Caste and the Indian Economy †

KAIVAN MUNSHI*

Caste plays a role at every stage of an Indian's economic life, in school, university, the labor market, and into old age. The influence of caste extends beyond private economic activity into the public sphere, where caste politics determines access to public resources. The aggregate evidence indicates that there has been convergence in education, occupations, income, and access to public resources across caste groups in the decades after independence. Some of this convergence is likely due to affirmative action, but caste-based networks could also have played an equalizing role by exploiting the opportunities that became available in a globalizing economy. Ethnic networks were once active in many advanced economies but ceased to be salient once markets developed. With economic development, it is possible that caste networks will cease to be salient in India. The affirmative action programs may also be rolled back and (statistical) discrimination in urban labor markets may come to an end if and when there is convergence across caste groups. In the interim period, however, it is important to understand the positive and negative consequences of caste involvement across a variety of spheres in the Indian economy. (JEL G22, J15, J71, O15, O17, Z13)

1. Introduction

Caste plays a role at every stage of an Indian's economic life. His caste will determine the type of school he has access to, the way he is treated by his teachers, and his interactions with his classmates. In young adulthood, his caste will determine whether or not he benefits from affirmative action in higher education and (later) access to government jobs. Over the course of his working life, his caste will determine how he is assessed by potential employers, while, in parallel, networks organized around his caste help him find new jobs. Overlapping caste-based networks will provide him with credit, help him start a business if he is endowed with entrepreneurial talent, and provide him with insurance against income shocks and major contingencies into old age.

The influence of the caste extends beyond private economic activity. Electoral politics

in postindependence India has had a strong caste component, with political parties targeting public resources to particular castes in return for their votes. Caste politics extends down to the local level. While it is well known that local politicians target their own castes, new research shows, in addition, that cooperation within the caste can lead to an increased supply of non-excludable public goods.

The importance of caste in India's economy and polity should not come as a surprise, given the structure of its society. Hindu society is stratified into four hierarchical classes, or varnas, with a large subpopulation of untouchables excluded from this system entirely. Within each of these classes, and among the untouchables, are thousands of castes or *jatis*. The central rule in Hindu society is that individuals must marry within their own caste. Recent genetic evidence indicates that this rule has been strictly followed for over two thousand years. Each caste consists of approximately 250,000 members spread over a wide area covering many villages (and, sometimes, urban locations). Spatial segregation within the village results in a high degree of local social connectedness, with caste clusters in distant villages linked to each other through ties of marriage over many generations.

The unusual social architecture described above has given rise to economic networks organized around specific castes (*jatis*) that are exceptional in their size and scope. While these networks smoothed the consumption of their members in the face of income fluctuations for centuries, they later expanded the domain of their activity to the urban labor market and into business when new opportunities became available under colonial rule. Caste networks have continued to play an important role in the Indian economy after independence in 1947. This period witnessed, in parallel, the establishment of the world's most aggressive affirmative

action program, whose stated objective was to eliminate the historical inequities of the caste system. Internal cooperation within castes supports economic networks, which compete with each other in the market without regard to their position in the social hierarchy. The much more visible affirmative action program, in contrast, was explicitly targeted at large groups of historically disadvantaged castes, giving them preferred access to higher education, government jobs, and political representation. These two forces have worked in tandem to shape the evolution of the modern Indian economy.

The aggregate evidence indicates that there has been convergence in education, occupations, income, and access to public resources across caste groups in the decades after independence. Some of this convergence is likely due to affirmative action, but caste-based networks could also have played an equalizing role by exploiting the opportunities that became available in a globalizing economy. It is possible that the caste-based affirmative action programs will no longer be needed at some point in the future. The available evidence indicates that caste discrimination, at least in urban labor markets, is statistical, that is, based on differences in socioeconomic characteristics between upper and lower castes. This type of discrimination will disappear if and when there is convergence between these broad caste groups. Caste networks and caste politics will also lose their relevance once the market economy and the democratic political system begin to function efficiently. In the interim period, however, it is important to understand the positive and negative consequences of caste involvement across a variety of spheres in the Indian economy.

2.   *Caste in Indian Society*

The caste system is a system of social stratification in India that goes back as far

as 1500–500 BCE. The caste system comprises four hierarchical classes, or varnas, the Brahmins, Kshatriyas, Vaisyas, and Shudras. Certain population groups, known today as *Dalits*, were historically excluded from the varna system and were regarded as untouchables. Within each varna, and among the *Dalits*, are hundreds of castes or *jatis* (I will use these terms interchangeably in the discussion that follows).

The view among scholars of the British colonial era, as articulated by Dumont (1970), was that the hierarchical aspect of the caste system, characterized by the varnas, was largely symbolic. Upper castes maintained higher levels of ritual purity than lower castes, for example, by being vegetarian, but these differences had no direct economic consequences. Other research has contested this view, arguing that social status allowed the upper castes to economically exploit the lower castes (Gupta 2000, Dirks 2001). In recent decades, the affirmative action program put in place by the Indian government has given the varna system a new role in the Indian economy. This affirmative action program reserves positions in institutions of higher education, the central government, and local politics for individuals belonging to historically disadvantaged castes. These castes comprise the former untouchable castes, which are collectively classified as scheduled castes (SCs), socially and economically marginalized indigenous ethnic groups that are classified as scheduled tribes (STs), and, more recently, another group of castes that are referred to as other backward castes (OBCs).

When evaluating the affirmative action program, it is appropriate to compare broad caste categories, that is, SC, ST, OBC versus the other (forward) castes. This may also be appropriate when studying discrimination in urban labor markets, where employers may care about the socioeconomic background of job candidates, which is proxied by caste

category rather than their specific caste. However, any analysis of economic networks in the Indian economy must be based on the caste or *jati* and not on the broad caste group (SC, ST, OBC) or the village. As discussed below, a number of recent studies on networks in India have ignored this critical requirement.

Unlike the varnas, the castes were explicitly established to serve an economic function. Gupta (2000) posits that they were initially occupational guilds, tracing their emergence to the feudal period (seventh to twelfth century). Other scholars have claimed that *jati* identity and the rigid boundaries between castes that exist today were a creation of the British colonial government (de Zwart 2000, Dirks 2001). Although the British were not the first to enumerate the Indian population by caste—this goes back to Mughal rule (Appadurai 1993)—it is true that they assigned government jobs to particular castes, possibly creating a new *jati* identity or reinforcing an existing identity (Bayly 2001). Recent genetic evidence, however, places the origins of the *jati* system much further back than was previously assumed. This research indicates that marriage within castes, the lynchpin of the caste system, was rigidly adhered to long before the feudal period. While there was free mixing in the Indian population in ancient times, there was an abrupt change 1,900–4,200 years ago, after which strict endogamy within narrowly defined population groups became the norm (Moorjani et al. 2013).[1] The Indian population today consists of 4,000 distinct genetic groups, each of which is a caste (or its non-Hindu equivalent kinship group),

[1] Bidner and Eswaran (2015) propose an economic explanation for the emergence of the caste system that is based on the complementarity between a woman's productivity and her husband's skill. Their theory can explain many features of the caste system, including its hereditary nature, the key rule of marital endogamy, and its hierarchical structure.

making India one of the most genetically diverse countries in the world (Singh 1992).[2] The number of genetic groups that have been identified is in line with the number of castes (approximately 3,000) listed in the colonial censuses.[3] Thus, while colonial policies may have enhanced caste affiliations, they did not create them.

With 4,000 castes and a national population of one billion, each caste comprises 250,000 members on average. Although castes are spread over a wide area, sometimes an entire state, a single caste will not have a presence in all villages in the area that it covers. It will cluster, instead, in select villages that are often located far from each other. This clustering shows up clearly in the most recent round of the Rural Economic and Demographic Survey (REDS) conducted in 2006. The REDS is a nationally representative survey of rural Indian households that has been administered by the National Council of Applied Economic Research at multiple points in time over the past four decades. Each survey round includes a census of all villages that are covered, followed by a detailed survey of a sample of households in each village. The census data from the 2006 REDS indicate that the mean number of castes per state is sixty-four, while the mean number of castes per village is twelve. With 340 households on average in a village, this implies that a caste will have about thirty households in the select villages where it locates. This is a large enough number to support a local caste community within the village, while marriage ties link the caste across villages. Marriage in India is endogamous (within caste) and patrilocal, that is, the bride moves into her husband's household. Most women leave their natal village when they marry. A dense web of marriage ties thus forms over many generations, allowing caste networks to span a wide area while maintaining their connectedness.

Although the population census of India has not collected information on caste affiliation since 1931, recent rounds of nationally representative surveys such as the 1999 REDS and the 2005 India Human Development Survey (IHDS) report that over 95 percent of Indians continue to marry within their caste (or equivalent kinship community). This is not surprising, given the genetic evidence that has recently become available, and the same patterns of strict endogamy have been reported by respondents in other sample surveys conducted in rural and urban India (Munshi and Rosenzweig 2006, Munshi 2011, Luke and Munshi 2011). Figure 1 uses retrospective information from the 1999 REDS to report rates of outmarriage, that is, marriage between members of different castes, in rural India for the children and siblings of household heads over the 1950–99 period. Outmarriage is just around 5 percent of all marriages—this number appears with remarkable consistency across all the studies listed above—and has remained stable over time.[4]

In addition to marital endogamy, the distinctive feature of the caste system is the

---

[2] Muslims marry within *biradaris* and Christians continue to marry within their original (pre-conversion) *jatis*. Other religious communities such as the Parsis and the Jains also practice strict endogamy.

[3] The first colonial census to enumerate the population by caste was in 1871 and the last was in 1931. The Indian government has not collected this information since independence in 1947. The most recent census round in 2011 did collect this information, but it has not been made publicly available.

[4] It has been claimed in previous research, for example Cassan (2015), that caste identity can be manipulated when there is an economic incentive to report an affiliation with another *jati*. While this type of manipulation may show up in government documents and census statistics, it does not imply that caste identities have actually changed. The genetic evidence and the remarkably stable patterns of marriage obtained across many surveys suggest that caste boundaries have remained rigid for a very long period of time. Indeed, Christians continue to marry within their original *jatis* in rural India, despite having changed their religion many generations ago (Luke and Munshi 2011).
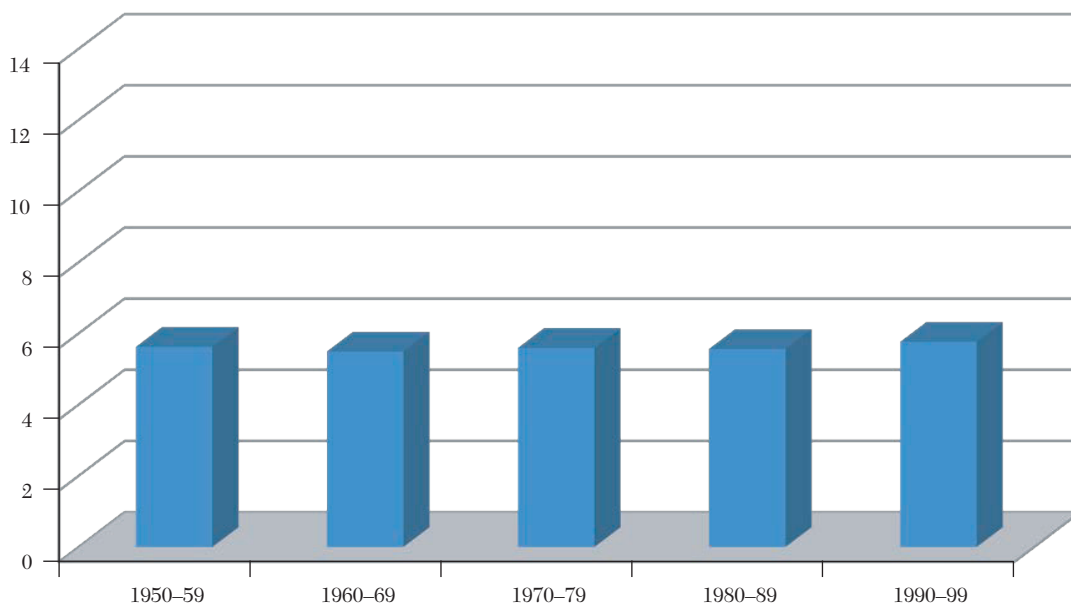
Figure 1. Marriage outside the Caste in Rural India

Note: Percentage of marriages by women in the household head's family (sisters and daughters) outside their caste.

Source: REDS 1999.

restrictions it places on social intercourse between castes, not just between upper castes and lower castes, but also between castes occupying the same position in the social hierarchy (Ghurye 1932). These restrictions result in spatial segregation, with members of a caste typically living together in a particular area of the village; indeed, streets in Indian villages are often referred to by their caste names. Munshi and Rosenzweig (2015) use data from the 2006 REDS to describe the spatial distribution of castes within villages, as well as social and economic interactions in the village. Focusing on the 91 percent of REDS villages for which information on local government (panchayat) wards is available, each caste makes up, on average, 6 percent of the population of a village. Within the ward, which

is a smaller spatial unit, the average caste's share increases to 14 percent, indicative of the spatial clustering along caste lines that characterizes the Indian village.

Given the strong social ties within castes and the weak social interactions between castes, we would expect informal economic transactions in the village to be concentrated within the caste. The 2006 REDS village census asked each household head to list the names of three individuals in the village that he would approach for food, in the event of a temporary shortfall, and for a short-term loan. If these individuals were approached without regard to caste affiliation, we would expect, on average, that 6 percent of the individuals providing food transfers and loans would belong to the respondent's caste. Based on the identity of the first listed

individual, the REDS data indicate that as much as 50 percent of food transfers and loans are from individuals belonging to the same caste. The REDS data also tell us that more than half of caste loans (by value) originate outside the village. If we included those loans, the fraction of individuals within the caste that would be approached for a loan would increase well above 50 percent. Informal transactions are evidently concentrated within the caste in the Indian village.

