so brides prefer to match with them at high value dowries (even if the dowry they recieve decreases as their numbers increase). Since this is an increasing proportion of marriages, average dowry increases. However, as the right tail takes on increasing mass, this can

actually decrease the dowry paid to higher quality grooms. That is because there is now a higher probability of a bride meeting a high quality groom if she rematches, and so the groom has to give up more of his surplus to convince her to match with him. Such a model is consistent with our empirical results and is consistent with the decline in high-end dowries: these high quality men can longer command princely sums because there are more of them around. Future drafts of the paper will also discuss this possibility.

4.5 Additional Theories

The previous sections gave empirical evidence on prominent theories of dowry. This section contains brief discussions of other possibilities. For example, one possible mechanism causing a rise in dowry is changes in the price of commodities typically given as part of dowry. This is rationalized by a model of dowry in which dowry has a social signaling value, and there is anchoring on the quantity rather than value of a good. For example, suppose members of one's community recently gave a 10 grams of gold as dowry and due to "keeping up with the Joneses", there is a social cost to not giving at least 10 grams, rather than the monetary equivalent. If there is a rise in the price of dowry commodities (in this case, gold) relative to inflation, this could increase the real value of dowry.

The most obvious commodity in which such inflation could occur is gold, which is given in over 90% of marriages with dowry. Bhalotra et al. (2016) find a positive relationship between gold prices and dowry size, though that is only an auxiliary portion of their analysis, which is predominantly focused on how dowry price affects sex selection. They also use the 1999 REDS data, but restrict their analysis to daughters of the household head, for a sample size of between 2239 to 4201 marriages between 1970 to 1999. They find a strong positive relationship between gold prices and dowry, and in some specifications, are unable to reject a 1:1 relationship between increases in gold price and the corresponding increase in dowry. Yet international gold prices were almost completely stable between 1945 and 1967, and then increased in the post-1970 period. This is the opposite of the pattern of dowry increases.

Another possibility is that increased urbanization may have lead to increases in dowry size. Urbanization may lead to changes in norms and adoption of new practices, such as dowry. It may also be related to changes in the composition of the workforce and pool of grooms/brides. Taking

district level urbanization data between 1911-1991, we find no relationship between urbanization and dowry adoption or dowry size (appendix table A2). Another theory is that marriage markets may have expanded over time, increasing competition for high quality grooms. However, we find no change in the distance between the natal homes of brides and grooms over time, which is inconsistent with this hypothesis.

Finally, it is possible that the causes of changes in dowry are non-economic in origin, such as changes in the relative gender status of men and women, shifts in social norms, and increased desire for social signaling. Such explanations are difficult to test quantitatively due to the lack of historical data on such factors. It is also scientifically unsatisfying to posit spontaneous shifts without specifying a reason – for example, why should norms change at this point in time, and not before? It could be that economic changes facilitated changes in norms, and thus are the root cause: for example, it may be that changes in incomes/groom differentiation drove changes in norms around taking/giving dowry. Unfortunately we lack the data for a more rigorous test of this hypothesis.

5 Conclusions

In this paper, we provide the first quantitative documentation of the prevalence and evolution of dowry in India, as well as a rigorous set of tests of the underlying factors behind dowry. We are able to reject a number of major theories of dowry, as well as accepting a theory that can rationalize the key empirical patterns. This theory has implications for design of anti-dowry policy, suggesting that efforts to improve the position of women relative to men, particularly on characteristics that may be relevant on the marriage market, will lead to lower dowry payments. Increasing female labor force participation may be particularly effective in this regard.

The results of this paper are suggestive of many interesting paths for future research. First, in considering policy aimed to fight dowry, future work could consider the distributional consequences of removing dowry. It may be that dowry encourages investment in male children that is a net positive for the economy, particularly given low rates of female labor force participation. On the other hand, if there are low returns to investing in female children, this could reinforce inequality between sexes and the aforementioned low female participation in the formal labor force. Second,

it may be that norms based approaches would be more effective in discouraging payment of dowry. Future work could consider running such campaigns and understanding how these norms operate, particularly how they may diffuse within a social network (e.g. campaigns focused on parents of male children versus female children). Third, the distributional consequences of dowry are relevant for considering how to construct anti-dowry political coalitions. If dowry continues to persist in India, as our findings suggest that it will, much more work remains to be done in design of policy responses to mitigate its most negative consequences.

References

Alfano, Marco, "Daughters, dowries, deliveries: The effect of marital payments on fertility choices in India," *Journal of Development Economics*, 2017, *125*, 89–104.

Anderson, Siwan, "Why Dowry Payments Declined with Modernization in Europe but Are Rising in India," *Journal of Political Economy*, 2003, *111* (2), 269–310.

_ , "The Economics of Dowry and Brideprice," *Journal of Economic Perspectives*, 2007, *21* (4), 151–174.

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

_ **and Chris Bidner**, "Property Rights Over Marital Transfers," *The Quarterly Journal of Economics*, 2015, *130*, 1421–1484.

Anukriti, S., "Sex Ratios, Schooling and the Marriage Market in India," 2013.

Anukriti, S, Sungoh Kwon, and Nishith Prakash, "Dowry: Household Responses to Expected Marriage Payments," *Working Paper*, 2017.

Arunachalam, Raj and Suresh Naidu, "The Price of Fertility: Marriage Markets and Family Planning in Bangladesh," *Unpublished Paper*, 2008.

_ **and Trevon D Logan**, "On the heterogeneity of dowry motives," *Journal of Population Economics*, 2016, *29* (1), 135–166.

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune, "Marry for what? Caste and mate selection in modern india," *American Economic Journal: Microeconomics*, 2013, *5* (2), 33–72.

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

Behrman, Jere, Andrew Foster, Mark Rosenzweig, and Prem Vashishta, "Women's Schooling, Home Teaching, and Economic Growth," *Journal of Political Economy*, 1999, *107* (4), 682–714.

Bhagwati, Jagdish N and Thirukodikaval Nilakanta Srinivasan, "Foreign trade regimes and economic development: India," 1975.

Bhalotra, Sonia, Abhishek Chakravarty, and Selim Gulesci, "The Price of Gold: Dowry and Death in India," *IZA DP No. 9679*, 2016.

Bhaskar, Venkataraman, "The Demographic Transition and the Position of Women: A Marriage Market Perspective," *Working Paper*, 2016.

Billig, Michael S, "The marriage squeeze on high-caste Rajasthani women," *The Journal of Asian Studies*, 1991, *50* (02), 341–360.

_ , "The marriage squeeze and the rise of groomprice in India's Kerala state," *Journal of Comparative Family Studies*, 1992, pp. 197–216.

Bloch, Francis and Vijayendra Rao, "Terror as a bargaining instrument: A case study of dowry violence in rural India," *American Economic Review*, 2002, *92* (4), 1029.

Botticini, Maristella and Aloysius Siow, "Why Dowries?," *American Economic Review*, 2003, *93* (4), 1385–1398.

Caldwell, John C, Palli Hanumantha Reddy, and Pat Caldwell, "The Causes of Marriage Change in South India," *Population Studies*, 1983, *37* (3), 343–361.

Caplan, Lionel, "Bridegroom Price in Urban India: Class, Caste and 'Dowry Evil' Among Christians in Madras," *Man*, 1984, *19* (2), 216–233.

Choo, Eugene and Aloysius Siow, "Who marries whom and why," *Journal of Political Economy*, 2006, *114* (1), 175–201.

Corno, Lucia, N. Hildebrandt, and Alessandra Voena, "Age of Marriage, Weather Shocks, and the Direction of Marriage Payments," *Working Paper*, 2017.

Dalmia, Sonia, "A hedonic analysis of marriage transactions in India: Estimating determinants of dowries and demand for groom characteristics in marriage," *Research in Economics*, 2004, *58* (3), 235–255.

_ **and Pareena G Lawrence**, "The Institution of Dowry in India: Why It Continues To Prevail," *The Journal of Developing Areas*, 2005, *38* (2), 71–93.

Deolalikar, Anil B and Vijayendra Rao, "The demand for dowries and bride characteristics in marriage: empirical estimates for rural south-central India.," in "Presented at the Annual Meeting of the Population Association of America, Denver, Colorado, April 30 - May 2, 1992" [Unpublished] 1992. Presented at the Annual Meeting of the Population Association of America Denver Colorado April 30-May 2 1992. 1995.

Dugar, Subhasish, Haimanti Bhattacharya, and David Reiley, "Can't Buy Me Love? A Field Experiment Exploring the Trade-off Between Income and Caste-Status in an Indian Matrimonial Market," *Economic Inquiry*, 2012, *50* (2), 534–550.

Edlund, Lena, "The Marriage Squeeze Interpretation of Dowry Inflation: A Comment," *Journal of Political Economy*, 2000, *108* (6), 1327–1333.

_ , "The Price of Marriage: Net vs. Gross Flows and the South Asian Dowry Debate," *Journal of the European Economic Association*, 2006, *4* (2/3), 542–551.

Eika, Lasse, Magne Mogstad, and Basit Zafar, "Educational assortative mating and household income inequality," *Working Paper*, 2014.

Epstein, T. Scarlett, *South India: Yesterday, Today and Tomorrow*, Macmillian, 1973.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos, "Marry your like: Assortative mating and income inequality," *The American Economic Review*, 2014, *104* (5), 348–353.

Jensen, Robert and Rebecca Thornton, "Early female marriage in the developing world," *Gender & Development*, 2003, *11* (2), 9–19.

Kapadia, Karin, "Marrying Money: Changing preference and practice in Tamil Marriage," *Contributions to Indian Sociology*, 1993.

Logan, Trevon D. and Raj Arunachalam, "Is There Dowry Inflation in South Asia?," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 2014, *47* (2), 81–94.

_ **and** _ , "On the heterogeneity of Dowry Motives," *Journal of Population Economics*, 2015.

Rajaraman, Indira, "Economics of Bride-Price and Dowry," *Economic and Political Weekly*, 1983, *18* (8), 275–279.

Rao, Vijayendra, "Dowry Inflation in Rural India: A Statistical Investigation," *Population Studies*, 1993, *47* (2), 283–293.

_ , "The Rising Price of Husbands: A Hedonic Analysis of Dowry Increases in Rural India," *Journal of Political Economy*, 1993, *101* (4), 666–677.

Rosenzweig, Mark R and Oded Stark, "Consumption Smoothing , Migration , and Marriage : Evidence from Rural India," *Journal of Political Economy*, 1989, *97* (4), 905–926.

Sautmann, Anja, "Partner Search and Demographics: The Marriage Squeeze in India," *Working Papers*, 2011.

Shafir, Eldar, Peter Diamond, and Amos Tversky, "Money Illusion," *The Quarterly Journal of Economics*, 1997, *112* (2), 341–374.

Singh, Vir Bahadur, *Economic History of India, 1857-1956*, Allied Publishers, 1965.

Smith, Herbert L., Sharon J. Ghuman, Helen J. Lee, and Karen Oppenheim Mason, "Status of Women and Fertility," 2000.

Srinivas, M.N., *The Remembered Village*, University of California, Berkeley Press, 1976.

_ , *Some Reflections on Dowry*, OUP India, 1984.

Srinivas, Mysore Narasimhachar, "A note on Sanskritization and Westernization," *The Journal of Asian Studies*, 1956, *15* (04), 481–496.