The insider–outsider dichotomy, characterized by close intragroup relations and weak (or even adversarial) intergroup relations, is a feature of all traditional societies.[5] This dichotomy is especially pronounced in Indian society due to its unique caste-based structure, as discussed above, and is also evident from cross-country comparisons based on the World Values Survey (WVS), documented by Munshi and Rosenzweig (2015). Figure 2 reports cross-country results from wave five of the WVS (conducted between 2005–09 in different countries) on questions relating to trust and tolerance of outsiders. Restricting the sample to countries with a population in excess of twenty million that are classified by the World Bank as low, lower-middle, or upper-middle income (with the addition of Russia), we see in figure 2, panel A, that India ranks close to the top of the list with regard to trust in neighbors. Given the spatial segregation that characterizes the Indian village, the Indian respondents are essentially reporting that they have a high degree of trust in their fellow caste members living nearby. However, the ranking reverses when we look at tolerance of outsiders—speaking a different language or following a different religion—in figure 2, panels B and C, respectively. India now lies at the very bottom of the list and, based on

the preceding discussion, we would expect Indians to be similarly intolerant of neighbors belonging to a different caste (this is certainly reflected in how they live).

The sections that follow will examine the role played by caste communities in a wide variety of economic activities. The insider–outsider dichotomy will be a common thread running through the discussion in each section. While internal cooperation allows castes to support the economic activity of their members when market institutions are missing or function imperfectly, the disregard (and even active hostility) toward outsiders creates new distortions in the economy. There is no perfect substitute for a well-functioning market, and we will see that while informal caste-based networks may increase economic efficiency in the short run, they can restrict the occupational and spatial mobility of their members in the long run.

3. *Mutual Insurance*

3.1 *Caste-Based Insurance*

Consider an economy with two risk-averse individuals, A and B. Each individual earns either 1,000 or 2,000 per period with equal probability. Income realizations are independent across individuals, resulting in the following income states: (1,000, 1,000), (2,000, 2,000), (1,000, 2,000), (2,000, 1,000), where the first term in parentheses refers to individual A's income and the second term refers to individual B's income and each income state occurs with probability 0.25.

Given that individuals are risk averse, they will benefit from a mutual insurance arrangement that smooths their consumption over time. Individuals who receive a positive income shock make transfers to temporarily less fortunate individuals in these arrangements, with the roles reversing in later periods when the income shocks

[5] Akerlof (1976) and Choy (2018) provide game theoretic foundations for why internal cooperation coupled with the exclusion of outsiders can be sustained in equilibrium.

Panel A. Trust in neighbors
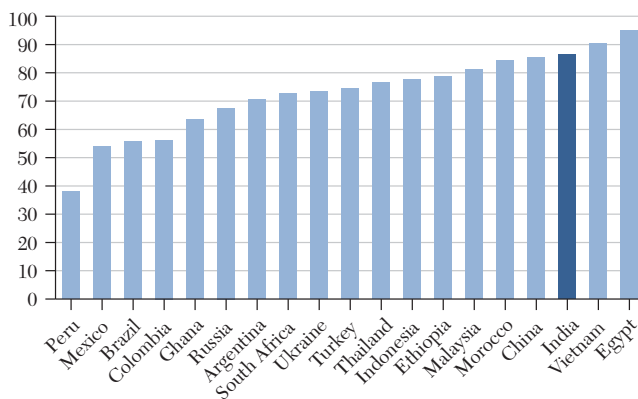
Panel B. Tolerance of neighbors speaking a different language
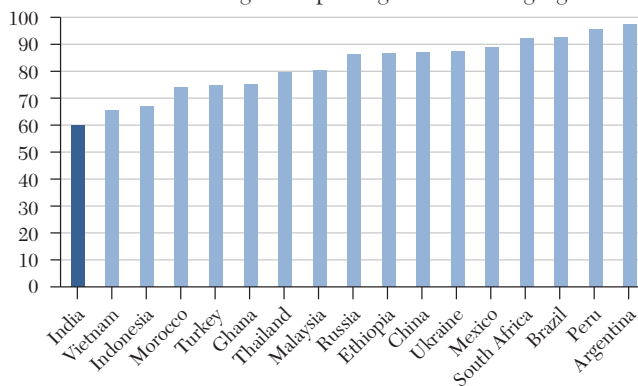
Panel C. Tolerance of neighbors following a different religion

Figure 2

Notes: Panel A reports the percentage of respondents who trust people in their neighborhood. Panel B reports the percentage of respondents who do not have a problem with their neighbors speaking a different language. Panel C reports the percentage of respondents who do not have a problem with their neighbors following a different religion.

Source: WVS, wave five.

are reversed, leaving everyone better off. An alternative (equivalent) way to characterize mutual insurance arrangements when preferences are logarithmic is that the 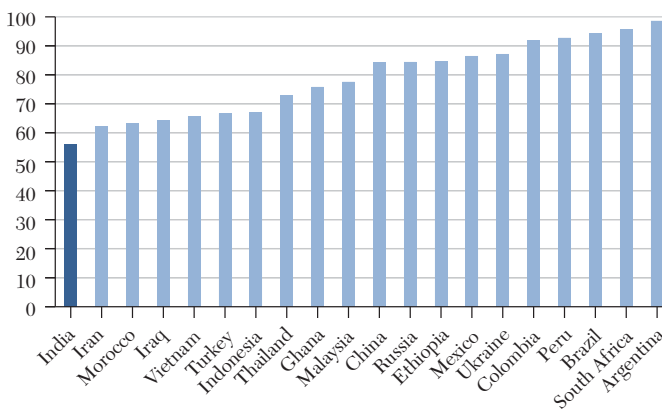total income of all participants is pooled in each state or period and then distributed on the basis of a prespecified income-sharing rule,[6]

$$(1) \qquad c_i^s = \frac{\lambda_i}{\sum_i \lambda_i} Y^s,$$

where c_i^s is the consumption of individual i in income state s, Y^s is total income (over all participants in the insurance arrangement) in state s, and $\lambda_i/\sum_i \lambda_i$ is individual i's share of the pot, which is constant across all states. In our simple example, the income-generating process is the same for the two individuals, and so it is reasonable to assume an equal income-sharing rule: $\lambda_A = \lambda_B = 1/2$. The usual assumption in models of mutual insurance is that there is no storage or savings, which results in three consumption states: $(1{,}000, 1{,}000)$ with probability 0.25, $(2{,}000, 2{,}000)$ with probability 0.25, and $(1{,}500, 1{,}500)$ with probability 0.5.

The preceding discussion serves as the basis for a test of full risk sharing. Taking logs in equation (1),

$$(2) \quad \log(c_i^s) = \log(Y^s) + \log\!\left(\frac{\lambda_i}{\sum_i \lambda_i}\right).$$

The preceding equation can be used to derive two testable predictions:

(i) Replacing states, s, with time periods, t, individual consumption will co-move with total income (or total consumption)

in the group around which mutual insurance is organized.

(ii) Conditional on total income, Y^s, notice that individual income, y_i^s, has no bearing on consumption, c_i^s. This is evident from our simple example; consumption is equal to 1,500 in states with unequal incomes, regardless of whether the individual earns 1,000 or 2,000.

In our example, income shares are equal. More generally, individuals earning higher incomes on average will have higher income shares. This is accounted for in the test of full risk sharing by including individual fixed effects, which capture the last term in equation (2). Conditional on group income in a given period, shocks to individual income, that is, deviations from average income, should have no effect on individual consumption with full risk sharing.

Starting with Townsend's (1994) seminal contribution set in village India, studies from across the developing world have documented extremely high levels of risk sharing (close to full risk sharing) in rural communities, for example Grimard (1997); Fafchamps and Lund (2003); and Angelucci, De Giorgi, and Rasul (2015). These communities are evidently able to smooth income risk extremely effectively, despite the fact that private credit and government safety nets are largely absent in the economies within which they operate. While this research has shed new light on the important role played by informal nonmarket institutions in smoothing income risk in developing economies, it suffers from two limitations.

The first limitation is that the social group within which risk sharing is assumed to take place is often chosen on the basis of convenience. For example, Townsend (1994) treats the village as the relevant social group, whereas our characterization of Indian society would suggest that the caste (within and across villages) is the social group around which

---

[6] More generally, with full risk sharing, the ratio of marginal utilities between any two members of the arrangement is constant in all income states. With log preferences, this implies that the ratio of consumptions are constant. Summing over all network partners, this implies that each individual receives a constant fraction of the total income that is generated in each state.

mutual insurance arrangements are organized in rural India. The available evidence supports the view that the caste or *jati*, rather than the village, is the relevant group around which mutual insurance arrangements are organized in rural India. Caldwell, Reddy, and Caldwell (1986) surveyed nine villages in South India after a two-year drought and found that nearly half (46 percent) of the sampled households had taken consumption loans during the drought. The sources of these loans (by value) were government banks (18 percent); moneylenders, landlords, or employers (28 percent); and relatives and members of the same caste community (54 percent), emphasizing the importance of caste loans for smoothing consumption.

The REDS, which provides detailed information on gifts and loans for a sample of villages covering all the major Indian states, can be used to corroborate these findings. Focusing on the 1982 REDS round, which included a detailed classification of the source and purpose for each loan, the key findings reported in Munshi and Rosenzweig (2016) are that (i) caste loans are second in overall value after bank loans, and (ii) caste loans are the primary source of support for meeting major contingencies such as illness and marriage. Gifts make up the bulk of transfers within the caste (67 percent by value). If we included gifts, the caste would be by far the most important source of support for smoothing consumption and meeting major contingencies in rural India. As noted by Munshi and Rosenzweig, much of this support comes from caste members outside the village. The same pattern holds up in later REDS rounds, such as the 1999 round, and in other nationally representative surveys such as the 2005 IHDS. Providing the most direct evidence that the caste, rather than village, is the relevant risk-sharing unit, Mazzocco and Saini (2012) use International Crops Research Institute for the Semi-Arid Tropics (ICRISAT) data

to establish full risk sharing at the level of the caste but not the village.

Recent research continues to ignore the central role played by the caste in smoothing consumption in rural India (among the exceptions are Mobarak and Rosenzweig 2012 and Munshi and Rosenzweig 2016). Morten (2019) treats the village as the risk-sharing unit, completely ignoring caste, just as Townsend did. Ambrus, Chandrasekhar, and Elliot (2014) partition the village into two broad caste groups: scheduled caste/scheduled tribe (SC/ST) and general merit/other backward caste (GM/OBC). As discussed previously, the SC, ST, OBC categories group castes with similar social status for the purpose of affirmative action programs implemented by the Indian government. *Jatis* within a broad caste group do not share a common history nor do they share a special social bond. The test of full risk sharing can be implemented with data from any subset of individuals belonging to the relevant social group. However, it will erroneously reject full risk sharing if it includes individuals outside the group, such as individuals outside the caste residing in the same village. An additional limitation of previous studies is that they do not collect information on network links outside the village. This limitation is especially relevant when the objective is to measure the household's centrality in its insurance network, as in Ambrus, Gao, and Milán (2017). As noted, the caste network spans a large number of villages.

The second limitation of the standard test of full risk sharing is that it does not account for other important aspects of consumption smoothing. Returning to our simple example, suppose now that there are N individuals in the mutual insurance arrangement. As before, income realizations are independent across participants, and each individual earns 1,000 with probability 0.5 and 2,000 with probability 0.5. The income-generating process is the same for all individuals, so we

continue to assume an equal income-sharing rule.

It is straightforward to show that mean consumption with insurance, M_I, is equal to mean income or, equivalently, mean consumption under autarky, M_A, which is 1,500 in our example:

$$(3) \quad M_I = E\left(\frac{1}{N}\sum_i y_i^s\right) = \frac{1}{N}(NM_A) = M_A.$$

Although mean consumption is independent of the number of participants, the variance of consumption with insurance, V_I, is decreasing in N:

$$(4) \quad V_I = V\left(\frac{1}{N}\sum_i y_i^s\right) = \frac{1}{N^2}(NV_A) = \frac{V_A}{N}.$$

Extremely high levels of risk sharing have been documented throughout the developing world. However, risk sharing does not map one for one into consumption smoothing. There is full risk sharing in our example with two individuals, but consumption varies substantially across states. The existing literature has failed to recognize this distinction, focusing exclusively on risk sharing and, more recently, on deviations from full risk sharing. Given that the deviations from full risk sharing are generally small, what makes caste-based networks exceptional in their ability to smooth consumption (which is what we ultimately care about) is their size and scope. A caste consists of thousands of individuals spread over a wide area, whereas informal networks in other countries are often organized around an extended family residing in the same village; see, for example, Angelucci, De Giorgi, and Rasul (2015). Large N reduces the variance of consumption, and the wide spatial scope of caste networks diversifies income shocks, reducing the variance of consumption even further. Without adequate access to private credit and government safety nets, the exceptionally well-functioning caste-based insurance

network provides a vital and irreplaceable service (in addition to other services when markets function imperfectly that will be discussed in the sections that follow).[7] This may explain the continuing relevance of the *jati* in the Indian economy and society. To the extent that marriage within the *jati* is a prerequisite for access to caste-based networks, this also explains why most Indians continue to marry within their caste.

3.2 *Caste-Based Insurance and Restrictions on Mobility*

The preceding view of caste-based insurance is that it has emerged, and persisted, because market credit and government safety nets are inadequate. While this view would suggest that caste-based insurance networks are efficiency enhancing, Munshi and Rosenzweig (2016) show that they can inadvertently generate distortions in other markets, specifically the labor market. Their research is motivated by two stylized facts: (i) the rural–urban wage gap in India is exceptionally high relative to other large developing countries and has remained this way for many decades, and (ii) the rate of permanent male migration has, nevertheless, remained exceptionally low over this period. This implies that there must be some friction that prevents rural workers from moving to the city and taking advantage of what appears to be an arbitrage opportunity.

Munshi and Rosenzweig posit that the source of this friction is the loss in (rural)

---

[7]Data from the 1982 REDS indicate that loan terms are substantially more favorable for caste loans, on average. Caste loans are significantly more likely to be interest free than loans from banks, employers, and moneylenders. They are also significantly more likely to be collateral free than loans from banks. Among the caste loans received in the year prior to the 1982 survey, 20 percent by value required no interest payment and no collateral. The corresponding statistic for the alternative sources of credit was close to zero, except for loans from friends where 4 percent of the loans were received on similarly favorable terms. Similar results are obtained with the 2005 IHDS.

network insurance when one or more members of a household migrates permanently to the city. Households with migrant members will have reduced access to rural caste networks for two reasons. First, migrants cannot be as easily punished by the network, and their family back home in the village now has superior outside options (in the event that the household is excluded from the network). It follows that households with migrants cannot credibly commit to honoring their future obligations at the same level as households without migrants. Second, an information problem arises if the migrant's income cannot be observed. If the household is treated as a collective unit by the network, it always has an incentive to misreport its urban income so that transfers flow in its direction. If the resulting loss in network insurance from migration exceeds the income gain, then large wage gaps could persist without generating a flow of workers to higher-wage areas. This distortion is paradoxically amplified when the informal insurance networks work exceptionally well, as in the Indian case for reasons listed above, because rural households then have more to lose by sending their members to the city.

The simplest test of the hypothesis that the potential loss in network services restricts mobility in India would be to compare migration rates in populations with and without caste-based insurance. This exercise is infeasible, given the pervasiveness of caste networks. What Munshi and Rosenzweig do instead is look within the caste and theoretically identify which households benefit less (more) from caste-based insurance. They then proceed to test whether it is precisely those households that are more (less) likely to have migrant members.

The literature on mutual insurance is concerned with ex post risk sharing, taking the size of the network and the sharing rule as given. To derive the connection between insurance networks and permanent migration, however, it is necessary to take a step back and model the ex ante participation decision and the optimal design of the income-sharing rule. Under the assumption that there is full risk sharing, households can either remain in the village and participate in the insurance network or send one or more of their members permanently to the city, increasing their income but losing the services of the network. The sharing rule that is chosen in equilibrium determines which households benefit more from the insurance network and, therefore, choose to stay.[8]

To derive the sharing rule, and associated migration, we begin with the household's decision rule. With log preferences, the expected utility from consumption can be derived as an additively separable function of mean consumption and its coefficient of variation (squared). With this specification, a household will choose to remain in the rural insurance network and forego its urban opportunities if the following condition is satisfied:

$$(5) \qquad \log(M_I) - \frac{1}{2}\frac{V_I}{M_I^2}$$

$$\geq \log(M_A) - \frac{1}{2}\beta\frac{V_A}{M_A^2} + \epsilon.$$

As defined above, M_I, V_I refer to the mean and variance of consumption with insurance. Here, M_A, V_A are the mean and variance of the household's income, that is, its consumption

---

[8] In related research, Morten (2019) derives the relationship between *temporary* migration and mutual insurance. Her model extends Ligon, Thomas, and Worral (2002) by allowing households to enhance their income in low rural-income states by sending migrants temporarily to the city. As in all models of ex post risk sharing, the size of the network is exogenously determined and there is no (permanent) exit. Morten follows Ligon, Thomas, and Worral by allowing the income-sharing rule to adjust in particular income states to satisfy the enforcement constraint. However, the baseline income-sharing rule (in the absence of these constraints) is exogenously determined.

in autarky when all its members remain in the village. This income would increase to $M_A(1 + \tilde{\epsilon})$ if one or more household members moved to the city, where $\epsilon \equiv \log(1 + \tilde{\epsilon})$ measures the household's urban opportunity (in utility units), which is unobserved by the rest of the rural network.

While migration will increase the household's income, this increase must be weighed against the additional consumption risk that a household with migrants must face. If a household remained in the village but did not have access to network insurance, then its consumption risk, which would be the same as its income risk, would be $(1/2)\left(V_A/M_A^2\right)$. The risk parameter β reflects both the decline in income risk through the diversification in household income that accompanies migration as well as the availability of nonnetwork insurance. We thus expect that $\beta < 1$. Nevertheless, the key assumption of the model is that $V_I/M_I^2 < \beta\left(V_A/M_A^2\right)$. This is the wedge that permits large wage gaps to exist in equilibrium. Retaining the assumption that the rural income-generating process is the same for all households, which implies an equal income-sharing rule and, hence, from equation (3) that $M_I = M_A$, a household will only send migrants to the city if $\epsilon > \beta\left(V_A/M_A^2\right) - \left(V_I/M_I^2\right)$. Since we have fixed the income-sharing rule, for the time being, we only need to solve for equilibrium participation, that is, the threshold ϵ_I above which households will send migrants to the city.

Once we introduce heterogeneity in rural incomes, the income-sharing rule and migration in each income class must be simultaneously determined. Munshi and Rosenzweig (2016) assume that the community's objective is to maximize the total surplus generated by the insurance arrangement (placing equal weight on all income classes), taking account of participation in each income class. With diminishing marginal utility, the total surplus

generated by the insurance arrangement can be increased by redistributing income so that relatively poor households consume more than they earn on average, that is, $M_I > M_A$. This gain from redistribution must be weighed against the cost to the members of the network from the accompanying decline in its size, since relatively wealthy households will now be more likely to leave and smaller networks are less able to smooth consumption. The first prediction of the model is that the income-sharing rule will nevertheless be set so that there is some amount of redistribution in equilibrium. This implies that relatively wealthy households within their caste benefit less from the network and so will be more likely to have migrant members ceteris paribus.

The second prediction of the model is that households facing greater rural income risk will be less likely to have migrant members. With full risk sharing, all households in the insurance network face the same consumption risk, regardless of their income risk. This implies that households facing greater income risk benefit more from the network and so will be more likely to stay. This prediction is especially useful in distinguishing Munshi and Rosenzweig's hypothesis, based on the rural insurance network, from alternative explanations for the large wage gaps and low migration in India. One alternative explanation is that a small number of castes fortuitously managed to establish networks in the urban labor market, extracting wage rents for their members and excluding outsiders. There are also alternative explanations available for redistribution within the caste and the increased exit from the network by relatively wealthy households. However, none of these explanations imply that households facing greater rural income risk should be less likely to have migrant members.

Munshi and Rosenzweig begin their assessment of the theory by showing that there is substantial redistribution of income

within castes, using data from the most recent ICRISAT panel survey and from the 2006 REDS. We know very little about the distribution of resources within castes, and this new finding will be corroborated below with data from urban India. Munshi and Rosenzweig then show, using the 2006 REDS village census, that relatively wealthy households within their caste are significantly more likely to report that one or more adult male members have permanently left the village. Highlighting the role that rural income risk plays in the migration decision, they also find that households with a higher coefficient of variation in their (rural) income—who benefit more from the rural insurance network—are less likely to have migrant members. Additional results, using the 1982 and 1999 REDS rounds, directly support the key assumption of the model, which is that migration should be associated with a loss in network services (and increased marriage out of the caste).

Having found evidence consistent with the theory, Munshi and Rosenzweig proceed to estimate the structural parameters of the model. Migration and the income-sharing rule (the λs in equation (1)) are determined jointly in the model. With logarithmic preferences, the λs map directly into relative consumption across income classes. The estimated λs match closely with relative consumption, which is available from the REDS but not used to estimate the model. This close match effectively validates the model, providing empirical support for the assumptions that preferences are logarithmic and that the community's objective is to maximize the surplus generated by the insurance arrangement (placing equal weight on all income classes). The substantial redistribution that is uncovered appears to be driven entirely by an attempt to equate marginal utilities across income classes. Counterfactual simulations that quantify the effect of formal insurance on migration, leaving the rural insurance

network in place, indicate that a 50 percent improvement in risk sharing for households with migrant members (which is still some way from full risk sharing) would more than double the migration rate from 4.3 percent to 9 percent. In contrast, doubling the rural–urban wage gap from 18 percent to 36 percent without any change in formal insurance increases the migration rate by just 1.5 percent. These results emphasize the central message of this research, which is that exceptionally well-functioning (caste-based) rural insurance networks coupled with inadequate access to formal insurance, rather than wage differentials as commonly assumed in models of migration, may explain much of the low mobility in India.

One way to circumvent these restrictions on mobility would be for members of the rural community to move to the city as a group. Members of the group could monitor each other and enforce collective punishments, solving the information and commitment problems described above. They would also help each other find jobs at the destination. As described in the sections that follow, the history of industrialization and urbanization in India is indeed characterized by the formation and evolution of caste-based urban networks, sometimes over multiple generations. We will see that a limitation of this strategy is that a sufficiently large (common) shock is needed to jump-start the new network at the destination, and such opportunities occur relatively infrequently. Thus, while members of a relatively small number of castes with (fortuitously) well-established destination networks can move with ease, most potential migrants will lack the social support they need to move. A second strategy to reduce the information and enforcement problems that restrict mobility is to migrate temporarily. The principal limitation of the temporary migration strategy is that it will not fill the large number of jobs in which there is firm-specific or task-specific

learning and where firms will set permanent wage contracts. Nevertheless, seasonal migration has been increasing rapidly over time in India (Morten 2019). The anecdotal evidence, for example Dhillon, Iversen, and Torsvik (2013), suggests that caste networks are once again playing an important role in this migration, and so this would appear to be an important and open area for future research.

4. *Business*

4.1 *Historical Background*

Business success in a developing economy requires connections, know-how, and access to capital. This effectively shuts out most potential entrepreneurs unless they have the good fortune to be born into a business family. The obvious question to ask, then, is how first-generation entrepreneurs enter business in the first place, particularly in a country like India. In Weber's (1958) view, the stability of the caste hierarchy served as a natural barrier to occupational mobility and entrepreneurial activity. This view has been contested by historians such as Rudner (1994) and Chandavarkar (1994) who take the position that mobility is possible even in a caste-based society, but that this mobility involves the movement of the entire group. We will see below that caste or *jati* networks have historically supported, and continue to support, the movement of first-generation entrepreneurs into business. Our objective will be to examine the conditions under which such transitions take place.

One varna, the Vaisyas, traditionally controlled moneylending and trade in India, with castes drawn from this varna active in different regions of the country. Non-Hindu communities such as the Jains, Bohris, and Khojas were also involved in trade from ancient times. Some of these castes (communities) had migrated overseas, to East Africa and the Persian Gulf, by the seventeenth century (Mehta 2001). However, mercantile opportunities were relatively limited until the arrival of the British in the eighteenth century. Commercial activity under the British was concentrated around the ports of Bombay and Calcutta.[9] Parsis, Gujarati Banias, Bhatias, and Bohris dominated long-distance trade out of Bombay from the early nineteenth century (Gupta 2014), whereas the Marwaris were only able to establish a foothold in the Calcutta trade (against stiff competition from the British) toward the end of that century (Timberg 1978). Over time, these communities moved from trade to manufacturing, as early as the middle of the nineteenth century in Bombay (Nafziger 1978) and from 1914 onward in Calcutta (Lamb 1955).

Apart from providing new opportunities in trade and manufacturing, British colonial rule also encouraged the settlement of castes in overseas destinations. The Nakarattar Chettiars, for example, served as moneylenders for the British throughout Southeast Asia, directly providing investment capital to farmers, plantation owners, and mine operators from 1870–1930 (Rudner 1994). Starting from the late nineteenth century, the *patidars* (Patels) of Charotar in Gujarat moved to British colonies in East Africa where they set up business and, in time, dominated commercial activity (Mehta 2001). While the Chettiars were forced to return to India in 1930, the Patels remained in East Africa till the 1970s, when they too were forcibly evicted (many of them subsequently moved to the United Kingdom where they now run successful businesses).

Although some of the castes listed above, such as the Marwaris and Chettiars, were historically engaged in business-related activities, colonial rule provided the first

---

[9] These cities were recently renamed as Mumbai and Kolkata, respectively. However, we will retain the historical names throughout to maintain consistency.

opportunity for other communities to move into business. These opportunities for group-level mobility have continued after independence in 1947, particularly after the liberalization of the Indian economy in the 1990s, with the emergence of a new business class drawn from a variety of agricultural castes and from castes that were historically concentrated in the bureaucracy and white-collar professions (Damodaran 2008). The striking feature of Indian business history is that the narrative is entirely about castes or their non-Hindu equivalent kinship communities; as Gadgil (1959, p. 16) puts it, "the history of the rise and growth of a modern business class in India is largely the history of the activities of members of certain groups."

Although occupational mobility did occur on occasion, the absolute number of castes and non-Hindu communities that moved into business over time is small. Timberg (1978) reports that twenty-three of the thirty-seven largest North Indian-owned industrial houses listed in the *Monopolies Inquiry Commission Report* of 1964 were Marwari or Gujarati Bania. More recently, a *Times of India* article (October 20, 2006) estimates that Gujarati-promoted companies account for 17 percent of the market capitalization of the Bombay Stock Exchange (BSE) BSE-500 index, followed by Marwari-promoted companies with 11 percent and Parsi-promoted companies with 8 percent. Public sector units, including banks and oil companies, account for 25 percent and all other companies, including multinationals, just 39 percent of the market capitalization.[10] Firms that are listed on the stock exchange, particularly those that are successful enough to be listed on the BSE-500, are operating outside the boundaries of their community network, almost by definition. What these

statistics tell us is that the pool from which these successful firms are drawn is not particularly socially diverse, that is, that businesses in India are drawn from a relatively small number of castes and communities. This is because once a community-based business network has formed, it will strengthen rapidly over time, bringing in hundreds of new entrepreneurs. Those entrepreneurs will support their children, which implies that the community network is less essential from the next generation. The critical challenge appears to be jump-starting the community network. The model that follows, which is a simplified version of Munshi (2011), characterizes conditions under which business networks will form and grow over time.