A Appendix Figures and Tables

A.1 Marriage Decision Tables and Figures

A.1.1 Selection of Husband

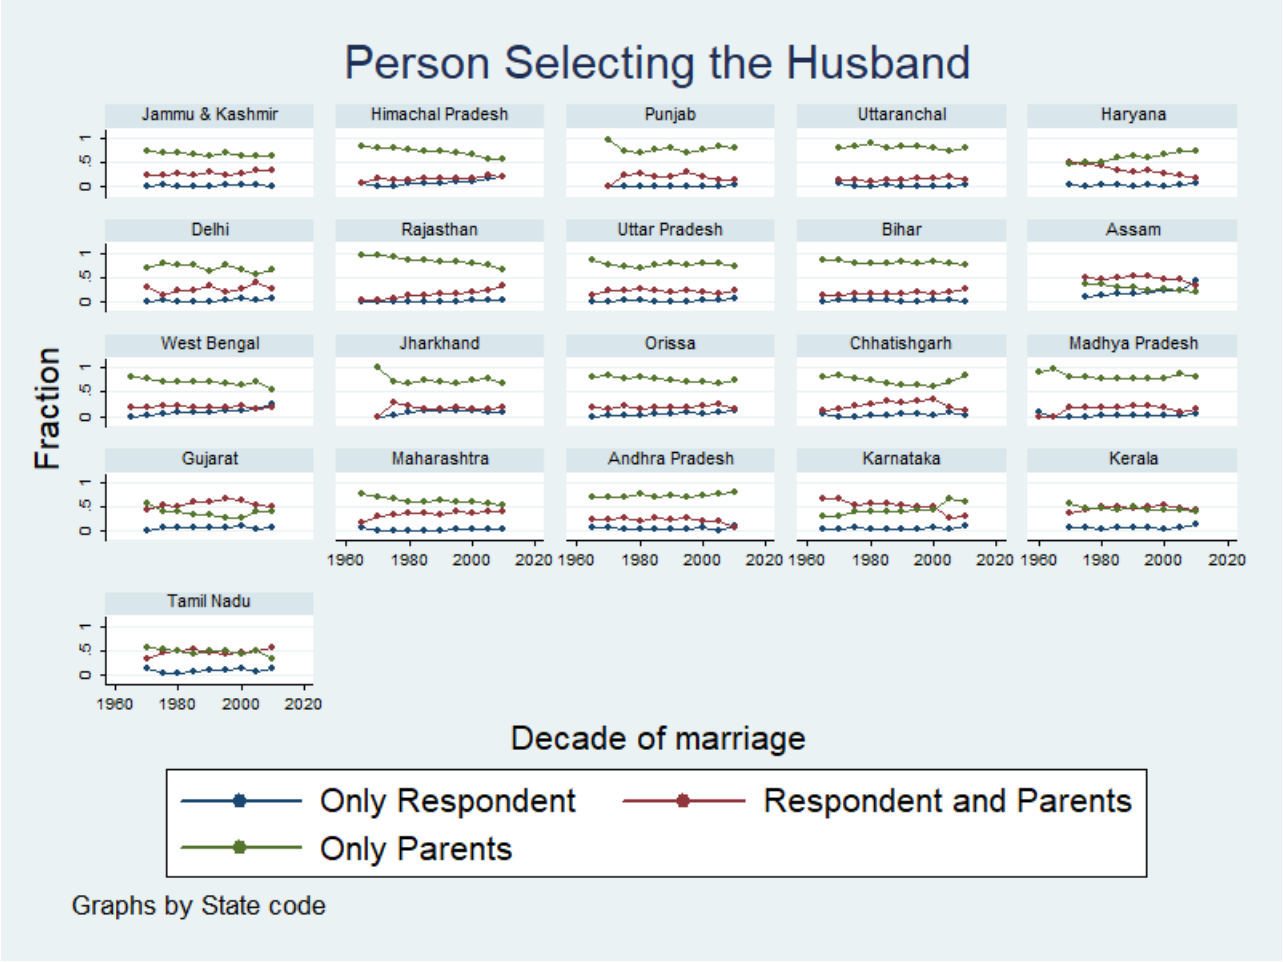

Figure A1: Person Selecting Husband, by State (IHDS, 2005)

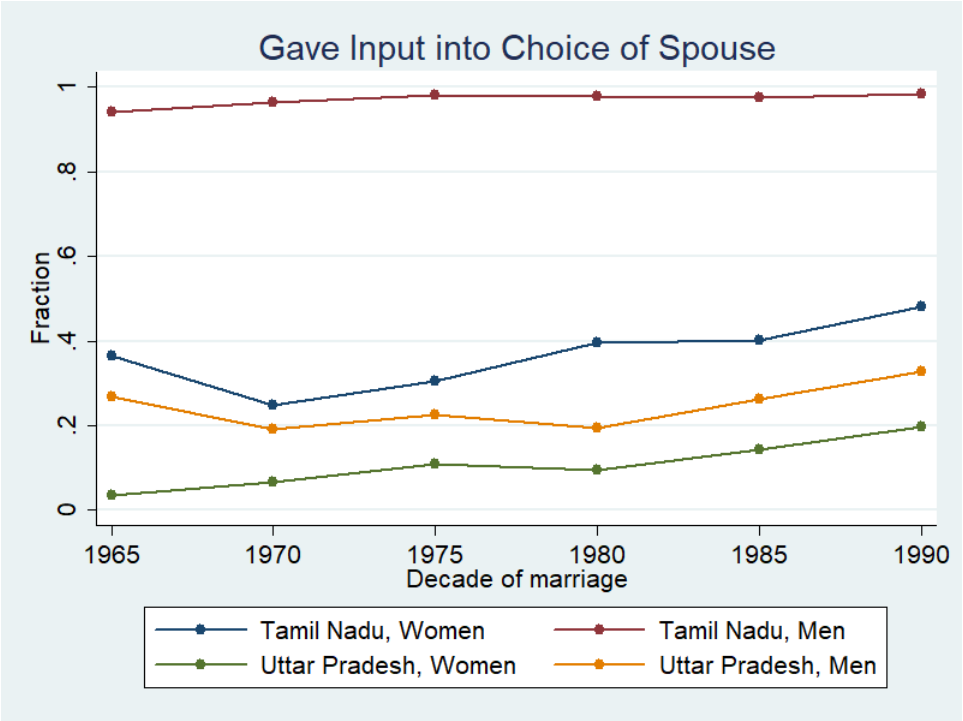

Figure A2: Influence over Choice of Spouse (SWAF)

In the SWAF data, both husbands and wives are asked if they gave input into who was selected as their partner, while in the IHDS data, only the wives are asked. Between 1960 and 2005, only 5-10% of women were the main decision maker in the choice of their husband, with little change over the period. In 73.7% of marriages, the parents were the sole decision-maker, and 17% were decided upon by the child and parents jointly. There are large differences depending on location and child gender. In Tamil Nadu, 98% of men and 37% of women have at least some say into the choice over their partner, while in Uttar Pradesh, this is true for 24% of men and 11% of women.

A.1.2 Time of Meeting Husband

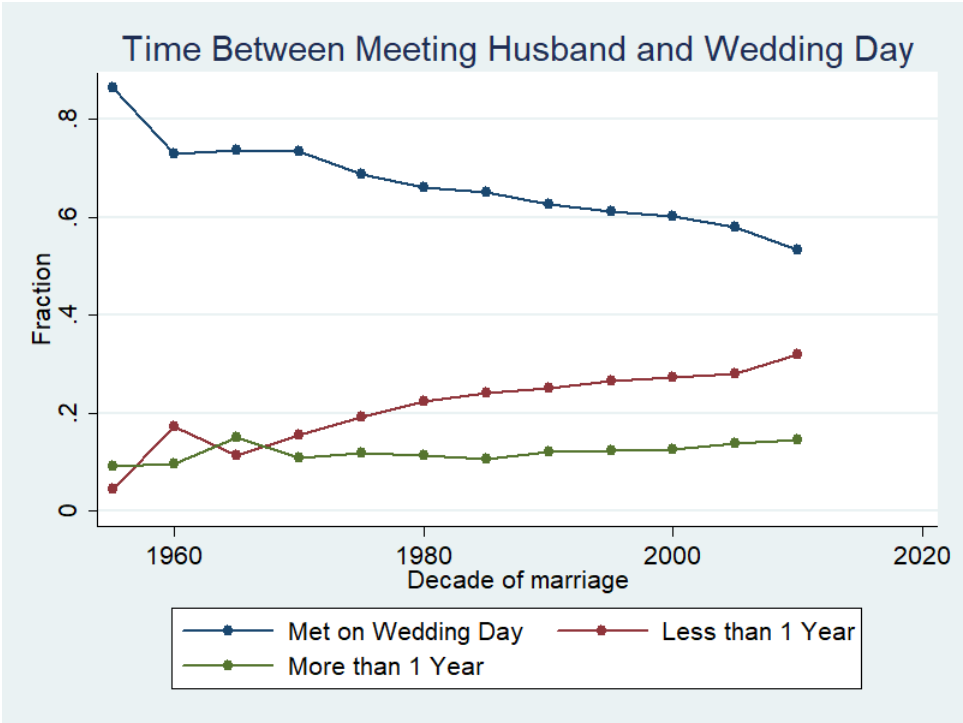

Figure A3: Bride and Groom First Met on Wedding Day (IHDS)

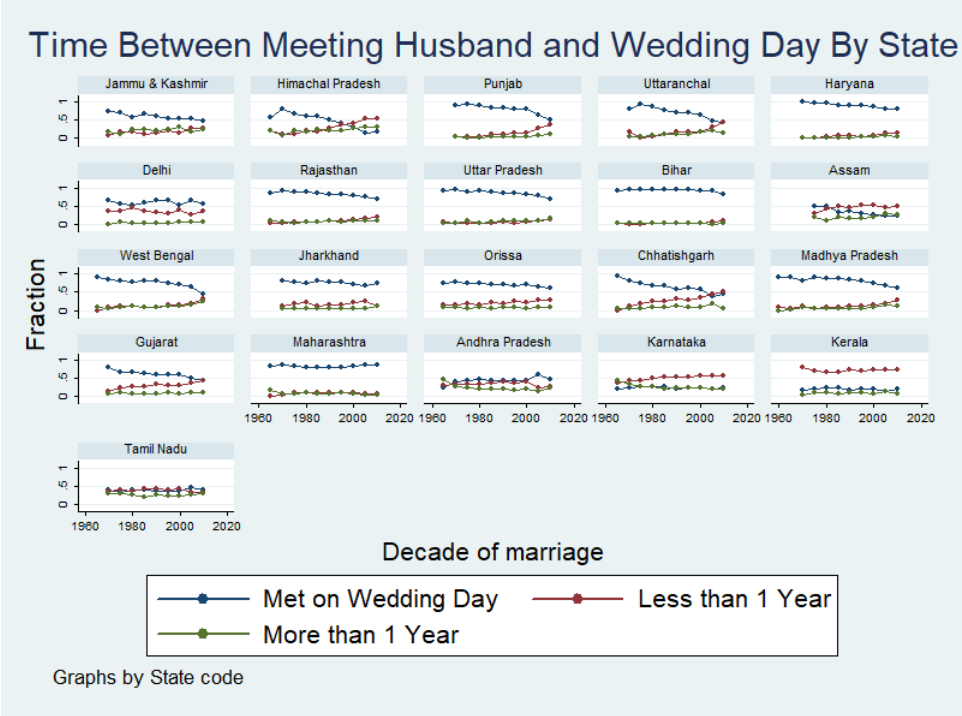

Figure A4: Time Between Meeting Husband and Wedding Day, By State (IHDS)

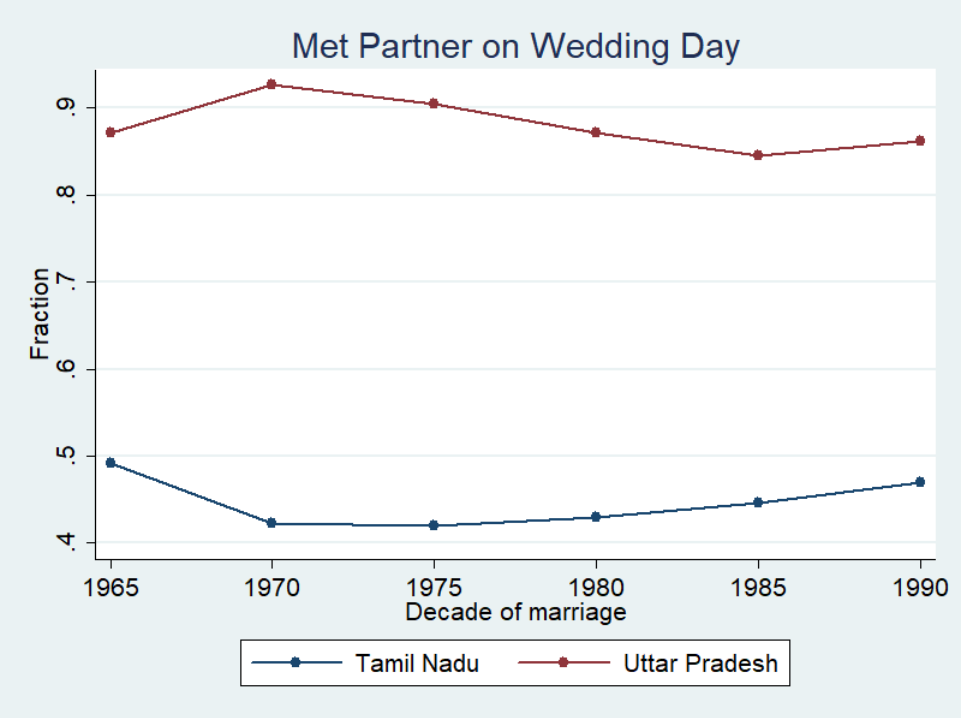

Figure A5: Bride and Groom First Met on Wedding Day (SWAF)