4.2 *Caste-Based Business Networks*

Consider an economy with a continuum of workers of unit mass entering the labor force in each period. There are two occupations in this economy: a traditional occupation and business. The payoff from the traditional occupation is V for all individuals. The payoff from business for individual i who enters the labor force in period t is determined by his ability, ω_i, and the size of his business network, which we measure by the stock of community members who are already in business; $\sum_{\tau=0}^{t-1} \Delta \omega_\tau$, where $\Delta \omega_\tau$ is the flow of individuals into business in period τ. The basic idea here is that the larger the existing network is, the greater is the support available to new entrants. Individual i will then select the business occupation if the following condition is satisfied:

$$(6) \qquad \omega_i + \beta \sum_{\tau=0}^{t-1} \Delta \omega_\tau \geq V.$$

As in Munshi and Rosenzweig's (2016) model of insurance and migration, the choice is between the location or sector where the network is active and going it alone. In their model, the insurance network at the rural

origin discourages individuals from moving independently to the city. In the current model, the business network at the destination encourages mobility; in each period t, there is a threshold ability, $\underline{\omega}_t$, above which individuals select into the business occupation, with the threshold declining over time as the network grows in size. Note, however, that there is no inconsistency between the two models; the general point is that community networks support group mobility but discourage the (independent) movement of individual members, and this will show up again in the subsequent section. Where the two models do differ is in the specification of preferences—we assume from now on that individuals are risk neutral—and that network effects occur with a lag (because individuals need to have some business experience before they can support others).

To initiate the network dynamics, $\Delta\omega_0$ individuals are moved exogenously into the business occupation in period 0. Assume that ability is distributed uniformly on the unit interval, $\omega \sim U[0,1]$. The flow of individuals into business in any subsequent period is then $\Delta\omega_t \equiv 1 - \underline{\omega}_t$. This allows us to derive the relationship between $\Delta\omega_t$ and the history of flows from equation (6),

$$(7) \qquad \Delta\omega_t = \beta\sum_{\tau=0}^{t-1}\Delta\omega_\tau - \alpha,$$

where $\alpha \equiv V - 1$. The assumption here is that $V > 1$, which ensures that no individual would choose the business occupation without the exogenous push in period 0. Equation (7) can be solved recursively to derive a closed-form solution for the flow of individuals into business in each period,

$$(8) \qquad \Delta\omega_t = (\beta\Delta\omega_0 - \alpha)(1 + \beta)^{t-1}.$$

The first implication of the preceding equation is that the business network will only grow over time if $\beta\Delta\omega_0 - \alpha > 0$, that is

if the initial push, $\Delta\omega_0$, is sufficiently large. This may explain why a relatively small number of caste communities have successfully managed the transition into business historically; it is not that they lack business potential, but that they were not fortunate enough to receive a large enough common shock that would have initiated subsequent group mobility.

Up to this point, the model has implicitly assumed that there is a single community (network) in the economy. In practice, multiple communities will coexist distinguished by their traditional occupations, which determine V, and their social connectedness, which determines β, the mapping from network size to business payoffs. The second implication of equation (8) is that once it does take off, the network will grow relatively fast in socially connected communities (with large β) and in communities with worse outside options (with small α and, hence, small V). This last result is especially interesting because it tells us that policies that attempt to jump-start business networks could be both growth enhancing and equalizing if they target historically disadvantaged communities.

Munshi (2011) tests the implications of the dynamic model with data from the Indian diamond industry. This industry occupies an important position in the Indian economy, accounting for roughly 14 percent of India's total merchandise exports, and has competed with textiles, and more recently with computer software, as the country's top export industry over the past decades. For the purpose of our analysis, the most interesting feature of the diamond industry is its social heterogeneity. India does not produce rough diamonds and so the diamonds must be imported, for the most part from the Antwerp market, and then cut and polished in domestic factories before being sold in Bombay to foreign buyers or directly abroad. Two traditional business castes—the

Marwaris and the Palanpuri Jains—controlled the business end of the industry from its inception in the mid-1960s, leaving the cutting and polishing to the Kanbi Patels, a community of lower-caste agricultural laborers known informally in the industry as the Kathiawaris (Engelshoven 2002). The industry structure changed in the late 1970s with the discovery of massive diamond deposits (one-third of the world diamond supply by volume) in Australia's Argyle mines. This supply shock allowed Kathiawari firms to enter business, and today all three communities account for a substantial share of the industry. Our interest is in characterizing the precise conditions under which a sufficiently large number of Kathiawaris were able to enter initially and then to compare the subsequent trajectory of the three community networks.

Because of the difficulty in enforcing legal contracts, the diamond industry is associated with a high degree of community networking throughout the world (Coleman 1988, Richman 2006). Over 96 percent of Indian diamond export firms are drawn from just three caste communities, indicating that networks must be active in the Indian industry as well. In the Indian industry, the role of the network appears to be to facilitate supplier credit. Most exporters visit the Antwerp market for a few days, once a month or every couple of months, to source rough diamonds. They tend to specialize in stones of a particular size, and while each exporter will have a small number of regular suppliers, he will also want to buy stones from other suppliers from one trip to the next. Given the high value of the diamonds, most Indian exporters must rely on supplier credit. What the network does is allow its members to receive credit from suppliers with whom they do not have long-term connections; other members of the community who do have established relations with those suppliers stand guarantor for them, with the threat of social sanctions ensuring that the beneficiaries of the referrals do not renege on their commitment.[11] An alternative strategy for sourcing rough diamonds, typically followed by established (second-generation) entrepreneurs, is to set up a branch in Antwerp. This allows the firm to build a reputation in the Antwerp market and operate independently of its caste network, both as an exporter and as a rough diamond supplier.

Retrospective information from a survey of 800 Bombay-based diamond export firms conducted in 2004–05 indicates that there was a sharp increase in the number of Palanpuri firms with branches in Antwerp just around the time that diamonds were discovered in the Argyle mines. The narrative in the industry is that these Palanpuris unloaded some of the increased supply of rough diamonds that came on the Antwerp market by encouraging their Kathiawari labor contractors, with whom they had established long-term bilateral relations, to enter business. A sufficiently large number of Kathiawaris must have entered at that time, because hundreds of Kathiawaris entered business in the subsequent decades (presumably supported by an underlying network that was growing in strength over time). This discussion highlights the difficulty in jump-starting community networks; a large supply shock was needed, together with support from outsiders (who were no longer reliant on their community network). Such a confluence of favorable circumstances is unlikely to occur very often, and if this is what is needed to jump-start a network, then it would explain why a relatively small number of castes are active in business today.

Once the dynamics have been initiated, the model tells us that networks will strengthen most rapidly in socially connected

---

[11]This cooperative arrangement has many features in common with Greif's (1993) description of the Maghribi traders' coalition a thousand years earlier.

communities and historically disadvantaged communities. Given the circumstances, we would expect the Kathiawari network to have strengthened relatively rapidly, with an accompanying decline in the ability of entrants from that caste over time. We motivated a role for community-based business networks with the argument that these networks substitute for a family background in business. This implies that the family background of the Kathiawari entrants should be declining particularly steeply over time, and this is precisely what we observe with data from the survey (which includes the establishment year of the firm and the family background of the entrepreneur).[12] Education, which is an alternative (and standard) measure of ability, also declines relatively steeply over time for the entering Kathiawaris, after accounting for age-cohort effects.

Our interpretation of the weakening ability of the entering entrepreneurs over time, especially the Kathiawari entrepreneurs, is that they are being supported by a strengthening caste-specific network. Administrative data on firm-level exports over a ten-year period (1995–2004) allow us to validate this interpretation. The Kathiawaris keep pace with their Marwari and Palanpuri rivals despite the observed decline in their ability. Indeed, once this compositional change is accounted for with firm fixed effects, the Kathiawari export trajectory is significantly steeper than the corresponding trajectories for the other two communities (or its own trajectory without fixed effects). There is evidently an underlying community-level force that is giving a boost to the Kathiawari firms, resulting in a dramatic shift from

farming to business over the course of a single generation.[13] Similar transitions have no doubt occurred before, but the diamond industry is not representative, and so we know relatively little about the precise circumstances under which other transitions occurred. Given the importance of entrepreneurship in the development process and the potential role played by caste networks in stimulating entrepreneurship in India, further investigation along these lines could add much value. This could involve additional industry-level or caste-level case studies, or the use of comprehensive census or survey data to examine the evolution of private enterprise in its entirety (subject to the constraint that caste identifiers are available at the firm level).

4.3 *Caste Networks and Static Inefficiency*

Once a business network has formed, we have seen that it can transform the circumstances of an entire caste in a single generation. From the perspective of the insiders, these informal institutions are entirely positive; they generate surplus for the caste by efficiently allocating resources internally, and they are redistributive, with less able individuals benefiting from the support provided by more capable caste members. However, we must also consider the consequences of these networks for outsiders (belonging to a different network or without a network of their own). One obvious inefficiency that arises in an economy where business is community based is that outsiders are crowded out; capable individuals do not receive the appropriate level of resources or are excluded from business activity entirely. The examples that

---

[12] While Marwari and Palanpuri entrants are slighly less likely to report, over time, that their father was a businessman, this decline is particularly steep for the Kathiawaris; 70 percent of the entrants in 1970 (before the supply shock and the establishment of the network) have a family background in business, but this statistic declines rapidly and drops below 20 percent for the entrants in 2000.

[13] Closer inspection of the entrepreneurs' family backgrounds indicates that most of the intergenerational mobility was driven by a shift out of agriculture into business; 52 percent of the Kathiawaris report that their father was a farmer (and this figure is much higher for recent entrants).

follow document such inefficiencies in urban and rural India, respectively.

Banerjee and Munshi (2004) examine resource allocation within and across castes in Tirupur's garment export cluster in a paper that predates the misallocation literature in macroeconomics. Tirupur, a town in the South Indian state of Tamil Nadu, supplies 70 percent of India's knitted garment exports. The textile industry in Tirupur was initially controlled by a local trading community. However, after a prolonged period of labor unrest in the 1960s, it was taken over by the Vellala Gounders, a wealthy farming caste (Swaminathan and Jeyaranjan 1994). Details of the circumstances under which the Gounders entered the industry are unavailable, but as with the diamond industry, there appears to have been a fortuitous element to the occupational transition.

Tirupur continued to be dominated by the Gounders for the next twenty years, with the industry producing exclusively for the domestic market. Starting from the mid-1980s, however, the export of knitted garments from Tirupur started to grow extremely rapidly (presumably as a consequence of another fortuitous shock), and by the early 1990s, the annual growth rate of exports exceeded 50 percent. This generated an influx of entrepreneurs from other regions of the country. In 1996, when Banerjee and Munshi conducted a survey of firms in the industry, about half the entrepreneurs were Gounders and the rest belonged to traditional business castes and communities, including Gujarati Banias, Parsis, and Marwaris.

The survey collected retrospective information on investment and production over a four-year period for each firm. Banerjee and Munshi's objective is not to model the entry of firms over time, as we did with the diamond industry, but instead to identify the mismatch between entrepreneurial talent (ability) and capital that arises because credit fails to cross community lines. Their analysis is motivated by two stylized facts: First, the Gounders hold more capital stock than the outsiders at all levels of experience. Adjusting for differences in production, the Gounders use roughly twice as much capital per unit of production than the outsiders. Second, exports grow faster for the outsiders than for the Gounders at all levels of experience. This is not mechanical convergence; the outsiders do start with lower export levels, but they outstrip the Gounders after five years of experience.

These two stylized facts, taken together, can be shown to imply that there must be a misallocation of resources across communities in this industry. To see this, let the firm's exports, Y, be determined by entrepreneurial ability, ω, and capital stock, K. We make the assumption that the two inputs are complements,

$$(9) \qquad Y = \omega K^{\alpha},$$

where $\alpha < 1$ ensures that the production function has the standard neoclassical properties. Let r be the interest rate faced by the firm. The firm will choose K to maximize its profit, which yields the following expressions,

$$(10) \qquad K = \frac{\omega^2}{4r^2}, \qquad Y = \frac{\omega^2}{2r}.$$

If all firms faced the same interest rate, as they would if markets functioned efficiently, then entrepreneurs with higher ability would hold more capital stock *and* have higher exports. This is evidently not the case in Tirupur. Indeed, it can be shown that the two stylized facts necessarily imply that the Gounders face lower interest rates, $r_G < r_O$, and have lower ability, $\omega_G < \omega_O$. To see this, substitute the expressions for K and Y from equation (10),

$$(11) \qquad K_G > K_O \ \Rightarrow \ \frac{\omega_G^2}{\omega_O^2} > \frac{r_G^2}{r_O^2},$$

(12) $Y_G < Y_O \;\Rightarrow\; \dfrac{\omega_G^2}{\omega_O^2} < \dfrac{r_G}{r_O}.$

Note that $r_G/r_O > (r_G/r_O)^2$ from inequalities (11) and (12), which implies, in turn, that $r_G/r_O < 1$. It follows, from inequality (12), that $\omega_G/\omega_O < 1$.