A.1.3 Location of Partner's Family

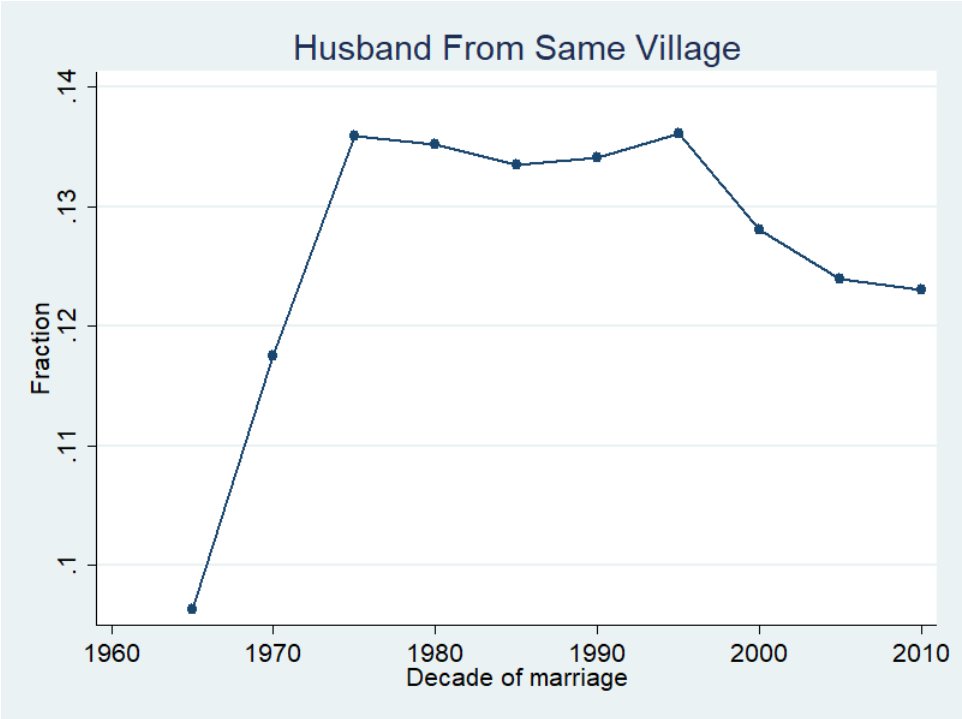

Figure A6: Location of Husband and Wife's Natal Homes (IHDS)

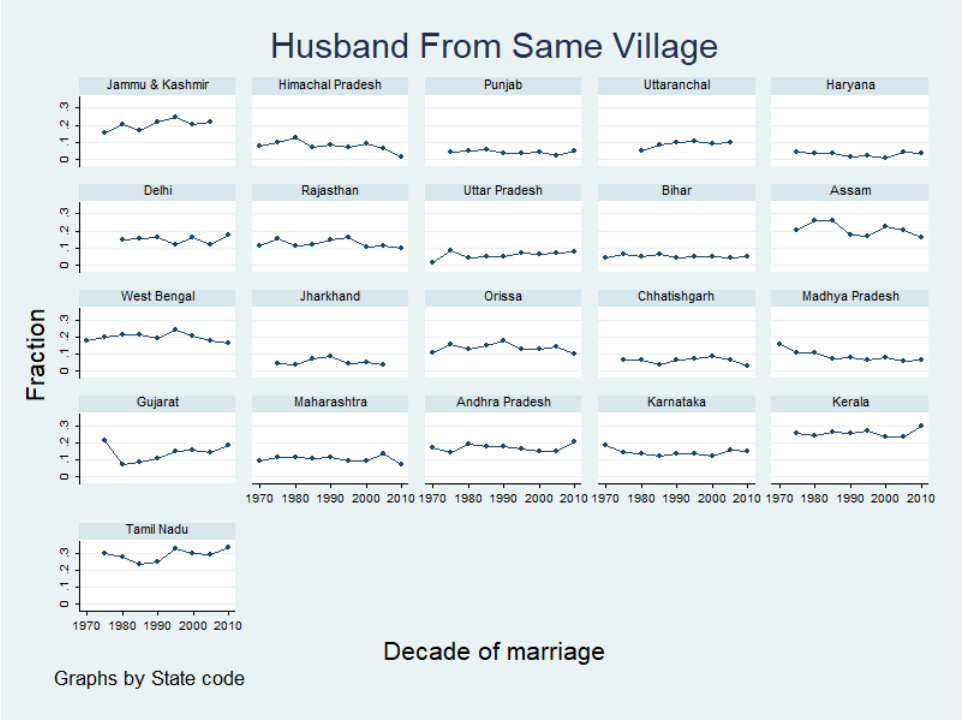

Figure A7: Location of Husband and Wife's Natal Homes (IHDS)

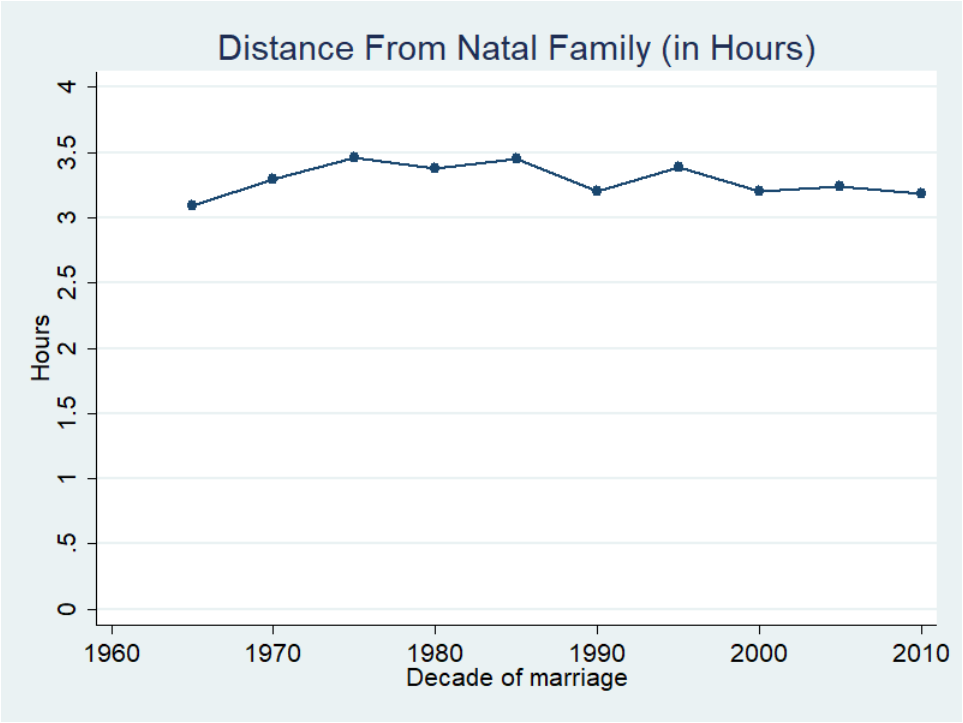

Figure A8: Distance between Natal Households of Brides and Grooms (IHDS)

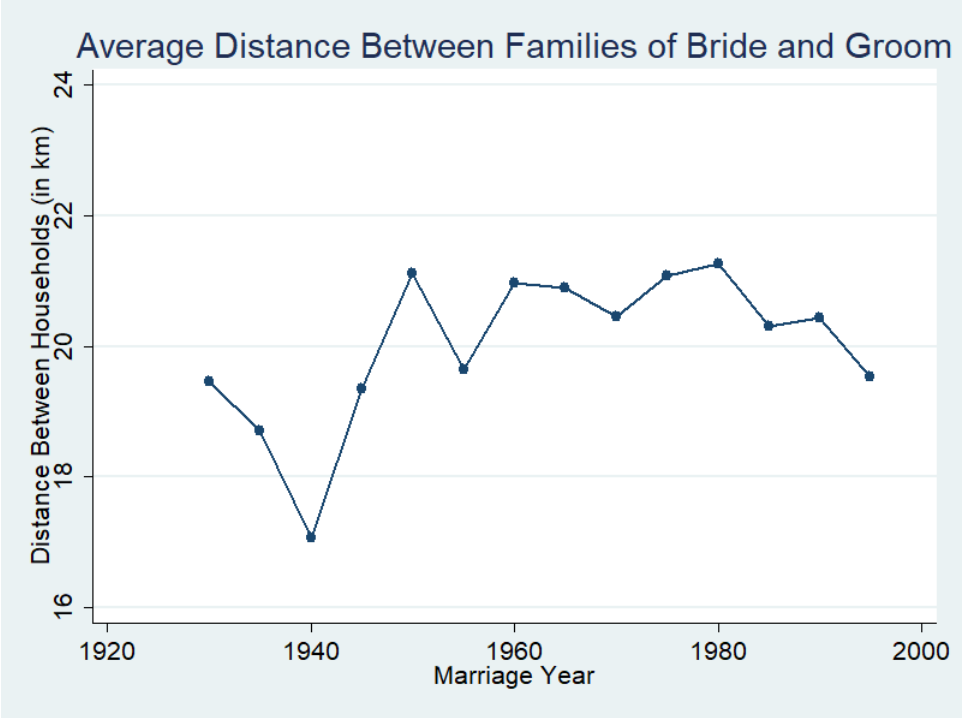

Figure A9: Distance between Natal Households of Brides and Grooms (REDS)

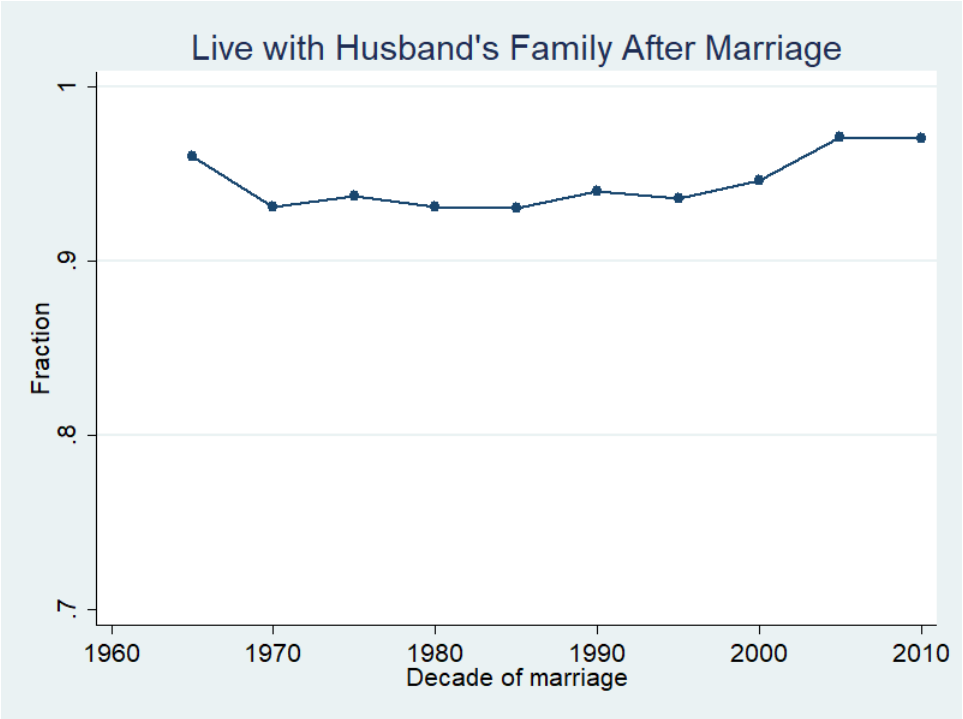

Figure A10: Living Situation of Couple After Marriage (IHDS)