The Gounders are a wealthy landowning caste, and garment manufacturing was their first activity outside agriculture. Given that they have few alternative uses for their capital, it makes sense that (internal) capital is cheaper for them than for the outsiders. Although we do not model selection into the industry, it also makes sense that they have lower ability on average, simply because the pool of potential entrepreneurs is weaker. The inefficiency that Banerjee and Munshi uncover is that cheap capital fails to cross community boundaries, from the Gounders to the more productive outsiders, presumably because commitment problems preclude such trade.[14]

The inability to trade across castes is not restricted to industrial activity. Anderson (2011) provides evidence that caste is an impediment to trade in groundwater, a critical agricultural input, in North Indian villages, despite the fact that castes have been coexisting in these villages for centuries. Anderson's analysis exploits the fact that there are two types of villages in her sample; an upper caste is dominant, that is it owns the majority of the land, in 48 percent of Hindu villages, whereas a backward agricultural caste (BAC) is dominant in the remaining villages. Upper castes are almost entirely absent from the BAC-dominated villages. A comparison of the two types of villages yields the key empirical finding, which is that lower castes—BACs, OBCs, and SCs—have substantially higher incomes and crop yields in BAC-dominated villages. Anderson's explanation for this finding is that low-caste water buyers in the BAC-dominated villages have better access to this input because they can trade more easily with water sellers of their own caste.

One limitation of Anderson's analysis, which she acknowledges, is that all the BAC, OBC, and SC *jatis* are lumped together in a monolithic group.[15] Identifying the misallocation is also more challenging for Anderson because she does not have "conflicting" stylized facts, as do Banerjee and Munshi, that can be used to uncover an underlying market inefficiency. Anderson notes that the caste composition of the villages in her sample has likely been stable for centuries. What has changed is the pattern of land ownership in those villages where upper-caste Brahmins were absentee landlords in the colonial period. Land reforms in the 1950s, shortly after independence, transferred land ownership to lower-caste (BAC) cultivators, and it is those villages that are BAC-dominated today. Anderson provides evidence that the two types of villages are the same with respect to cultivated area, population, overall access to irrigation, and the functioning of the water market. But if this is true, then there is a comparability problem because the characteristics of the lower castes and the quality of their land will necessarily vary across the two types of villages.

---

[14] One alternative explanation for the two stylized facts is that ability and capital are substitutes, in which case higher-ability outsiders could grow faster even while holding less capital than the Gounders (with interest rates the same in both communities). But if this is true, then exports and capital stock should be negatively correlated *within* each community. However, within each community, where we expect entrepreneurs to face the same interest rates, firms holding more capital have higher exports (and export growth) consistent with the assumption that ability and capital are complements.

[15] Anderson notes that in her North Indian setting, the upper castes consist primarily of Brahmins and Rajputs, while the BAC group is comprised mainly of Yadavs. The results of her analysis, based on aggregate caste groups, are thus likely driven by underlying caste-specific interactions.

The lower castes were historically independent cultivators (with absentee landlords) in the BAC-dominated villages, and we would expect this autonomy to have had long-term effects; indeed, they have significantly higher education than lower castes in the upper-caste-dominated villages. A second source of variation is that lower castes in the BAC-dominated villages will own land across the quality spectrum, whereas we would naturally expect the upper castes to have captured the best land in the villages where they dominate. The available quality measure, based on land value, does not reveal systematic differences for the lower castes in the two types of villages. However, there is reason to treat this finding with caution. If yields (and incomes) are higher for lower castes in BAC-dominated villages, then this should necessarily be capitalized in the value of their land. The absence of such variation in land prices indicates that either the land market is functioning imperfectly, in which case the prices are not informative, or that there is measurement error in the reported prices, which prevents a meaningful comparison.

The preceding discussion indicates that a comparison of irrigation, yields, or income for the lower castes across the two types of villages will not identify a misallocation unless the comparability issues are addressed. However, the direct information on the caste-based restrictions to trade that Anderson provides is more compelling. The strongest empirical finding is that lower-caste water buyers have greater access to irrigation, with an accompanying increase in their yields and income, if they reside in a BAC-dominated village where water sellers are predominantly of the same caste group. The issue of comparability remains, but it is reduced because we are now comparing cultivators on the same side of the water market across villages. Given that access to irrigation for water buyers as a whole does not vary

across the two types of villages, it must be that upper-caste water buyers have greater access to irrigation in upper-caste-dominated villages (where water sellers are predominantly upper caste). This symmetry is reassuring, with the general finding that trade appears to be smoother when buyers and sellers belong to the same caste.

Taking this result as given, a natural question to ask is why potential buyers and sellers belonging to different castes cannot enter into long-term bilateral relationships. Social ties, and the social sanctions that go with them, are useful in supporting cooperation in a one-shot game. However, these farmers have been living together in the same village for centuries and, given the low mobility in rural India, would expect to live together long into the future. With a repeated game, caste ties would seem to be less relevant, with the potential termination of the long-term bilateral relationship being sufficient to deter deviation from the cooperative equilibrium. One possible explanation for the failure of buyers and sellers from different castes to come together is that they start with negative priors about the long-term stability of the relationship. Just as a sufficiently large common shock is needed to jump-start community networks, a positive shock is needed to create a long-term intercaste bilateral relationship. Anderson's results would suggest that such shocks occur fairly infrequently.

5. *Labor*

5.1 *Historical Background*

Castes were originally occupational units, and individuals remained within their caste's traditional vocation from one generation to the next prior to the arrival of the British. With colonial rule came overseas trade and, subsequently, industrial manufacturing. The accompanying growth of cities resulted in the demand for labor in nontraditional

occupations, in the mills, docks, and factories. A second new source of labor demand came from mines and plantations, both in India and overseas. Historical accounts indicate that networks organized around the *jati* played an important role in supplying labor to cities and plantations.

Commercial activity under the British was centered around the cities of Bombay and Calcutta. Bombay's textile-mill and factory owners appear to have adjusted their levels of output and employment to demand shocks, periodically bearing the cost of idle capacity (Chandavarkar 1994). Similar fluctuations in the demand for labor have been documented on the docks, in the shipping industry, and in general laboring jobs. These are precisely the conditions under which labor networks will emerge, with incumbent workers providing referrals for new hires from their social group. Social connections within the community provide the incumbent workers with information about the quality of new hires. The same connections ensure that new workers will not shirk, even if they are hired temporarily, to maintain the community's reputation in the labor market (and the incumbent worker's reputation in his firm).

In Bombay, the labor networks appear to have been organized around the *mukadam* or jobber, a foreman who was in charge of a work gang in the mill, factory, dockyard, or construction site, and, more importantly, also in charge of labor recruitment.[16] In an environment with frequent fluctuations in labor demand, we would expect the jobber to have recruited from his social group, which in the Indian context would be the *jati*, and this is precisely what appears to have happened. As Morris (1965, p. 76) argues, "I am using the term [caste] to refer to what are precisely called subcastes (*jati*) … for any

analysis of labor recruitment and the study of the behavior of rural recruits in an industrial environment … it is entirely inappropriate to lump into larger groups because of similarity of name, function, social status, or region-of-origin subcastes that are not endogamous."

The same pattern of caste-based recruitment appears to have emerged in Calcutta and in the mines and plantations. While business in Bombay was controlled by outsiders, labor was drawn mostly from the Maharashtrian hinterland. In contrast, labor in Calcutta's jute mills and factories was drawn from far away, from the western districts of Bihar, the adjoining eastern districts of Uttar Pradesh, and two Orissa districts, Cuttack and Balasore (Das Gupta 1981). Once again, a foreman or sardar was in charge of labor recruitment, and once again he appears to have drawn on his caste back home to fill new jobs. The same source districts in Bihar and Uttar Pradesh supplied the bulk of the labor to overseas sugar colonies—Mauritius, the West Indies, and Fiji—during the latter part of the nineteenth century, and while the recruitment of indentured labor to these colonies was heavily controlled by the colonial government, the concentration of particular castes among the migrants indicates that there was, nevertheless, an endogenous aspect to the labor supply.

Within India, caste-based recruitment emerged wherever labor was needed in large numbers. For example, caste-based recruitment was also observed in southern tea plantations. A particular feature of the caste system in South India was the institution of agrestic slavery in which the members of certain castes (and their progeny) were owned by upper-caste masters (Alexander 1989, Kooiman 1989). The abolition of slavery in India by the colonial government in 1861 coincided with the opening up of the plantations in South India and Sri Lanka. These plantations were located in

---

[16] The jobber was known as the *maistry* in the construction industry and the *sarang* in the shipping industry.

previously uninhabited hills and so labor had to be transported from the plains in Tamil Nadu. Initially, the British planters sent their own agents to recruit workers, but within a short time, workers began to migrate independently in gangs drawn from the foreman or kangany's *jati*. Our dynamic model tells us that once a network has formed, it will strengthen most rapidly in communities with the weakest outside options, which in this context would have been the former slave castes. Luke and Munshi (2011) surveyed a sample of South Indian plantation workers three generations later and found that 68 percent of the workers belonged to the former slave castes (Pallars and Paraiyars), highlighting the role played by caste networks in supporting mobility, not only into business but also into nonagricultural labor occupations.

Luke and Munshi show how caste networks in the plantations were refreshed and reinforced by new arrivals from the plains, through marriage, from one generation to the next. Networks that formed in the city were similarly reinforced over the course of many generations by the arrival of new workers from the rural origin. Focusing on a single caste, Dandekar (1986) traces the evolution of a network of Jadhavs (a particular caste) from a single village in interior Maharashtra to one jobber who went to Bombay in the 1930s, working first on the docks and later in the textile mills. In 1942, 67 percent of the Jadhav migrants from that village were working in the textile mills and 4 percent in other factories. In 1977, their descendants and new arrivals from the village over the next generation were still in the same occupations, with 58 percent in the textile mills and 10 percent in other manufacturing industries. De Haan (1997) similarly traces the evolution of a network of Khandait migrants in Calcutta to a single sardar. This individual brought an initial group of twenty workers with him from Orissa at the beginning of the

twentieth century, and this group was presumably large enough to jump-start a new network that persisted and grew in size over the next century.

The model of network formation and the preceding case studies emphasize the importance of initial conditions. When the model is extended to multiple destination locations or occupations, it generates the additional prediction that the community could get locked into those destinations where it (sometimes fortuitously) first established itself. Caste clusters have been documented among Bombay's dock workers (Cholia 1941), construction workers, and in the railway workshops (Burnett-Hurst 1925), the leather and dyeing industries, and in the Bombay Municipal Corporation and the Bombay Electric Supply and Transportation Company (Chandavarkar 1994). Castes sometimes found niches even within occupations; Gokhale's (1957) survey of textile workers in the 1950s found that particular castes were concentrated in particular departments within the mills, with different castes dominating a given department in different mills. De Haan (1997) documents similar caste segmentation in Calcutta's jute mills. Urban jobs did not necessarily relate to traditional rural occupations; indeed, castes that were traditionally weavers ended up in the spinning section of the jute mills by historical accident, emphasizing the importance of initial conditions in determining long-term outcomes when networks are active.

5.2 *Discrimination*

One side of caste networks is that they harness social ties within the community to solve market imperfections and increase economic efficiency. The other side is that they exclude outsiders once they are established. This discrimination is not a consequence of social prejudice, but is simply the outcome of economic competition between groups; the foreman would always prefer to recruit from

his own caste. In recent years, a number of studies have examined another aspect of caste discrimination in which upper-caste employers prefer not to hire lower-caste workers. This type of discrimination is not targeted at specific castes, especially in urban labor markets but is based on more broadly defined social status. While the *jati* may be the appropriate social unit when analyzing caste networks, the hierarchical aspect of the caste system, defined by the varnas, is relevant when analyzing discrimination. One approach that has been used to identify caste discrimination is the Oaxaca–Blinder decomposition. In this method, the earnings differential between upper castes and lower castes, which is attributed to discrimination, is divided into two components: differences in the estimated returns to endowments such as education and occupation for upper- and lower-caste groups and differences in the endowments themselves.

Banerjee and Knight (1985) apply the Oaxaca–Blinder decomposition to a sample of migrants that they surveyed in Delhi. Earnings are 17 percent lower for the SC migrants relative to the non-SC migrants. The two groups of migrants also differ in their endowments—education, the proportion who are salaried, and the proportion in unskilled manual jobs. The estimated earnings function coefficients reveal no difference between the two caste groups for education, age on arrival in Delhi, occupation, and the intercept. However, there is a greater premium for salaried workers versus daily wage earners among the SCs, a greater return to formal sector employment versus informal sector employment for the non-SCs, and a steeper earnings-experience profile for the non-SCs. Overall, differences in the estimated returns to individual characteristics favor the non-SCs but are small. Depending on whether SC or non-SC earnings function estimates are used, half to two-thirds of the earnings differential can be explained by differences in endowments (mainly education and occupation).

Madheswaran and Attewell (2007) apply the Oaxaca–Blinder decomposition to test for caste discrimination with National Sample Survey (NSS) data. The advantage of their data is that it is nationally representative and available at multiple points in time. SC and ST workers earn 5 percent less than other workers in 1983, 8.4 percent less in 1993–94, and 8.9 percent less in 1999–2000. Corroborating Banerjee and Knight's finding, the endowment component explains the bulk of the earnings differential (70–87 percent) over the three time periods. If occupational variables are included in the earnings function, the endowment effect increases even further. Kijima (2006) applies a variant of the Oaxaca–Blinder decomposition to explain the consumption gap between SC/ST and non-SC/ST households over multiple NSS rounds (1983–99). In contrast with the studies listed above, he finds that differential returns to endowments (education and occupation) explain half of the consumption gap, with the decomposition remaining stable over the sample period.

It is standard practice in this literature to interpret differences in the returns to endowments between groups as wage discrimination and differences in the endowments themselves as pre-market discrimination. This interpretation may, however, overstate the extent of current discrimination. The lower castes, particularly the SCs, were socially disadvantaged for centuries. They were originally assigned the lowest-status occupations, requiring little investment in human capital, and even when they managed to achieve occupational mobility under the British, they typically ended up in low-skill industrial jobs. We have noted that caste networks remain active in the Indian labor market. Once the upper-caste networks had established a presence in particular sections of the labor market, they would have ensured

that outsiders were subsequently excluded from those jobs. Persistent differences in occupations between SCs and non-SCs could thus be determined by historical circumstances rather than by discriminatory practices by upper-caste employers today.

Given historically low levels of human capital, and the importance of parental human capital in determining preschool human capital, the education deficit for the SCs could be similarly determined by historical circumstances that are slow to change rather than by current discrimination in education. Differences in the returns to endowments, particularly education, are also difficult to interpret as discrimination.[17] For example, well-qualified non-SC parents could transmit noncognitive skills and other traits that are complementary to education to their children. If these parental inputs are observed by the market but not the researcher, then the estimated returns to education would be lower for the SCs. Caste-based reservation in institutions of higher education would also affect the estimated returns to education. If affirmative action serves its intended purpose, then SC students enrolled in a given course in a given college will be less prepared than their non-SC classmates when they join. If they fail to catch up over the course of the program, then they will graduate with the same credentials but will be less competent. If this difference in competence is accounted for by employers, but unobserved by the researcher, then the estimated returns to education will, once again, be lower for the SCs. In general, it is difficult to disentangle discrimination from other channels, based on unobserved individual characteristics, through which intergroup earnings differentials are generated.[18]

Following Bertrand and Mullainathan's (2004) research design that was used to document racial discrimination in the United States, a field experiment in which caste affiliation is randomly assigned to otherwise identical job applications is the cleanest way to identify caste discrimination in the Indian labor market. Thorat and Attewell (2007) conducted such an experiment by replying to job advertisements in major English daily newspapers, sending three applications to each call: an applicant with an upper-caste name, an applicant with an SC name, and an applicant with a Muslim name. The applicants had identical educational qualifications and experience. No explicit mention of caste or religious background was made in any application; the only distinguishing feature was the name of the applicant. In addition, a "discordant" application was submitted to each advertisement; for jobs requiring an undergraduate degree, an applicant with an SC name and a master's degree was submitted, and for jobs requiring a postgraduate degree, an applicant with an upper-caste name and an undergraduate degree was submitted.[19]

The main finding of the study is that equally qualified applicants with SC or Muslim names were significantly less likely to be called for an interview or for testing. In the first discordant experimental treatment, the probability of success for the

---

[17] Banerjee and Knight detect no differences in the returns to education for SCs and non-SCs, whereas Madheswaran and Attewell's estimates indicate that the returns are consistently lower for the SCs.

[18] Deshpande and Newman (2007) compared the job-market expectations of SC and non-SC students in their final year of university in Delhi. The SC students expected a lower wage than the non-SC students, but this could be due to differences in endowments or perceived discrimination. The SC students performed worse than the non-SC students on the final university exam and had weaker family backgrounds (measured by parental occupation and education).

[19] Initially, one set of four applications was submitted for each advertisement. However, this was later increased to three sets to increase statistical power.

overqualified SC applicant was greater than (less than) the corresponding probability for an appropriately qualified SC (upper-caste) applicant. However, these differences are not significant at conventional levels. The second discordant treatment had greater statistical power, with under-qualified upper-caste applicants being significantly less likely (as likely) to be successful as appropriately qualified upper-caste (SC) applicants.

These experimental results indicate that unobserved individual characteristics that are correlated with caste identifiers may not entirely explain the findings from the Oaxaca–Blinder decompositions. It appears that private firms do select applicants explicitly by their caste status. The next question to ask is whether this discrimination is statistical; that is, firms use caste identifiers to proxy for applicants' unobserved characteristics that are relevant for their jobs, or taste-based, or in other words, based on preferences for a particular group. One way to dig deeper and answer this question is through qualitative research. Jodhka and Newman (2007) conducted in-depth interviews with managers responsible for hiring workers in twenty-five large private firms in the New Delhi area. The central finding was the view, expressed by virtually every respondent, that workers should be selected strictly on merit. However, merit was defined broadly enough to include family background, and while the respondents did not explicitly mention caste as a consideration, they did endorse regional stereotypes. In a situation where information on family background is missing, as in the field experiment above, they would very likely have used caste and regional characteristics as a proxy for family characteristics when making their hiring decision. Complementing this finding, Banerjee et al. (2009) conducted a resume audit study in Delhi that was restricted to software and call center jobs. They find no evidence of discrimination against SCs, STs,

or OBCs for software jobs. However, they do find large and significant differences in call-back rates between upper castes and OBCs (and to a lesser extent SCs) for call center jobs. These results are consistent with statistical discrimination, since unobserved family characteristics that are correlated with caste status are more likely to matter for the call center jobs. Statistical discrimination is not necessarily economically inefficient, but it does punish high-achieving members of underperforming groups, and it can lead to self-fulfilling stereotypical behaviors in equilibrium.[20]

5.3 *Caste-Based Labor Networks and Dynamic Inefficiency*

Although caste networks may reduce imperfections in the labor market, they can also generate inefficiencies of their own. We saw earlier that the failure of capital to cross caste lines results in a mismatch between entrepreneurial talent and investment in capital. Caste-based networks can similarly misallocate labor inputs by excluding competent outsiders from the market. Even the gain in efficiency within the caste through internal cooperation is not without cost. Munshi and Rosenzweig (2006) show that the same caste networks that supported the movement of their members from the village to the city historically discouraged the occupational mobility of their members in subsequent generations. This phenomenon is not unique to India. Heavily networked blue-collar communities, originally established by dynamic migrants, have been observed to be especially resistant to change

---

[20] There is no direct evidence that I am aware of that shows (conditional on observed credentials) that lower-caste workers are less productive. Deshpande and Weisskopf (2011) use data from the Indian Railways to show that an increase in the fraction of reserved workers does not appear to lower productivity. These results must be treated with caution, however, given the endogeneity of the inputs and the fraction of jobs that are reserved.

in other economies; for example, Gans (1962) and Kornblum (1974).

To provide economic foundations for this inertia, we return to the dynamic model of occupational choice that was described in the previous section. There was a continuum of ability levels in that model, which sorted individuals into business and the traditional non-business occupations. We now assume that there are two levels of ability, high and low, and two occupations, skilled and unskilled, in the labor market. Education plays an important role in the empirical application that follows by determining the type of job that the individual receives. We incorporate this in the model by assuming that high (H) education leads to the high-skill occupation, while individuals with low (L) education are channeled into the low-skill occupation. Ability sorts individuals into occupations once again, but it is now via the educational choices that they make. It costs $\overline{C}$ for low-ability individuals to attain high education, whereas the corresponding cost for high-ability individuals is $\underline{C} < \overline{C}$. We normalize so that the cost of attaining low education is zero.

Absent any other forces, high-ability individuals would select into the high-skill occupation if $W_H - \underline{C} > W_L$, and low-ability individuals would select into the low-skill occupation if $W_H - \overline{C} < W_L$. Once we introduce community networks, however, all individuals in a community could choose the same occupation, regardless of their ability, with these community-level choices persisting over many generations. A notable feature of the historical descriptions of urban caste-based networks in India is that they focus on working class jobs. This is a feature of labor market networks worldwide, possibly because the information and commitment problems that give rise to the networks are most acute in those jobs. In an early study set in Chicago, Rees (1966) found that informal sources accounted for 80 percent of all hires in eight blue-collar occupations and 50 percent of all hires in four white-collar occupations. Matching these patterns, 68 percent of blue-collar workers and 38 percent of white-collar workers reported having received help finding a job in Gore's (1970) study of migrants in Bombay. Thirty years later, Munshi and Rosenzweig (2006) document similar referral patterns in an area close to Gore's study, with 70 percent of blue-collar workers and 35 percent of white-collar workers reporting that they received help from a member of their caste in finding their first job. The model incorporates this feature of the labor market by assuming that the skilled wage, W_H, is exogenously determined, whereas the unskilled wage, W_L, is increasing in the (lagged) size of the community network in that occupation. In particular, members of the community from the previous generation in the low-skill occupation provide referrals, and increase wages, for the next generation.

Let $W_L(0)$ be the unskilled wage when no one from the individual's community selected that occupation in the previous period. Let $W_L(N)$ be the corresponding wage when all N members of the previous generation selected the unskilled occupation. If communities were concentrated entirely in either the skilled or the unskilled occupation in the initial period, then the following conditions ensure that there will be occupational persistence at the level of the community in all subsequent generations:

(C1) $W_H - \overline{C} > W_L(0),$

(C2) $W_H - \underline{C} < W_L(N).$

The first condition says that if the low-skill network was not active in the previous generation, then individuals of both types would

invest in high education in the current generation and end up in skilled jobs. The second condition says that if everyone in the community selected into the network in the previous generation, then individuals of both types would select low education and end up in the unskilled occupation. If conditions (C1) and (C2) are satisfied, it follows that communities will stay in the initial equilibrium from one generation to the next, with everyone either investing or not investing in education.

Economies will, however, restructure over the course of the development process, and one consequence of development is that the returns to education, and the skilled wage, will start to grow. When the inequality in condition (C2) is just reversed, high-ability individuals from communities that traditionally exclusively selected into the low-skill occupation will invest in high education and earn ϵ more in the skilled occupation than they would have if the entire community remained in the traditional occupation. If the fraction of high-ability individuals in these communities is non-negligible, then their exit from the network will result in a substantial (first-order) decline in the wage received by low-ability individuals. From a utilitarian perspective, welfare will decline, providing an economic rationale for community-based restrictions on mobility.

Individuals who select out of the traditional low-skill occupation will often move to a new location. The conventional punishment mechanisms that maintain cooperation within communities will then no longer be effective. Munshi and Rosenzweig argue that an alternative strategy to maintain cooperation in that case would be to instill a strong sense of community identity in childhood, which ensures that individuals remain tied to their community in adulthood.[21] The idea

that identity, and values more generally, are purposefully instilled to further community objectives is in line with previous work on this topic in economics; for example, Bisin and Verdier (2000) and Tabellini (2008). It also explains why heavily networked blue-collar communities tend to instill an especially strong sense of identity among their members in many economies.

The community identity described above aligns individual choices more closely with the social optimum. While this identity may thus be welfare enhancing when it is first put in place, it can result in a dynamic inefficiency if it persists in subsequent generations past the point where W_H has grown large enough that it is socially optimal for the high-ability individuals to exit. Cultural norms and social values are persistent by design, which explains why blue-collar communities often appear to stubbornly resist change (despite the fact that the same communities were extremely dynamic when they formed). Munshi and Rosenzweig build on this idea to document a dynamic inefficiency in Bombay's labor market, which was historically organized around caste-based networks. The starting point for Munshi and Rosenzweig's analysis is the economic and financial liberalization of the Indian economy in the 1990s. This increased white-collar wages (W_H in the model) substantially. The question is whether the occupational persistence associated with the blue-collar networks prevented the next generation of workers from taking advantage of these new opportunities.

However, the level of these transfers would need to be agreed upon long before individuals entered the labor market because investment in education is necessary for entry into the skilled sector. Reaching such an agreement ex ante is difficult in an economy where the (steeply increasing) wage trajectory is uncertain. The community would also need to credibly commit to making these transfers ex post, thereby creating potentially substantial income inequality among individuals doing the same job.

[21] Another possibility would be to provide compensatory transfers to those individuals to induce them to stay.

When the Indian economy restructured in the early 1990s, the next generation of workers would have still been in school. Munshi and Rosenzweig's strategy to test for occupational persistence exploits the fact that schooling choice is often a good predictor of future occupational outcomes. Schooling in Bombay, the setting for their research, can be either in English or Marathi (the local language). Marathi schooling channels the child into networked blue-collar jobs, whereas more expensive English schooling substantially increases the probability that the child will obtain a high-wage white-collar job in the future. The economic restructuring thus effectively increased the returns to English schooling. If networks discourage individual mobility, then parents in some castes could, nevertheless, continue to make schooling choices for their children that channel them into the traditional blue-collar occupation.

Munshi and Rosenzweig test for this intergenerational occupational persistence by examining schooling choices for twenty cohorts of Maharashtrian children who entered first grade between 1982–2001, in one area of Bombay city consisting of thirty schools (twenty Marathi and ten English). The key result is that the probability that a boy will be schooled in English is decreasing in the proportion of men in his caste employed in blue-collar (working class) jobs. This result is robust to the inclusion of parental education and household income as regressors. One concern is that children from lower castes, who were historically concentrated in the blue-collar occupations, have lower preschool human capital, which leads their parents to favor Marathi schooling. The parental and household variables control for this to some extent, but Munshi and Rosenzweig also implement an alternative approach that relies on the fact that the blue-collar networks were largely restricted to men. Male occupational choice in the previous generation should not affect girls'

schooling choices, and this is indeed what they find (once parental and household controls are included in the regression). Pooling boys and girls and including caste fixed effects, the second (complementary) result is that the gender gap in schooling choice is increasing in male occupational choice in the previous generation. More importantly, this relationship remains strong and stable over the course of the 1990s, even as the returns to English grew, providing microlevel evidence of the lock-in into specific traditional occupations.

At the aggregate level, this lock-in is reflected in the absence of convergence in schooling choice between upper-caste and lower-caste boys. Upper-caste boys who entered school in the 1980s are much more likely to be schooled in English than lower-caste boys, reflecting the fact that upper castes historically had access to administrative and professional jobs, while lower castes were concentrated in blue-collar occupations. While there is a steep increase in the fraction of boys enrolled in English schools in the 1990s, there is no convergence between the castes. This can be explained by the stable intergenerational correlation in occupational choice over the course of the 1990s that we noted above. In contrast, while there is also a caste gap in schooling choice for the girls in the 1980s, this gap narrows dramatically in the subsequent decade. Although lower-caste networks may have historically supported the movement of their members from the village to the city, with accompanying occupational mobility, many generations later they appear to hold their descendants back.