A.1.4 Inter-caste Marriage

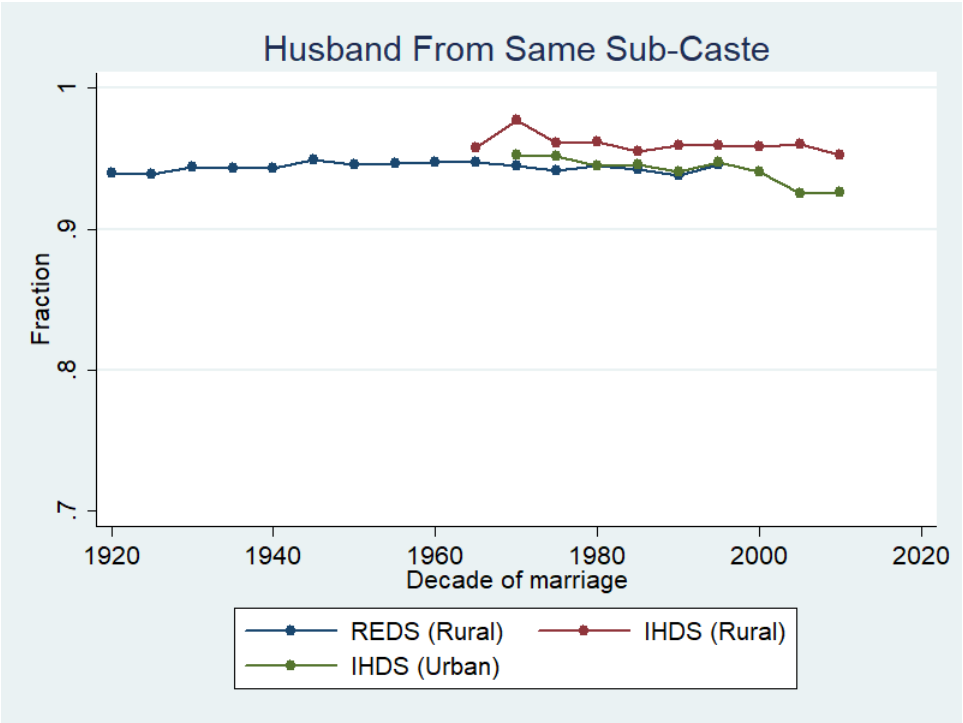

Figure A11: Prevalence of Inter-jati Marriage from 1960-2005

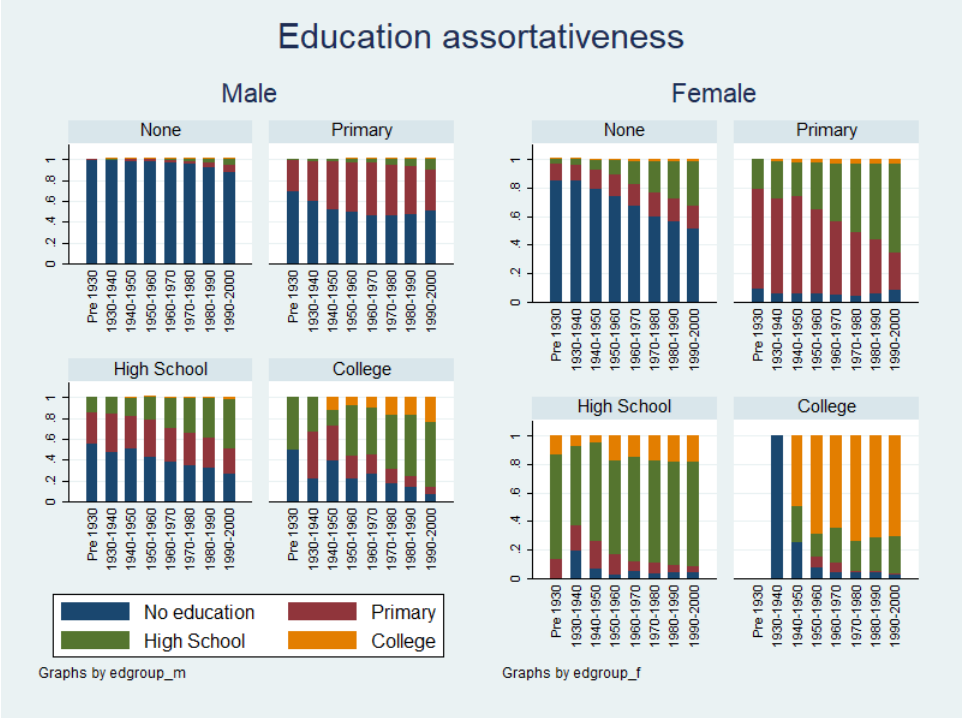

Figure A12: Assortative matching across education groups

A.2 Dowry-Related Tables and Figures

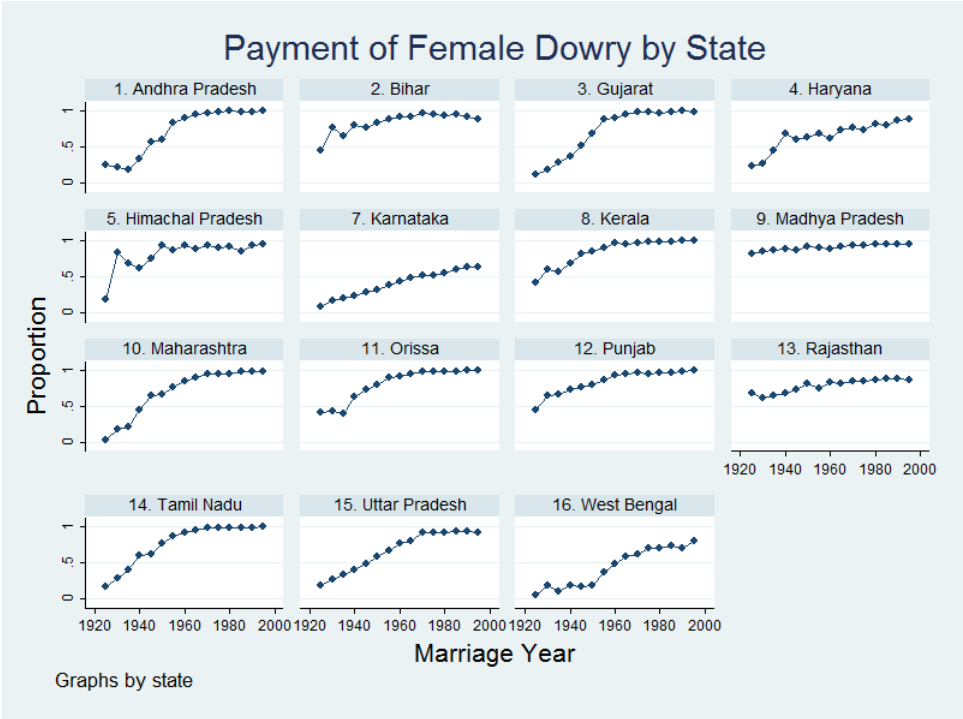

Figure A13: Prevalence of Dowry Across States

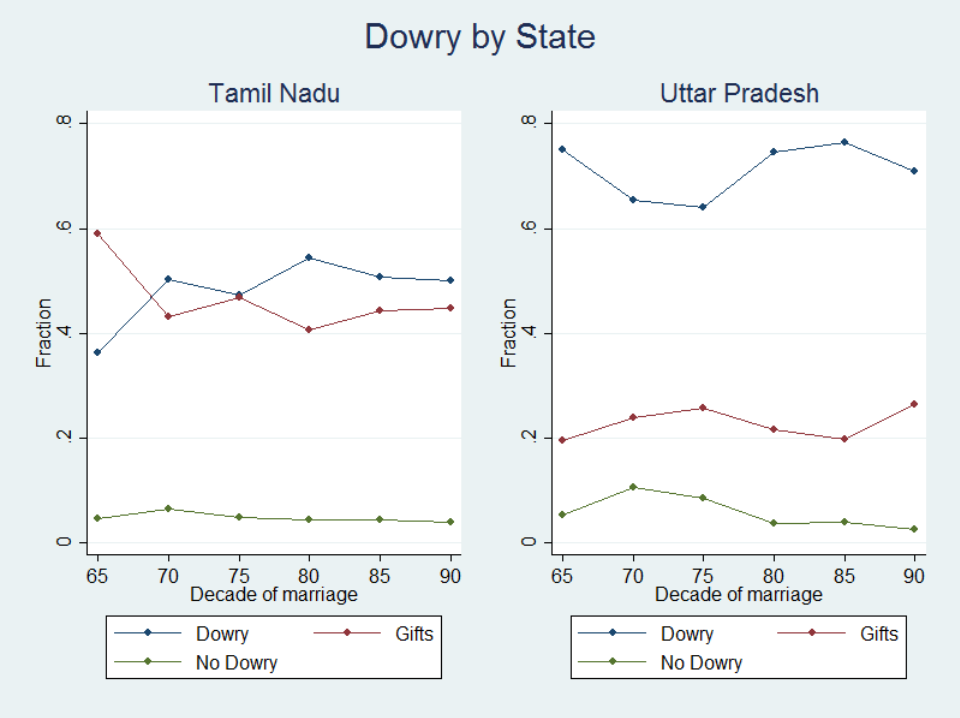

Figure A14: Prevalence of Dowry in the SWAF

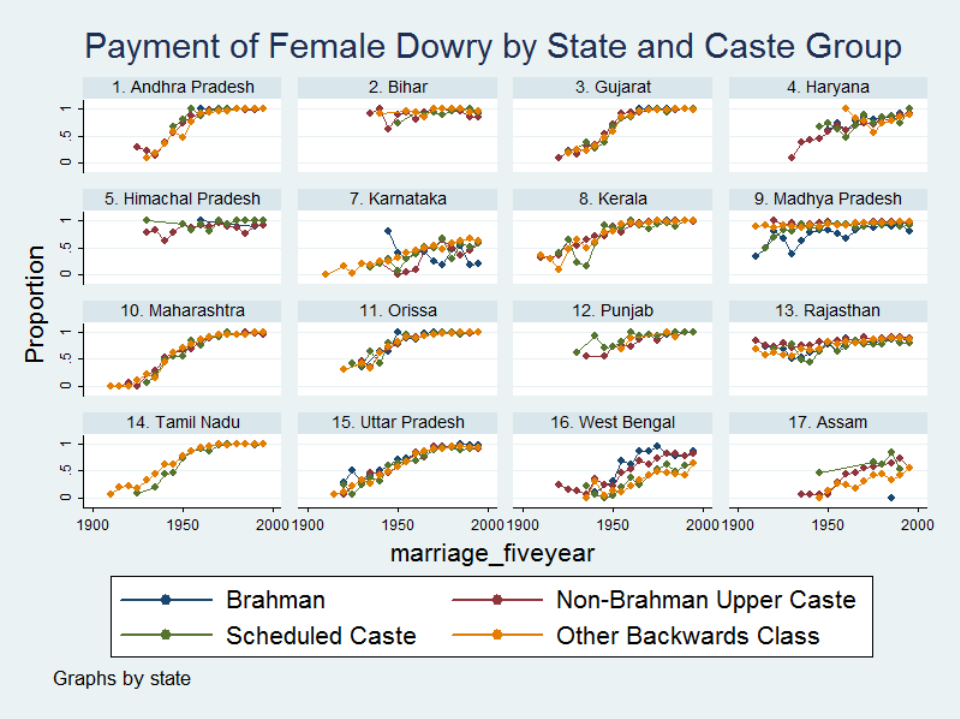

Figure A15: Prevalence of Dowry by Decade and Caste, Across States

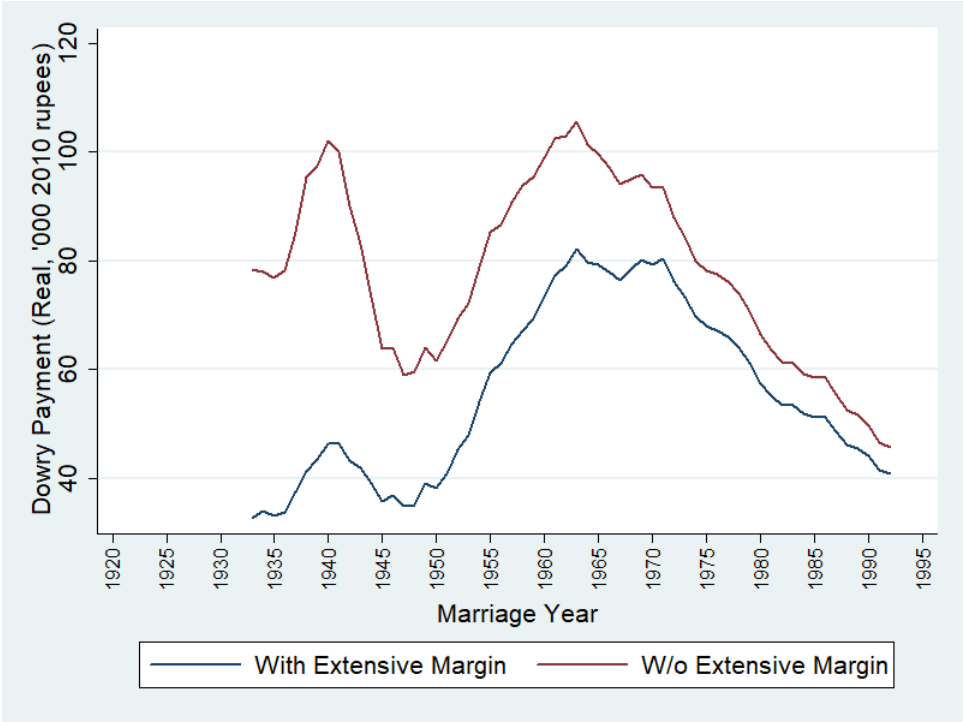

Figure A16: Mean Dowry Payment By Year of Marriage

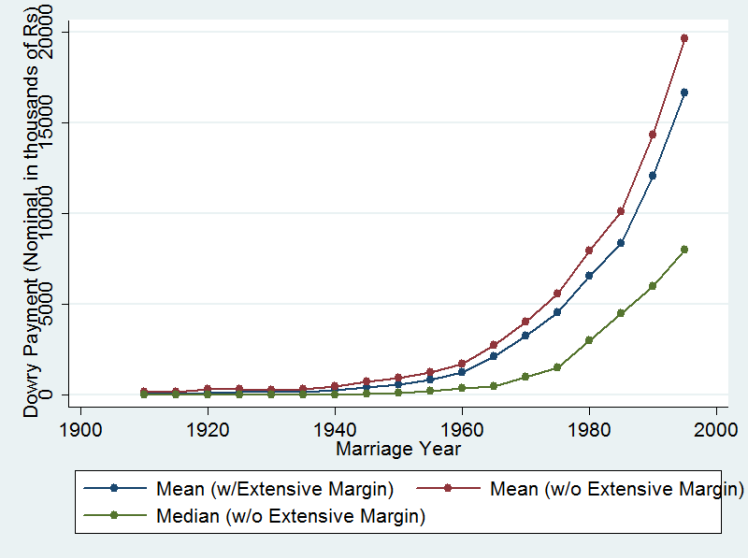

Figure A17: Mean Nominal Dowry Payment By Year of Marriage

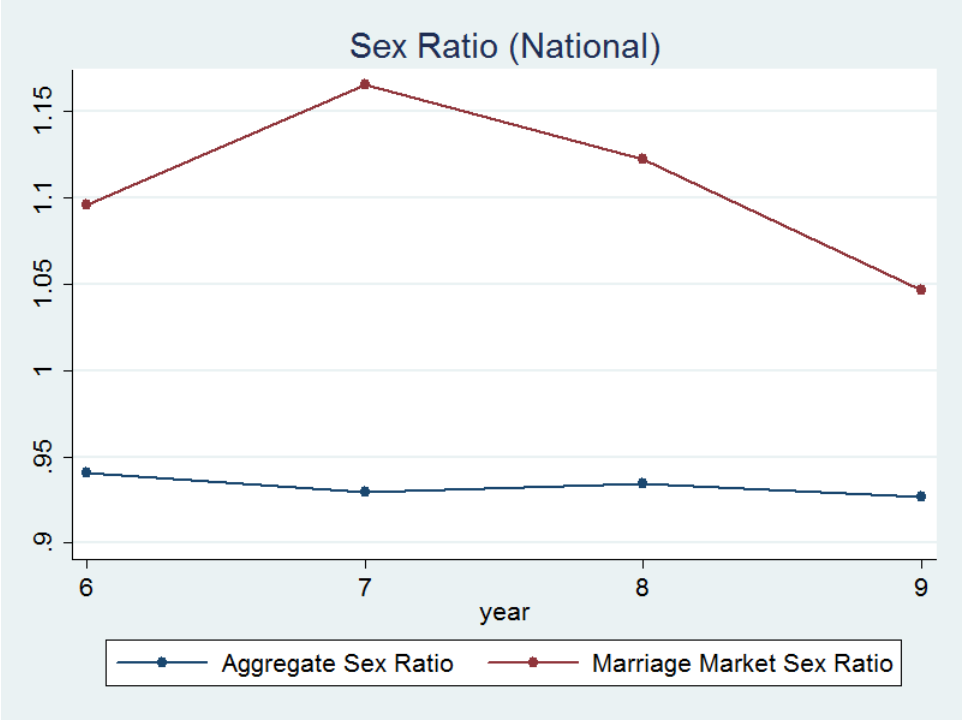

Figure A18: Sex Ratio Over Time (Census)

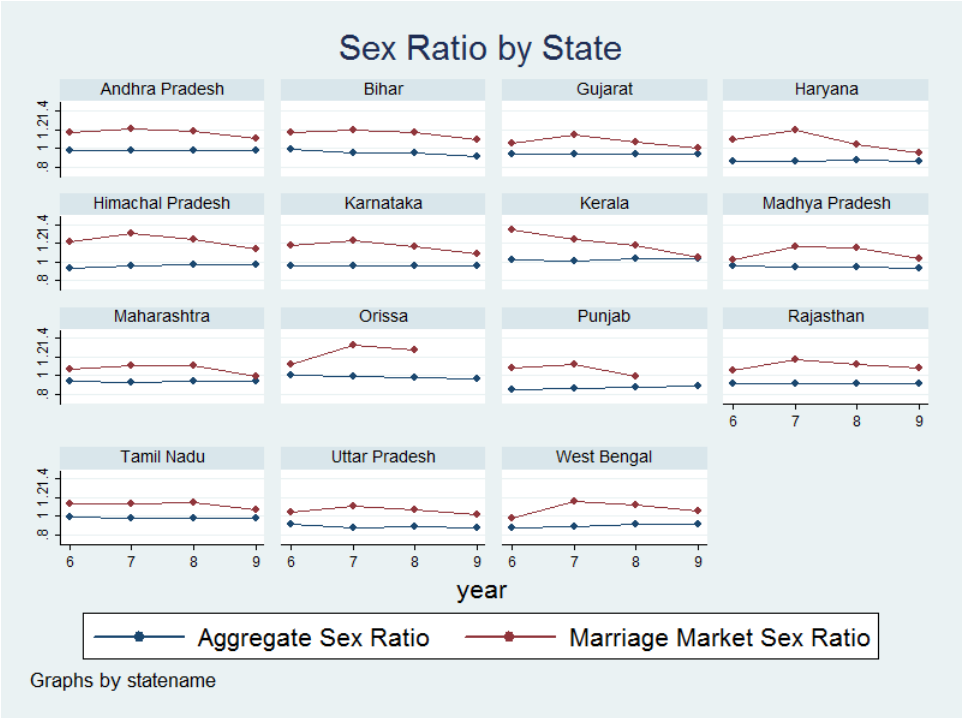

Figure A19: Sex Ratio Over Time by State (Census)

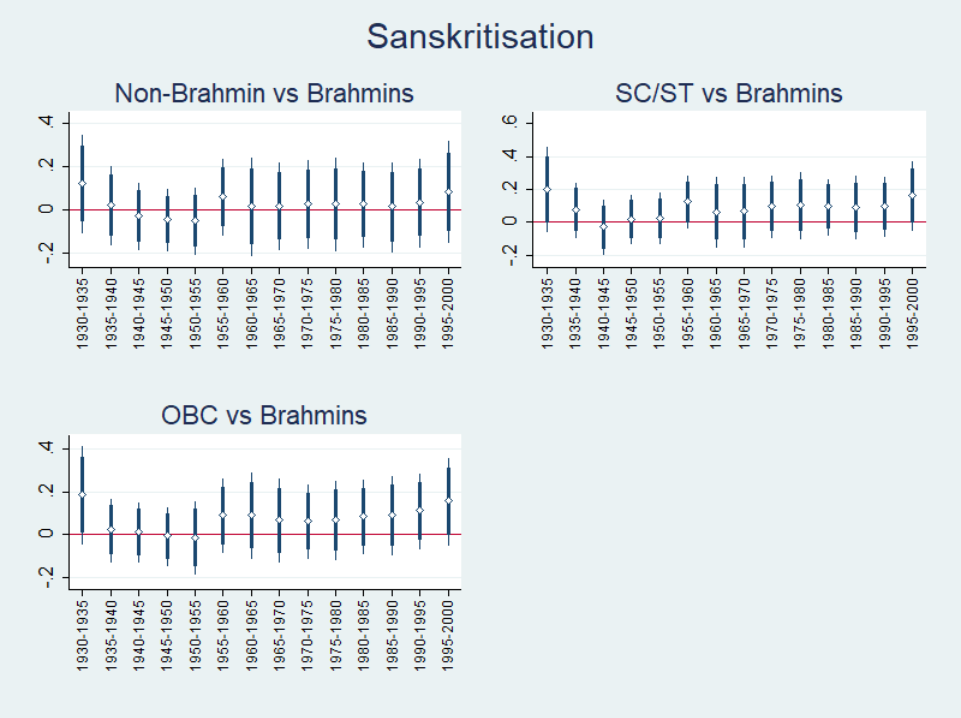

This figure plots the coefficients from the regression

$$D_{hct} = \alpha + \sum_{c} \beta_c 1_C(Caste_{hc} = c) + \sum_{t} \gamma_t 1_T(Year_{ht} = t) + \sum_{c} \sum_{t} \delta_{ct} 1^C_{hc} \text{ x } 1^T_{ht} + \theta_d + \epsilon_{hct} \quad (5)$$

where the dependent variable D_{hct} is a dummy that takes the value 1 if a household from caste group c paid any dowry in year of marriage t and 0 otherwise. $1_C(Caste_{hc} = c)$ is a dummy that takes the value 1 if the household is of caste c and 0 otherwise; $1_T(Caste_{ht} = t)$ is a time dummy that takes the value 1 if the marriage was in year t is 0 otherwise. θ_d are district fixed effects that control for underlying variation in dowry prevalence across districts that does not change with time. The δ_{ct} coefficients report the differential rates at which a non-Brahmin caste group was likely to practice dowry *relative* to the Brahmins across time and these are reported in figure . Consistent with figure 7, we cannot reject the hypothesis that these coefficients are statistically different from 0, indicating no differential prevalence of dowry across Brahmin and non-Brahmin caste groups over time, thus contradicting the predictions of the Sanskritization theory.

Figure A20: Prevalence of Dowry Across Caste Groups (Regression Estimates)

	(1) Dowry
Fraction of Men with At Least Primary Education	82.64* (40.32)
Observations	39776
Year FE	Yes
State-Caste FE	Yes

Standard errors clustered at the state level

Table A1: Aggregate Educational Attainment and Dowry

	Paid Dowry	ln(Dowry)
urbanization	-0.0736 (0.219)	-0.818 (1.225)
Observations	44408	35356
Controls	Yes	Yes
District and Year FEs	Yes	Yes

Table A2: Urbanization and Dowry

B Data Details

B.1 Recall Bias

Given the lack of extant historical data on dowry, it is necessary to use retrospective data. However, it is possible that respondents are unable to properly recall dowry transactions, particularly given that many occurred well before the time of the survey. In this context, there are a number of reasons to think that recall will not be a major problem. Marriage is one of the main focuses of life in rural India, and events around that time are likely to be particularly salient. It is particularly unlikely that an individual would not recall whether or not they had paid dowry, even if they were not certain of the exact figure. But given the import of dowry payments for households, it is likely that respondents will be able to recall them; it would be like asking a home owner in the US what they paid for their house, a number likely to be remembered with ease. Even if errors occur, as long as they are not systematic, this will not cause problems for our estimates.

Another issue is that death of potential respondents could bias the estimates. This would be a particular problem if poorer households died at an earlier age, and so the sample of households in the earlier data were systematically different from the households in later data. Fortunately,

households are asked about all relevant marriages, even if those individuals have died, and so we have information on the full sample. For the majority of our time period of interest, there is a living member of the couple. Even for marriages occurring between 1940 and 1950, 47% of the relatives on whom the respondent is reporting are still alive, while that figure is 72% for marriages between 1950 and 1960, and 89% for marriages between 1960-1970.