As described in the previous section, caste-based business networks are continually forming in the dynamic Indian economy. We would expect labor market networks to be similarly forming as new opportunities arise, with high-ability individuals leading the way into new occupations. At the same

time, the evidence presented in this section indicates that other caste networks will simultaneously be decaying, although social frictions could substantially delay this process. Munshi and Rosenzweig show that relatively high-ability children (where ability is proxied by the father's years of schooling) are the first to exit the traditional blue-collar occupation, providing a useful counterpoint to the patterns of selection that are observed with caste network formation. A fruitful area for future theoretical and empirical research would be to bring the twin processes of network formation and network decay together.

6. *Education*

6.1 *Aggregate Trends*

Given their historically disadvantaged circumstances, the lower castes lagged far behind the upper castes with respect to education and other socioeconomic indicators at independence in 1947. Article 46 of the Indian Constitution put into place one of the largest and most aggressive affirmative action programs in the world to rectify this historical disadvantage. Seats in universities and central government jobs are reserved for SCs and STs to match their share of the population, 15 percent for the SCs and 7.5 percent for the STs. More recently, reservation has been extended to OBCs, bringing the overall share of reserved seats close to 50 percent. In an economy where university seats are heavily oversubscribed and unemployment is high, the affirmative action program has naturally been highly controversial. Much of the public debate has centered on whether the SCs and STs have actually benefited materially from the reservations and, if so, whether privileged individuals within those groups have captured the bulk of the targeted resources.

Despite the importance of these questions, there have been few broad-based evaluations of the affirmative action program and, more generally, examinations of the relative circumstances of the SCs and STs in the decades after independence. As with labor market discrimination, it is appropriate to treat all castes among the SCs, and all tribes among the STs, as a single group if the objective is to evaluate the average effect of affirmative action on these historically disadvantaged groups. In the most comprehensive study to date, Hnatkovska, Lahiri, and Paul (2012) use data from five successive rounds of the NSS, covering the 1983–2005 period, to measure differences in education, occupations, wages, and consumption between SC/ST and non-SC/ST households. The key finding is that there is significant convergence in all of these outcomes over the period under consideration that extends well past the point in the early 1990s when the Indian economy started to liberalize. Decomposition analysis indicates that much of the convergence in wages and consumption is driven by convergence in education between the SC/STs and non-SC/STs. Moreover, the convergence is observed at all points in the respective distributions.

In a follow-up paper, Hnatkovska, Lahiri, and Paul (2013) compare the intergenerational mobility of SC/STs and non-SC/STs using the same data. Their analysis takes advantage of the fact that 60 percent of households in the NSS sample include members from different generations. Complementing their previous findings, intergenerational mobility of SC/STs was lower than that of non-SC/STs in 1983 but rises faster over the subsequent decades and is roughly equal by 2005. For example, the probability that an SC/ST child's education differs from his parent's education increases from 0.42 to 0.67 over the 1983–2005 period, whereas the corresponding probabilities for the non-SC/STs are 0.57 and 0.67. Similar convergence in intergenerational mobility is observed for income and occupation. Once

again, these gains in mobility are observed at all points in the distribution. In particular, the largest intergenerational changes in education for SC/STs are at the bottom of the distribution, out of illiteracy into primary and secondary school.

While these results are heartening, there are a couple of important caveats. First, the convergence cannot be entirely attributed to affirmative action (which targets higher education). Some of this catch-up could be mechanical, since the SC/STs started behind the non-SC/STs. Moreover, the government has taken steps to expand basic education as described below, and caste networks, which are active in many economic domains, could also have played a role in this convergence. Recall that networks will strengthen more rapidly, once they form, in historically disadvantaged castes with weaker outside options. Second, while there may have been convergence, substantial gaps in the *level* of education between SC/STs and non-SC/STs remain. Using representative NSS data, Hnatkovska, Lahiri, and Paul show that this gap is as large as two years of schooling. The discussion below provides micro-foundations for the observed convergence and the persistence of the caste gap, describing the challenges faced by SC/ST students in school, the benefits they receive through affirmative action in university, and the role played by caste networks in determining the demand for education.

6.2 *Access to Schooling*

It is well known that school participation is particularly low among the SC/STs. Drèze and Kingdon (2001) use data from the Public Report on Basic Education (PROBE) survey, covering randomly selected villages in the northern states of Bihar, Madhya Pradesh, Rajasthan, Uttar Pradesh, and Himachal Pradesh to compare school attendance and grade attainment for SC/ST and non-SC/ST children. Controlling for parental education, household wealth, measures of school quality, and teacher inputs, SC/ST children (and OBC children) are significantly less likely to attend school.

One explanation for this finding is that the school-quality and teacher-input variables do not fully capture the (lower) quality of education that lower-caste children receive. The government of India has made it a priority to provide a school within walking distance of each rural household. Given the segregation by caste within villages, this policy has resulted in lower-caste students being separated from upper-caste students, denying them the positive peer effects they would otherwise have enjoyed (Kochar 2008). The additional effect of this policy is that the average lower-caste student resides in a neighborhood of smaller size and will, therefore, be enrolled in a smaller school (with worse infrastructure) than the average upper-caste student (Rawal and Kingdon 2010).

Even if upper-caste and lower-caste students were mixed in school, the available evidence indicates they would not be treated the same. Rawal and Kingdon use data on primary school children in rural Bihar and Uttar Pradesh to show that student performance increases substantially when the teacher belongs to the same caste. Hanna and Linden (2012) conducted an experiment in which teachers were recruited to grade examinations. Candidate characteristics— caste, gender, and age—were randomly assigned to each cover sheet to ensure that there was no relationship between these observed characteristics and exam quality. The main finding is that exams assigned to lower castes are scored lower by examiners, although the effects are extremely small in magnitude. Interestingly, this bias is accounted for entirely by low-caste graders. It appears to be generated by statistical discrimination—a prior belief in that group of graders that lower-caste students are worse on average—because the bias disappears

over time as the link between demographic characteristics and test performance is revealed (this would not be the case with taste-based discrimination).

In addition to differential treatment by teachers, lower-caste students must also confront negative stereotypes when they interact with higher-caste students in the classroom. A large literature in social psychology finds that individuals belonging to a group associated with a particular stereotype perform worse on tasks associated with that stereotype when observed by others. In the Indian context, this implies that lower castes, who were historically perceived to be intellectually inferior, would perform worse on cognitive tasks if the results of their efforts were made public. Hoff and Pandey (2006) test this hypothesis by conducting an experiment in which low-caste and high-caste sixth and seventh graders were assigned the task of solving mazes in six-person groups. When the participants' caste status was kept anonymous, low-castes and high-castes performed equally well on the task. However, when caste identity was revealed, in homogeneous or mixed-caste groups, there was a sharp decline in low-caste performance. This result is obtained from a stylized lab-in-the-field experiment that may be subject to its own biases. For example, the high castes perform equally well when their caste status is kept private or is revealed in mixed-caste groups. However, their performance declines substantially when their caste status is revealed in high-caste-only groups. This curious result is difficult to explain, and it suggests that more evidence is needed, with more comprehensive data, in support of the experimental findings. Indeed, the available evidence indicates that conditional on attendance, grade attainment does not vary by caste (Kingdon 1998 and Drèze and Kingdon 2001). Where caste differences in performance do appear is in higher education, where affirmative action has substantially increased college attendance for (less qualified) lower castes in recent decades (Darity, Deshpande, and Weisskopf 2011).

6.3 *Affirmative Action*

As discussed in Frisancho Robles and Krishna (2012), there are three issues associated with affirmative action in higher education: (i) targeting—whether the beneficiaries are truly disadvantaged or whether the resources are captured by wealthy elites within the historically disadvantaged groups[22]; (ii) catch up—whether there is convergence between reserved groups, who start behind by design, and unreserved groups over the course of the university degree; and (iii) mismatch—whether the beneficiaries of affirmative action would have actually done better by not taking advantage of their legally mandated opportunity, perhaps because the college or the major they chose turned out to be too challenging.

Bertrand, Hanna, and Mullainathan (2010) address the targeting and mismatch issues with data on all individuals who took the engineering entrance exam in one Indian state in 1996.[23] The minimum score for admission, out of a total of 900, is 480 for the upper castes, 419 for the OBCs, and 182 for the SCs, confirming that affirmative

---

[22] A related concern is that individuals belonging to the upper castes could misreport their caste in order to take advantage of affirmative action. It is commonly believed that such identity manipulation does occur, although there are no statistics on the prevalence of this practice (for obvious reasons).

[23] As discussed below, Bagde, Epple, and Taylor (2016) use similar data with a different identification strategy to assess whether reservation improves college performance for beneficiaries. In contrast, Kochar (2010) and Frisancho Robles and Krishna (2012) focus on catch-up by examining student progress through college. A limitation of both studies is that they are restricted to a single college and lack information about the entire pool of candidates (differential selection by caste group into college is thus a major concern).

action does indeed bring in weaker students at the margin from historically disadvantaged castes. To gather more information on the socioeconomic background of the candidates and their post-college labor market outcomes, the research team conducted a household survey between 2004–06 in which the sample was based on candidates on either side of the cutoff test score for their caste group.

Combining the survey data with NSS data, which is representative of the state population, the first finding of the study, related to targeting, is that although lower-caste applicants are positively selected from their population subgroup, affirmative action does redistribute educational resources to less wealthy households. Lower-caste applicants to the engineering colleges have household incomes that are four times larger than the average income of their caste-group in the state. Nevertheless, the lower-caste applicants who are just above the test score threshold that secures them a position in engineering college have household incomes that are just over half the income of the upper-caste candidates whom they displace (with test scores just below the threshold for their caste group that would have secured them admission).

The second finding of the study, related to mismatch, is that despite the fact that lower-caste entrants have much weaker skills than upper-caste entrants (by virtue of their substantially lower scores on the entrance exam), they still gain by attending engineering college. However, this gain is more than offset by the loss in income for the high-caste candidates they displace. Thus, even if there is catch-up over the course of the engineering program, the two groups do not converge with respect to their earnings potential by the time the program is completed.

Bertrand, Hanna, and Mullainathan (2010) estimate the gain from affirmative action by comparing the income of lower-caste students above and below the entrance exam cutoff for their group. They estimate the location of the cutoff using a standard structural break approach; that is, they select as the cutoff the score that minimizes the sum of squared residuals in a regression of engineering college attendance on a dummy equal to one if an individual's test score is above a particular score threshold, zero otherwise. Lower-caste students just above and just below the cutoff can be assumed to have effectively the same intrinsic characteristics, and so a comparison of the two groups of students provides an estimate of the gain from affirmative action for the marginal student who benefits from that policy. In practice, Bertrand, Hanna, and Mullainathan sample applicants with scores in a relatively wide range around the threshold. While the two groups—above and below the threshold—do have the same observed characteristics on average, differences in unobserved characteristics are now a concern. Their point estimates are also now interpreted as the effect of affirmative action on a group (of students) that is less clearly defined.

A more recent examination of affirmative action in higher education by Bagde, Epple, and Taylor (2016) uses similar data but a different identification strategy to estimate the effect of reservation in engineering colleges in one Indian state on the performance of lower-caste (SC, ST, and OBC) students one year into the program. Their approach takes advantage of the fact that all entrants to the 215 colleges in the state must take a common entrance exam and, once admitted, sit for another common exam at the end of the first year. Bagde, Epple, and Taylor also obtain high school test scores for most (89 percent) of the candidates that provide additional information on intrinsic ability.

Instead of focusing on the marginal entrant, Bagde, Epple, and Taylor estimate the effect of affirmative action on admitted students over the entire ability distribution.

They do this by constructing two ranks: the absolute rank of each individual, based on the performance on the entrance exam, and the effective rank, which adjusts for the caste reservation. Conditional on the absolute rank or, equivalently, the score on the entrance exam, which measures intrinsic individual ability, the effective rank determines the effect of reservation on both entry into college and first-year performance (conditional on entry). High school performance provides an additional measure of ability and, in practice, latent ability is specified to be a polynomial function of the score on the college entrance exam and the high school exam.

The results of the analysis are as follows. First, the probability of college attendance is (not surprisingly) increasing in effective rank. Second, conditional on latent ability as measured above, performance on the first-year exam is increasing in effective rank. Note that the effective rank contains no information about individual ability beyond what is contained in the entrance exam score, that is, the absolute rank. This result thus provides an unbiased estimate of the effect of affirmative action on college performance. An augmented specification interacts effective rank with caste–gender dummies. The coefficients on the interaction terms are insignificant, indicating that all eligible caste groups benefit equally from affirmative action.[24]

The preceding results tell us that affirmative action increases the college attendance of reserved groups and improves their performance (conditional on attendance). However, they do not provide a mechanism through which these benefits are generated.

One natural mechanism is that affirmative action improves the quality of colleges to which reserved students have access, and this improves their performance. College quality, measured by the mean entrance exam rank of men in the unreserved (open) category, is increasing in effective rank. Instrumenting for college quality with effective rank, first-year performance is found to be increasing in college quality.[25] Contrary to the mismatch hypothesis, access to better colleges by less prepared lower-caste students significantly improves their outcomes.

Bagde, Epple, and Taylor conclude their analysis by quantifying these improvements. For this exercise they use the estimated coefficients from the analysis described above, replacing effective rank with absolute rank. The counterfactual analysis indicates that affirmative action increases engineering college attendance for men and women from reserved castes at all points in the entrance exam distribution, although the effects are strongest in the middle of the distribution. Despite these gains, the lower castes are underrepresented in the engineering colleges relative to their share of the population. With regard to the impact of reservation on test scores, average gains for SC and ST men are 0.12 and 0.18 standard deviations respectively, with similarly sized effects for SC and ST women. These magnitudes are roughly in line with the estimated effect from large-scale education interventions in the United States and other developing countries, including India.

Although caste-based affirmative action applies to institutions of higher education and government jobs, it can affect the choices that forward-looking parents make for their

---

[24]With a more parsimonious specification that includes effective rank and its interaction with the gender dummy, effective rank has a positive and significant effect for both men and women, although the effect is slightly weaker for women.