We run a number of direct tests of recall bias using our data. The first test takes advantage of the panel nature of the REDS survey. Respondents were interviewed in 1999 and 2008, and asked similar questions about dowry payments at the time of marriage. If there were a systematic bias that emerges over time, we would expect it to emerge in comparisons between the 1999 and 2008 waves of the survey due to the gap in survey administration timings. The reported nominal dowry is shown over time for both waves of the survey.[33] We find that until 1955, the two surveys give nearly identical responses on average. A gap emerges during the 1940s, but this may be because the set of marriages being recorded in the 2008 wave does not include some during that time period. This can occur if the identity of the head of household transfers to a younger member of the household between the time of surveys.[34] However, we are encouraged that recall is fairly similar over most of time period, inconsistent with recall bias.

As a test for systematic bias, we take the average nominal dowry size for each state within five year bands for each wave of the survey (e.g. average for Maharashtra between 1960-65), and regress the 1999 wave average on the 2008 wave average. Table A3 shows that although the two are not identical, it is not possible to reject the null hypothesis that the two are on average the same.

	Dowry (1999)
2006 dowry_f	1.005***
	(0.0414)
Observations	286
State FE	Yes

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table A3: Dowry Payments by State-Year across the REDS rounds

As second check of the validity of retrospective dowry data, we turn to the SWAF. It conducts

[33]We prefer using nominal dowry for the comparison to avoid dealing with price deflators.

[34]The reason is that households are surveyed about the marriages of the parents, brothers, sisters, sons and daughters of the household head. If the identity of the household head moves to that of a younger member (e.g. due to death of an older household head), this will result in fewer people being asked about the earlier time period. If poorer households are systematically more likely to experience this, then a bias could emerge.

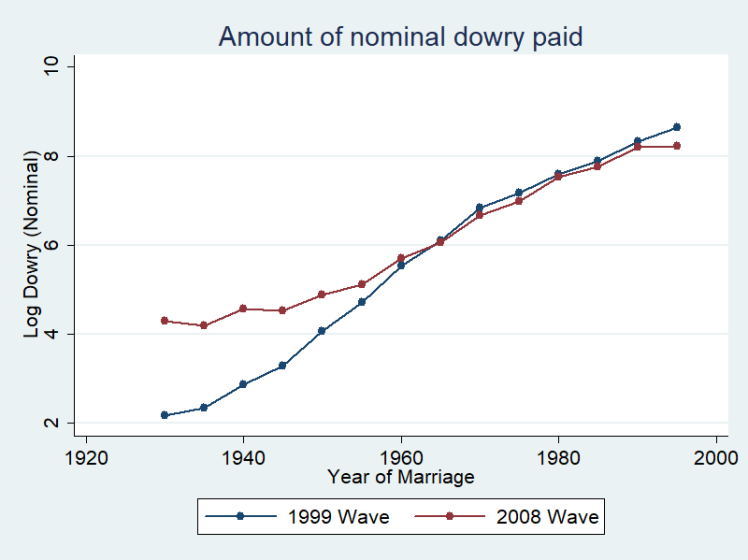

Figure A21: Dowry Recall Amount Between Waves (REDS)

separate interviews with husbands and wives, but asks a series of identical questions about whether particular items were given as part of the dowry in their marriage.[35] If recall is poor, then we would expect that the answers of the two parties would be poorly aligned: in the extreme where they were guessing on each response, they would match 50% of the time. Instead, their responses are nearly identical, matching in 87.8% of cases.

We can also check whether recall of dowry transactions seems to decline over time by checking how match rate varies depending on time elapsed since the marriage. Figure A22 plots the match rate at five year intervals after 1960. It is particularly notable that there is nearly no decline in recall matches between couples whose marriages occurred between 1970-1975 and those whose marriages were within a year of the survey, for whom the mismatch presumably does not stem from memory issues. There is a small decline for marriages occurring between 1965-1970, but the match rate remains exceptional. This increases our confidence in the validity of using retrospective dowry about dowry: while recall might be poor for less significant economic transactions, it does not appear to be so for one as important as dowry.

[35]The items included are land, jewelry, cash, a vehicle such as a car, a TV, furniture, a radio, utensils, a bicycle, livestock and clothing

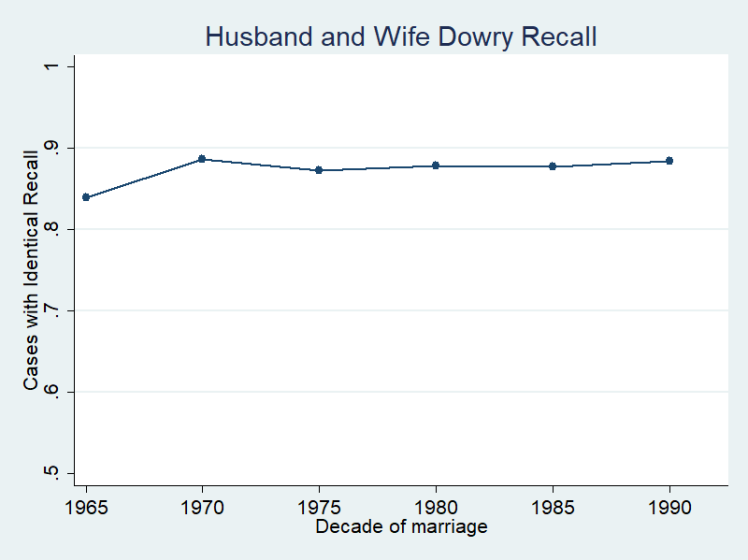

Figure A22: Dowry Recall in the SWAF

B.2 Comparison of the 1999 and 2008 waves of the REDS survey

In our paper, we use the REDS data, which aside from one other paper (Anukriti et al., 2017), has not previously been used to document changes in dowry practices over time. We use data from the 1999 wave survey rather than the 2008 wave as in Anukriti et al. (2017). This has some advantages and disadvantages. The major advantage is in the coding of non-dowry payment. In the 2008 wave of the REDS data, non-payment of dowry is coded as a missing value rather than zero, so it is impossible to distinguish cases where respondents were unable to recall dowry and paid zero dowry. That does not create problems for Anukriti et al. (2017), since they focus on the period after which dowry payment is nearly universal, and the focus of their paper is not on changes in dowry over time. However, for our paper, understanding the extensive margin is a critical component, and so the 1999 wave of the survey is a better choice. Second, since the 1999 wave of the REDS survey was conducted around 10 years before the 2008 wave, there is better coverage of the time period from 1930-1960. Since this is when many of the changes in dowry practices occurred, this is advantageous for this paper. Furthermore, if there are recall issues, it is helpful to get information from closer to the dates of marriage. The disadvantages of the 1999 wave of the REDS data is that in a few states, there is some inconsistent coding by surveyors. These challenges are discussed in section B.3.

B.3 Dowry Payments

As would be expected, recall of past dowry payments is imperfect, and so we encounter some missing dowry information in the data. This is complicated by state to state variation in the administration of questions, and in particular, how surveyors coded non-response and payment of zero dowry. In some cases, it is not possible to distinguish between zero and missing dowry payments even in the 1999 REDS data. In this appendix, we discuss these challenges and the coding decisions that we made, as well as the implications of taking different decisions.

Figure A23 show the data on female dowry payments by state. The dowry value from each marriage is coded as either missing, zero, or non-zero. Based on this, the states can be divided into four categories in terms of usability of dowry data. The first category is those states with a low and fairly constant proportion of missing data: Bihar, Haryana, Himachal Pradesh, Kerala, Madhya Pradesh, Punjab, Rajasthan, Uttar Pradesh, West Bengal, and Assam. Jointly, these account for approximately 59% of the population of India, and can be used without caveat.

The next category is those states without usable information on dowry amount. This fortunately applies only to Maharashtra, where 74% of dowry values are missing. Given the time trends and after comparing with data from the 2008 wave of the survey, it is clear that surveyors in the 1999 wave in Maharashtra were unwilling or unable to elicit numerical dowry payments from respondents. As a result, when respondents made a dowry payment but did not say how much the payment was, this was coded as a missing value.

The third category, consisting of Karnataka has a different problem: when administering the survey, surveyors in Karnataka did not use a different code for zero dowry payments and non-response/missing data. This is the same coding as in the entirety of the 2008 wave of the REDS survey (see section B.2). For this data, we can estimate the proportion of truly missing information (i.e. uncertain whether was or was not payment of dowry) using the upper bound from other states, and then code the remaining as zero dowry payments.

Finally, the last category is those states with significant trends in missing data over time: Gujarat, Orissa and Tamil Nadu. In these states, we see that the proportion of missing values is initially low, but then rises sharply over time. Our interpretation of this data is similar to the data from Maharashtra, i.e. that surveyors successfully found that respondents had paid dowry, but

were unable to elicit the precise amount. This interpretation is supported by comparisons with the 2008 wave of the data from those states, which gives a similar trend over time to that seen in this data.

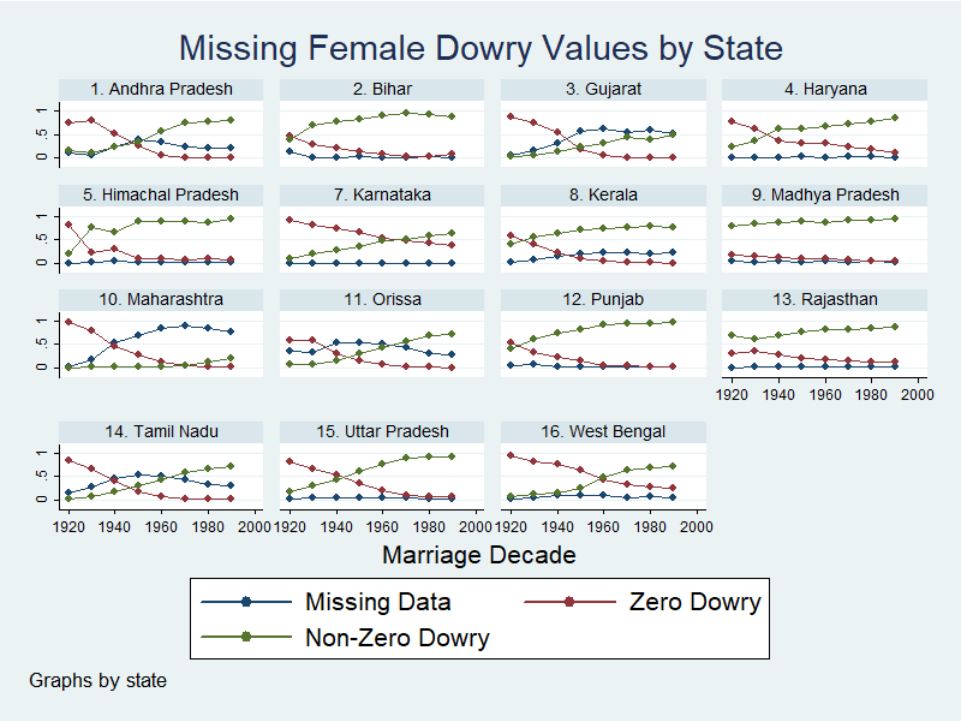

Figure A23: Missing Dowry Information by State

As a result, for the purposes of analysis on whether dowry was paid, we run a robustness check in which we: 1) code missing values in Maharashtra, Gujarat, Orissa, and Tamil Nadu as having paid dowry; and 2) code missing values in Karnataka as having not paid dowry. The right panel of table A24 plots the proportion of marriages with dowry with this change and finds that the qualitative patterns are the same as in the left-hand panel: a sharp increase in the post-1940 period, with a peak around 1970. Our tests of all of the theories also all go through after making these changes. For the analysis on size of dowry, we: 1) drop Maharashtra; 2) code missing values in Karnataka as a non-payment (payment of zero) for dowry; and 3) drop Gujarat, Orissa, and Tamil Nadu. We re-run all of the analysis with these changes, and again, our tests continue to hold.

One concern is that the pattern of missing data could affect some of our results on changes in dowry by caste over time, particularly in the Sanskritization analysis (e.g. individuals in some caste groups have worse memories). Below, we plot the rate of missing data by caste group over

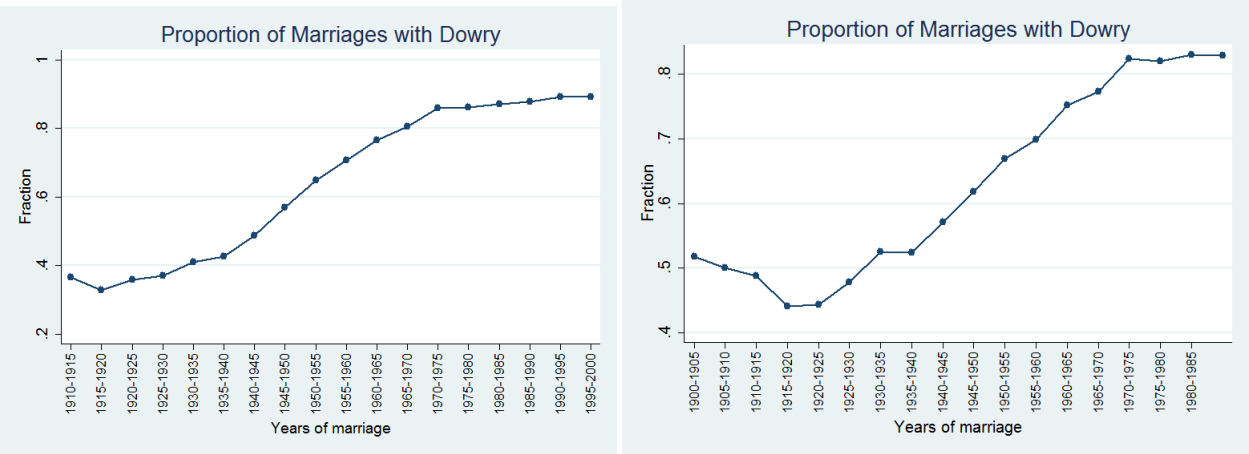

Figure A24: Comparison of Dowry Prevalence Graphs

time using states with consistently good data. Non-recall rates are low, and fairly consistent across caste. Note that this graph omits the states of Maharashtra and Tamil Nadu. This is because they have much higher prevalence of missing data, as seen in the previous data, as well as a much higher proportion of Brahmins. As a result, this makes it appear as if there is more missing data among Brahmins, when it is in fact a state effect.

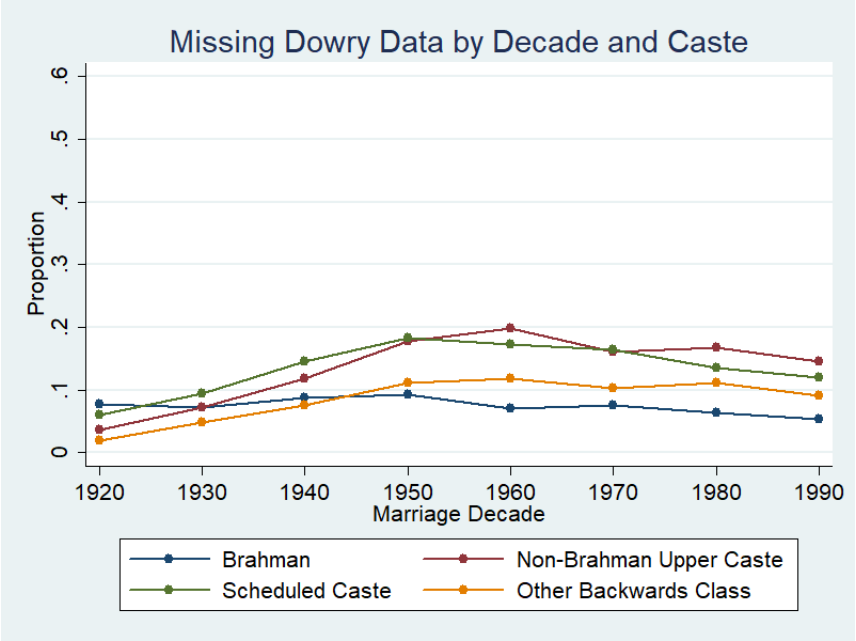

Figure A25: Missing Dowry Information by Caste

Another concern is that the pattern that we observe is also consistent with memory failure: it could be that the prevalence of dowry was roughly the same over time, but that people do not

remember dowry payments that occurred a long time ago. First, given the centrality of marriage and payment of dowry in rural India households, they are likely to intimately remember the details of past transactions. Second, when we plot data from the 2008 round of the survey in figure A26, the patterns are quite consistent, i.e. a rise in the prevalence of dowry beginning around 1930 and culminating in nearly universal dowry by the early 1970s. If recall bias were driving our results, we would expect for dowry payments to peak ten years later, around 1980, rather than in 1970.

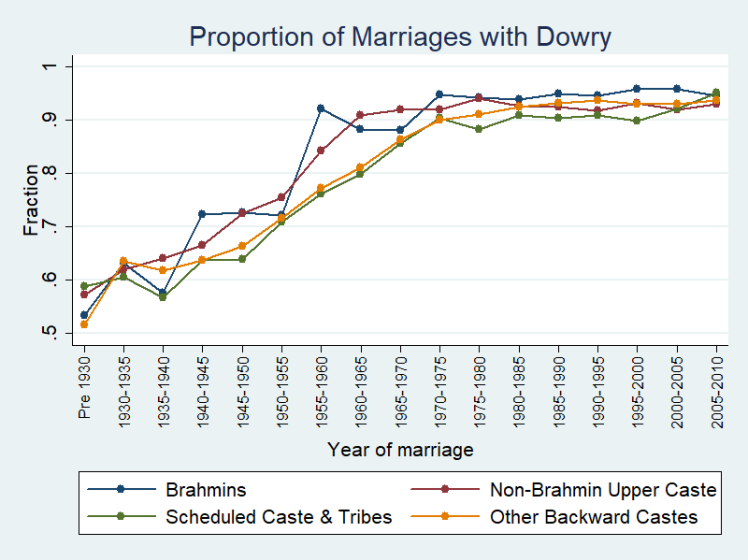

Figure A26: Marriages with Dowry/Gifts from the 2008 ARIS/REDS survey

B.4 Miscellaneous Data Notes

NSS Data In order to measure the average household income by state, we use data from the Indian National Sample Survey. We use information on the average rural agricultural wages by state for most years between 1956 and 1992, where the wage is denoted in contemporary prices. Since dowry in our data is also denoted in contemporary prices, we divide reported dowry payment by average rural agricultural wage multiplied by 300 (as the number of working days per year). This produces figure 6.

Price Deflation In order to properly plot the evolution of dowry over time, we need a standardized price index. We prefer to use the wholesale price index, since this measures inflation in prices of all goods and has a long time series. Other sources have used deflators based on the prices of gold

and rice. The price of gold fluctuates massively, and leads to unreasonable results when comparing before and after 1980, when there is a massive price spike. We are unable to find satisfactory rice price data for India for the pre-independence period, and so prefer the more consistent measure of WPI. This data is pieced together from a number of sources. For 1900-1950, we use wholesale price index data from pg 685 of Singh (1965). For the period 1939-1950, we cross-check this with data from other sources and find it to be consistent. For 1950 to 1970, the data is taken from the first section of Bhagwati and Srinivasan (1975), which comes from data from the Office of the Economic Adviser. For 1970-2013, we use data available online from the Reserve Bank of India.

C Sex ratio and Dowry

In this paper, we attempt to measure the sex ratio as experienced by a respondent when they participated in the marriage market. Ideally, we would measure the number of women between the ages of approximately 13-22 and men between the ages of 18-30 from the same jati and district as the participant when they are on the market. Unfortunately, there are no datasets that allow us to do this, so we must rely on approaches that approximate this sex ratio. Each of the three possible approaches has strengths and weaknesses, so to check the robustness of our results, we apply all three and check for consistency. By demonstrating that the results hold over each, we increase the probability that no one flaw drives our findings. We discuss each method and its advantages, with a summary in table A4.

Method 1: Historical Census Tables (1911-1991) The first method is to use district-level data from the Census of India rounds between 1911 and 1991.[36] Each census round has tables giving the current population of men and women in five-year age bins in each district at the time of each census (e.g. males and females aged 0-5, 5-10, 10-15, etc. in West Godavari district). While the pre-independence data has this additionally broken down by religion, this is no longer true after independence. In the pre-independence data, we only have data for 69/100 REDS districts. This is because some of the REDS districts were in princely states, for which we do not have colonial census data. Dropping the pre-independence period has no effect on our results. With this data,

[36]Excepting the 1941 census, which is known to be of questionable data quality, primarily due to British attention being diverted to World War II

we construct the sex ratio for marriages in a given five period and district as the number of women aged 10-20 divided by the number of men aged 20-30 at the start of that period.

Method 2: Census Age Tables (1991 Census) Our second method uses data from the 1991 round of the Indian Census. In this data, we observe the full age distribution of individuals by gender at the district level. For example, in West Godavari district of Andhra Pradesh, we observe that living in the district, there are a total of 37120 men and 36260 women who were born in 1975, 40850 men and 36240 women born in 1976, 35474 men and 35170 women born in 1977, etc. This data does not identify population by caste or religion. It is based on a 10% sample of individuals in major states and an 100% sample in smaller states for a total of 108.4 million records, or approximately 250,000 individuals per district. While technically not a complete census, it is a sufficiently large sample to get extremely precise estimates of the population by age group. With this data, we construct the sex ratio for marriages in a given year and district as the number of women aged 13-20 divided by the number of men aged 18-25 in that year.

Method 3: Pooling Nationally Representative Surveys The third approach takes multiple rounds of the National Sample Survey (rounds 38 [1983], 43 [1987-88], 50 [1993-4], 55 [1999-2000], 62 [2005-6]) and combines them to estimate the total population of men and women born in a given year. In this dataset, we observe the current state, district, caste, gender, religion, education, and birth year of a given individual. We reweight according to the associated weight file to produce population-valid values. Even though pooling the NSS data produces approximately 1.5 million observations, we can only generate estimates of sex ratio at the identity group-state level: doing so at a lower level of aggregation such as district would not have enough observations.

We divide the data into 11 religion-caste based identity groupings : Hindu-ST, Hindu-SC, Hindu-Other, Muslim-SC, Muslim-Others, Sikh-SC, Sikh-Others, Christian-ST, Christian-Others, and Other Religions. The Others caste grouping aggregates OBC and General caste individuals, since we don't observe OBC status for early rounds of the NSS. We then take sex ratio estimates for each identity groups at the state-level for any cell in which we have at least 5000 total observations. This ensures that there are enough observations to calculate an accurate measure of sex ratio for that state-identity group level in fine increments of year. After doing this, 88% of the identity

grouping cells in which we observe sex ratio are either Hindu (any caste) or Muslim-Others. Such a rule allows us to observe identity groupings that may be common in some states (e.g. Christian-STs in the state of Jharkhand), but not others.With this data, we construct the sex ratio for marriages in a given year, state, and identity grouping as the number of women aged 13-20 divided by the number of men aged 18-25 in that year, identity grouping and state.

C.1 Discussion of Methods

In an ideal world, we would use data to construct sex ratio that satisfies six criterion: 1) non-retrospective data; 2) full census; 3) jati-level; 4) if not jati-level, then caste-level; 5) district-level; and 6) age broken down at the yearly level (rather than in 5-year bins). Since no existing data set satisfies each of these criterion, we discuss each and then how each of the approaches does or does not satisfy them.

(1) Non-Retrospective Data: When measuring the population of men and women who were part of the same marriage market as a given individual, we use those within a certain age of the individual. It is strongly preferable to measure this population at the time close to, or even prior to, the age of marriage, rather than based on the population of men and women within that age band at a much later date (e.g. measuring the population of 40-50 year old men and 35-45 year old women in 1980 and inferring sex ratio in the marriage market in 1960 based on that). The first reason for this is differential mortality by gender between the time of being on the marriage market and measurement. In particular, women in India have high rates of maternal mortality in childbirth, and so the number of 35-45 year old women in 1980 might be much reduced from the population of 15-25 year old women in 1960. If there are some regions where maternal mortality during childbirth is particularly severe, it may appear that sex ratios at age of marriage are relatively skewed, when it is in fact due to a confounding factor.

The second reason is that introduces additional noise into the measurement. It is quite common for individuals in India, particularly in rural areas, to not know the exact year in which they were born. For those under the age of 20, this presents less of a challenge, since it is straightforward to figure out approximate age (e.g. spacing between siblings and age of most recent birth). For older individuals, this is more of a challenge, leading to clustering of ages. For example, it is much more common to observe individuals stating that their age is a quantity ending in 0 or 5, e.g. 40,

45, 50, etc). By averaging over a range of years (since we consider the marriage market to be the range of years over which an individual of a given gender tends to marry), we partially mitigate the problem, but it is clearly better to use more accurate measurements from younger individuals.

The third reason is that individuals may migrate between the time of marriage and later measurement. Fortuantely, long-term migration in India is relatively low relatively to temporary migration, and individuals often migrate within a district or state (such as to a larger metropolitan entity). Nonetheless, if this migration is differential by gender across areas (e.g. one district sends many male economic migrants to Delhi, while another does not), this can skew the measurement of population and sex ratio.

	Historical Census	Census Age Tables (1991)	Pooled National Surveys
Non-Retrospective	Yes	No	No
Large Sample/Census	Yes	Yes	No
Jati-level Data	No	No	No
Caste Grouping-level Data	No	No	Yes
District-level Data	Yes	Yes	No
Yearly Data	No	Yes	Yes

Table A4: Advantages and Disadvantages of Population Datasets

(2) Full Census: In one of our approaches, we combine nationally representative surveys to estimate sex ratios. These are relatively large samples (in the hundreds of thousands of individuals), but still are derived from less than 1% of the population. On the other hand, the Census of India measures all individuals within a given geographical boundary and age. By using a sample rather than the full population, this can introduce rather severe noise into the measurement of sex ratios, particularly given that sex ratios vary within a relatively small range. Conventional standard errors will not properly correct for this, since the estimated sex ratio is treated as being measured without noise.

(3) Jati or Caste-Level Data: As shown in the paper, marriage markets in India are at the jati-level: individuals rarely marry outside their jati and face significant social sanction, such as out-casteing, for doing so. Thus it would be ideal to measure the sex ratio at the jati level. Unfortunately, that is not possible in any of these datasets. The Indian Census no longer collects information at the jati level, or anything more detailed than whether an individual is from a

scheduled caste or tribe. While there are some nationally representative surveys with information on jati (e.g. National Family Health Surveys), their samples are too small for reliable inference on any jati-level statistics. Nationally representative surveys such as the National Sample Survey collect information on the broad caste group to which an individual belongs (general caste, other backwards class, scheduled caste, scheduled tribe) and religious affiliation, but that combines a large number of heterogenous jati groups. Nonetheless, this is likely to be more strongly correlated with jati-level sex ratios than the overall population.

(4) Yearly Data One possible advantage for a dataset is that it contains the number of individuals born in a given year, rather than in broad five year bins. This makes it more straightforward to determine the total number of men and women who would have been on the marriage market at the same time as a given individual.

(5) District-level data: As we show in the paper, marriage markets are concentrated within the district: 78.3% of marriages are within the district, and the average distance between the household of brides and grooms is less than 15 miles. It is thus advantageous to measure population at the district-level, rather than state or higher, since there is substantial variation in sex ratio across districts within a state.

C.2 Method 2: 1991 Census Age Tables

The tables for the first method can be found in the text of the paper. This table is the results from running the sex ratio regressions using sex ratio as defined by the second method, using district-level age tables from the 1991 census. We again do not find a relationship between sex ratio and either the intensive or extensive margin of dowry payment. For this, we restrict the sample to marriages from 1950 onwards (the oldest age group in this definition would be 66 years of age in 1991, at which point mortality starts to increase), but our findings are robust to restricting the sample to marriages from 1940 onwards, 1960 onwards, 1970 onwards, and 1980 onwards.

C.3 Method 3: Pooled National Surveys

Table A6 reruns the sex ratio regressions, but with sex ratio constructed from the pooled National Sample Survey. As in table A5, we restrict the sample to marriages from 1950 onwards.

	(1) Dowry Dummy	(2) Dowry Dummy	(3) Dowry Dummy	(4) Ln(Dowry)	(5) Ln(Dowry)	(6) Ln(Dowry)
sexratio	0.0174	-0.0154	-0.0000292	-0.471	-0.539	0.101
	(0.0439)	(0.0415)	(0.0366)	(0.358)	(0.329)	(0.284)
Observations	59589	59496	39610	40551	40263	21167
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Village FE	Yes	No	No	Yes	No	No
Household FE	No	Yes	No	No	Yes	No
Household-5 year FE	No	No	Yes	No	No	Yes

Standard errors clustered at the district level

Table A5: Sex Ratio Regressions (1991 Census Age Tables)

	(1) Dowry Dummy	(2) Dowry Dummy	(3) Dowry Dummy	(4) Ln(Dowry)	(5) Ln(Dowry)	(6) Ln(Dowry)
Sex Ratio	-0.0657	0.0213	0.0143	-0.793	0.251	0.145
	(0.0676)	(0.0578)	(0.0462)	(0.465)	(0.445)	(0.472)
Observations	51785	51701	34304	27876	27723	18276
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Village FE	Yes	No	No	Yes	No	No
Household FE	No	Yes	No	No	Yes	No
Household-5 year FE	No	No	Yes	No	No	Yes

Standard errors clustered at the district level

Table A6: Sex Ratio Regressions (Pooled National Sample Survey)

D Marriage Market Facts

D.1 Assortativity of Matching on Education

We follow two methods to examine how the assortativity of education has been changing over time. Section D.1.1 discusses the method proposed by Eika et al. (2014) while section D.1.2 discusses the method proposed by Greenwood et al. (2014). Both the methods suggest an increase in the assortativity of matching on education over time.

D.1.1 Eika et al. (2014)

We first discuss the method suggested by Eika et al. (2014). Intuitively, their method checks whether matches between partners of the same educational level occur more or less frequently than would be expected by chance. If they occur more (less) frequently, then matches are positive (negative) assortative. They divide men and women into relatively coarse educational categories and examine the proportion of marriages in which the bride and groom have the same level of education. They then divide this by the expected proportion of marriages in which the bride and groom would

have the same level of education if matches had been done randomly. For each education bin, an assortativity coefficient r_{ii} can be calculated, where the numerator is the proportion of marriages where the bride and groom both are of education level i, and the denominator is expected proportion of marriages in which the bride and groom would have education level i if matching were done randomly.

In our analysis, individuals are divided into education groups of no education, only primary education, middle-school education (4-8 years of school) and secondary education or higher. Figure A27 plots these ratios (r_{ii}), demonstrating that the patterns of assortativity have changed substantially over time. While those with a middle school education or higher have begun to match less assortatively over time, those with no education have become more likely to marry others with no education. Those with primary education initially matched less assortatively, but have begun to match more assortatively since the late 1970s. We also calculate an aggregate assortativity ratio. This is equal to the total number of marriages in which the bride and groom have the same level of education divided by the expected number of marriages in which bride and groom education would be equal if matching were done randomly. Driven by the increased assortativity of matching by those with lower levels of education, we find that the coefficient of assortative matching on education nearly doubles between 1940 and 2000.

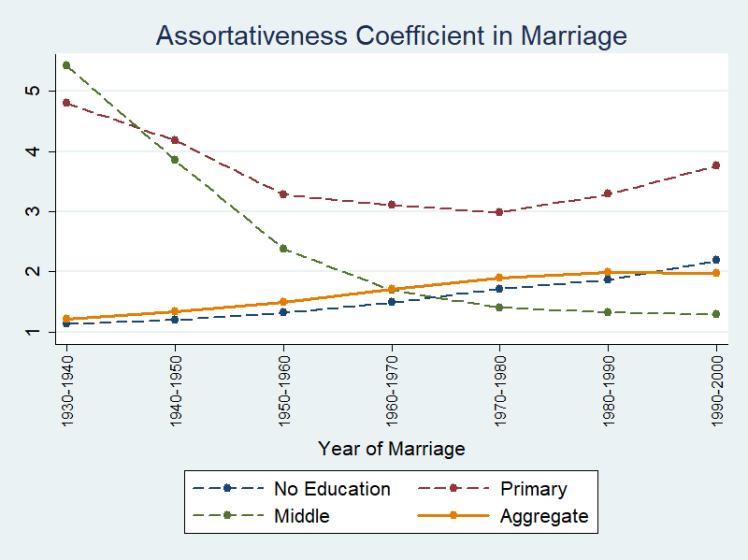

Figure A27: Assortativity in marriages

D.1.2 Greenwood et al. (2014)

In this section, we calculate assortativity using the method of Greenwood et al. (2014). To do this, for each marriage m in decade t, a regression of the following form is used to quantify the degree of assortativity across time:

$$Edu^f_{mt} = \alpha + \beta Edu^m_{mt} + \sum_{k \in \Gamma} \gamma_k \text{ x } Edu^m_{mk} \text{ x } D_k + \sum_{k \in \Gamma} \theta_k D_k + \epsilon_{mt} \tag{6}$$

where $t = \{1930, 1940, \ldots, 1990\}$; D_k is a dummy variable that takes a value of 1 when $k = t$ and 0 otherwise. We estimate the regression for both when education is measured in years as well as when education is categorized into no education, primary education, high school and college (as in figure A12 above). Figure D.1.2 plots the γ_t coefficients. As we can see from the coefficient plot, the positive assortativity of matching across education groups is increasing over time.

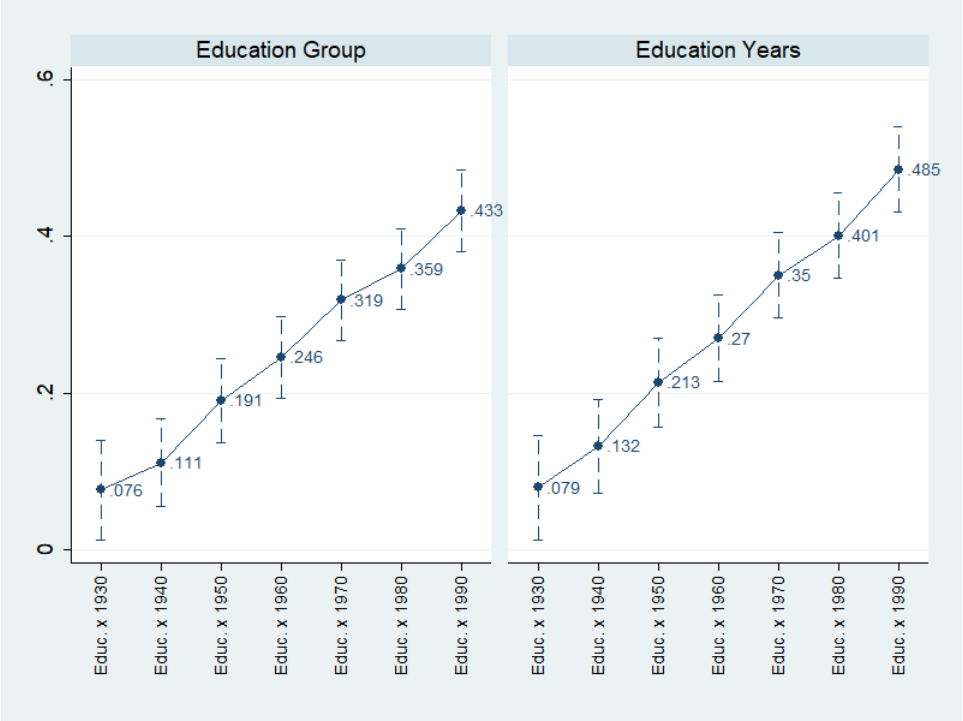

Figure A28: Education assortativity coefficient plot

Table A7: Education Assortativeness

	(1) Female Educ. (Group)	(2) Female Educ. (Years)
edgroup_m	0.139***	
	(0.0215)	
Educ. x 1930	0.0729***	0.0726***
	(0.0270)	(0.0280)
Educ. x 1940	0.0907***	0.108***
	(0.0240)	(0.0250)
Educ. x 1950	0.163***	0.180***
	(0.0228)	(0.0238)
Educ. x 1960	0.226***	0.245***
	(0.0222)	(0.0233)
Educ. x 1970	0.291***	0.315***
	(0.0220)	(0.0230)
Educ. x 1980	0.326***	0.364***
	(0.0219)	(0.0230)
Educ. x 1990	0.392***	0.444***
	(0.0222)	(0.0231)
Male Education		0.119***
		(0.0227)
Observations	66874	66874
R^2	0.535	0.573
Year FE	Yes	Yes
District FE	Yes	Yes
Jati FE	Yes	Yes