[25]Conditional on the latent ability measure, which subsumes the absolute rank, the effective rank can only affect first-year performance by changing the college that a student attends. The instrument thus satisfies the exclusion restriction.

children earlier in life. Cassan (2019) exploits two exogenous policy changes to estimate the effect of caste reservation on the educational choices of SC children. The first policy change redrew the borders of a number of Indian states in 1956. SC status was not harmonized within states at the time, resulting in the new residents of a state sometimes not having the same reservation status as the old residents, despite belonging to the same *jati*. This changed with the second policy change in 1976 that ensured that all members of a *jati* had the same reservation status in each state.

Cassan (2019) exploits the discrepancy in reservation status within *jatis* between 1956–76 to examine the effect of affirmative action, measured by SC status, on school enrollment. His difference-in-difference research design compares the schooling choices made for children after 1976 (the post period) with corresponding choices made before 1976 (the pre period) across two groups: the "new" SCs who were only reserved after 1976 and the "old" SCs who were always reserved. The baseline specification finds no effect of affirmative action; that is, the post-new-SC coefficient is close to zero and statistically insignificant. However, stronger results are obtained when the sample is partitioned by gender. The post-new-SC coefficient is now positive and significant for the boys, with weak and even sometimes negative effects for the girls, depending on the specification. Caste fixed effects are included in all specifications, and the identifying assumption is that new SCs and old SCs within a caste have the same intrinsic characteristics. If convergence is immediate, then this implies that new SCs and old SCs will have the same educational outcomes in the post-1976 period. To rationalize this with the estimated effects, it must be that the new-SC coefficient, which measures the difference in educational attainment between new SCs and old SCs in the pre-1976 period when only the old SCs had

access to reservation, is negative and significant for the boys but not the girls. This is indeed what Cassan finds.

6.4 *Demand for Education*

The preceding studies describe how different caste groups fare in the Indian educational system. While the convergence in education levels documented by Hnatkovska, Lahiri, and Paul may be generated, in part, by universal access to primary schooling and affirmative action in higher education, the differential *demand* for education at the narrower caste level may also have played a role in this convergence. Munshi and Rosenzweig (2006) document differences in schooling and occupational choice, by caste and gender, in Bombay city. In their urban setting, lower-caste men, who have been engaged in heavily networked blue-collar jobs for generations, are least responsive to the new opportunities that become available in the globalizing Indian economy. This inertia is reflected in the schooling choices made by lower-caste boys. Globalization will also affect schooling choices in rural India, where the majority of the population resides, but the new opportunities that arise will mostly be low-skill (urban) occupations that require a different type of education. Luke and Munshi (2011) exploit a unique setting—a group of tea plantations in the South Indian High Range—to assess how economic development will affect schooling choices, by caste, in the future. The key requirement when estimating differential effects by caste on schooling choices is to control for differences in household wealth across castes. As discussed below, Luke and Munshi's research setting satisfies this requirement.

The High Range, a mountainous area straddling the modern Indian states of Kerala and Tamil Nadu, was virgin forest until it was acquired by British planters and converted into tea plantations in the last quarter of the

nineteenth century. Because the plantation land was previously uninhabited, workers were brought to the High Range from the plains in Tamil Nadu. The workers on the tea plantations today are the third-generation descendants of those migrants, whose population has been supplemented by a fresh influx of new workers from the plains in each subsequent generation through marriage.

The workers in the High Range continue to be tied to their ancestral communities in rural Tamil Nadu, despite having lived on the tea plantations for many generations. Loans and transfers flow back and forth between the tea plantations and the origin communities, the children of the workers are often sent home to study, and many workers will buy land, build a house, and return to their ancestral villages when they retire. Perhaps the most distinctive feature of the South Indian kinship structure is marriage among close relatives (Kapadia 1995). Reinforcing existing network ties, many workers continue to marry their children to relatives from the ancestral location in the traditional fashion. The key parental decision is the extent to which children will be tied to their ancestral communities and the traditional economy. This decision will, in turn, determine the probability that children will marry among their relatives as well as their (optimal) educational attainment.

As in Munshi and Rosenzweig (2006), attachment to the traditional network lowers parents' incentives to invest in their children's education but for different reasons. First, education does not help the child (particularly the daughter) match with a better partner if she is married in the customary fashion to a relative. Second, attachment to the network reduces spatial mobility and accompanying job opportunities, as in Munshi and Rosenzweig (2016), with an accompanying decline in the returns to education. The heterogeneous caste composition in the tea estates allows Luke and Munshi to test this hypothesis. Recall that a particular feature

of the caste system in South India was the institution of agrestic slavery. Abolition of slavery coincided with the establishment of the tea plantations, and, not surprisingly, workers belonging to the former slave castes—the Pallars and the Paraiyars—are disproportionately represented in the High Range. The slave castes were historically spread over a very wide area, with their members having minimal social interactions with each other. We would expect their rural networks to be relatively weak and, thus, the expectation is that workers from those castes would have the greatest incentive to invest in their children's education.

Using data from a survey of plantation workers, augmented with administrative data on income over multiple years, Luke and Munshi show that workers belonging to the former slave castes have the same income and jobs in the tea plantations but are less attached to their ancestral communities and have higher levels of education than the non-slave castes.[26] It is worth noting that these statistics are very similar across the non-slave castes, despite the fact that they are drawn from very different positions in the social hierarchy.

Looking within the household, Luke and Munshi show that part of the difference in education across castes in the tea plantations can be explained by variation in the incentives of men and women to remain attached to their ancestral communities. Interactions with the ancestral community and the prospect of permanently settling there after retirement are especially burdensome for women belonging to the former slave castes. Thus, while both parents might benefit from

[26] The fact that Pallar and Paraiyar workers have the same jobs and earn the same income but have higher education is not necessarily inconsistent. Higher education for someone growing up in the tea plantations increases the probability of exit, with access to better job opportunities, in adulthood, but this outcome is not certain. What we observe are those individuals who stayed behind.

mutual insurance, old age support, and other forms of assistance organized around their ancestral communities, the cost of maintaining the social ties that are necessary to obtain this social support may be substantially larger for the women than the men in those castes.

The wider gender preference gap among the former slave castes is likely reinforced by the greater ability among the women in those castes to act on their preferences. Lower-caste women historically worked outside the home, and although they were typically employed in low-paying menial occupations, it has been argued that the income and independence that this work provided gave rise to a culture in which these women were relatively assertive within their households (Kapadia 1995, Chakravarti 1993). Given that culturally determined beliefs and practices are often slow to change, we expect these caste differences to persist in the High Range even though all women work in the tea estates, irrespective of their caste affiliation.

Luke and Munshi test for these gender differences by caste by estimating the effect of an increase in female income, net of total household income, on important decisions and outcomes for the children. Women have permanent full-time jobs as tea-leaf pluckers in the plantations, and they actually earn more than the men on average. Exploiting exogenous variation in male and female incomes generated by differences in rainfall across plantations and using a variety of independent estimation strategies, Luke and Munshi find that a relative increase in female income weakens the family's ties to the ancestral community as the woman gains bargaining power within the household. The children are less likely to marry a relative, be schooled in the ancestral location, and ultimately settle there. At the same time, an exogenous increase in relative female income increases the educational attainment of the children. Most importantly, significant female income effects are obtained for the former slave castes alone, accounting for part of the caste differences in schooling and marriage choice observed in both the parental and current generation.

The apparent willingness of the Pallar and Paraiyar women to act on their preferences does not, however, come without a cost. Previous evaluations of female income-generation programs in developing countries have documented a male backlash, as measured by the increased incidence of marital violence, when recently empowered women shift household decisions away from their husbands' preferences (Schuler et al. 1996). Similar patterns of marital violence have been documented in the international migration literature as well (Min 2001). Luke and Munshi provide evidence that a relative increase in female income among the former slave castes also leads to an increase in marital violence. In contrast, female income has no effect on household decisions and outcomes among the other castes in the tea estates and no effect on marital violence as well. Apart from providing independent support for the cross-caste results described above, this result also links marital violence, a serious social problem in India, to underlying caste-based institutions in a way that has not been done before.

Munshi and Rosenzweig's (2006) analysis of educational choices in urban India and Luke and Munshi's (2011) analysis of these choices in the tea plantations both indicate that historically disadvantaged caste groups may be especially responsive to new economic opportunities, leading them to invest in their children's education because they have fewer ties to the traditional economy to hold them back. Demand-side effects thus explain part of the convergence across castes that was documented by Hnatkovska, Lahiri, and Paul in recent decades. Luke and Munshi's analysis indicates, in addition, that disadvantaged individuals within those groups, notably the women, may be at the forefront of these

changes, providing an interesting and relatively unexplored gender component to the role played by caste in the changing Indian economy.[27]

7. *Politics and Access to Public Resources*

7.1 *State and National Politics*

Clientelism, a phenomenon by which politicians deliver benefits to specific individuals or groups of individuals in return for their votes, is pervasive in many developing countries. For a reciprocal arrangement of this nature to be sustained, each side must credibly commit to cooperate with the other. Politicians represent political parties, which have a long-term reputation to maintain with the electorate, and so will fulfill their promises even if they are elected for a single term and have personal preferences that do not align with the promises they made ex ante. While it is difficult to keep track of the voting behavior of individuals, well-organized groups can similarly credibly commit to honoring their reciprocal obligation to a political party (because they too have a long-term reputation to maintain). In India a natural social unit around which "vote banks" are organized is the caste or *jati*. We have seen in many different settings that social connectedness within the caste facilitates private economic activity by solving information and commitment problems. The same social ties can be used to support reciprocal

relationships between castes and parties, and it is thus not surprising that caste has played a central role in Indian democratic politics.

Yadav (1999) divides Indian politics after independence into three electoral systems. The first electoral system, covering the first four general elections from 1952–67, was dominated by the Congress Party. The opposition was regionally fragmented and often came from within the Congress. "The voter of course did not vote as an individual but rather as a member of a politicized *jati*" (Yadav 1999, p. 2394). For the caste to vote as a block, voters spread over a wide area must coordinate on a candidate. One way for the decentralized caste to coordinate in a given election would be with the help of a caste association. Rudolph and Rudolph (1960) note that caste associations did play a role in Indian electoral politics immediately following independence. However, relatively few Indians are members of formal religious or caste associations. For the caste to vote as a block, other coordinating mechanisms are needed.

A common feature of clientelist politics is the role played by brokers or middlemen who link voting blocks to parties. Srinivas (1951) describes how higher caste landowners and moneylenders served this function in the Mysore village that he studied. In other areas, representatives of the caste itself served this vital function; this is similar to the role played by foremen in linking workers from their caste to firms in the urban economy. The success of the Congress Party in the decades after independence has been attributed to its ability to establish a network of such middlemen (Kothari 1970, Kohli 1990). Although most Congress leaders came from the upper castes, the middlemen generated popular support for the party that was much broader than its leadership (Brass 1994).

Although the Congress Party may have been successful in organizing caste-based

---

[27] Given that lower-caste women face fewer restrictions on mobility and labor force participation, we would expect them to be more responsive to new economic opportunities than upper-caste women. However, an experiment providing basic financial literacy and business skills to women in Ahmedabad generated the curious result that upper-caste women but not lower-caste women responded to the intervention (Field, Jayachandran, and Pande 2010). The argument that the returns to training were for some reason higher from more restrictive groups in this setting does not square with the additional finding that there was no effect for the most restricted Muslim women.

voting banks, many castes were left out. Competition within the party for scarce resources also resulted in some castes benefiting more than others. According to Yadav, the second electoral system, starting in 1971, was characterized by the entry of middle-caste (OBC) politicians and the end of the complete domination of national politics by the Congress Party. As politics became more competitive, there was a consolidation of vote banks into state-wide multi-*jati* alliances, such as Kshatriya–Harijan–Adivasi–Muslim (KHAM) in Gujarat and Ahir–Jat–Gurjar–Rajput (AJGAR) in Uttar Pradesh. Nevertheless, the basic strategy of creating a multi-caste coalition of voters at the local level by delivering public resources to specific castes once elected continued to be followed by all political parties (Chhibber 2002).

This changed in the late 1980s with the onset of Yadav's third electoral system, characterized by the emergence of parties organized explicitly along (broad) caste lines. As discussed above, caste has always played an important role in Indian politics. But while the *jati* in its local setting constituted the voting block in the first electoral system and statewide alliances of *jatis* served this role in the second electoral system, the third electoral system witnessed the formation of a new caste identity based on *jati* groupings such as upper caste, OBC, or SC.

In Chhibber's view, the event that precipitated this change in the political equilibrium was the Mandal Commission report that expanded reservations in government jobs and educational institutions to the OBCs. The Congress Party was forced to distance itself from the report because forward, backward, and scheduled castes were part of its coalition. The social polarization created by the report thus opened the door for other parties to step in and represent the interest of broad caste groupings that were large enough to influence electoral outcomes at the state level. This reconfiguration of party politics explicitly on caste lines is most visible in Uttar Pradesh, India's largest state with a population of two hundred million, where the Bharatiya Janata Party (BJP) represented the upper castes, the Samajwadi Party (SP) represented the backward castes, and the Bahujan Samaj Party (BSP) represented the scheduled castes. Similar social cleavages in party politics also emerged in other states such as Gujarat, Maharashtra, and Bihar.

The advantage of clientelist politics is that it allows reciprocal relations between political parties and population groups to be sustained. The well-known disadvantage is that there will be an under-provision of public goods in favor of transfers targeted at the elected leader's vote bank (Bardhan and Mookherjee 2012). In the Indian variant of clientelist politics, as described above, the emergence of explicitly caste-based parties was accompanied by the formation of a new caste-group identity (Yadav 1999). Banerjee and Pande (2009) examine the (negative) consequences of this identity for the quality of elected leaders. Their basic idea is that if voters care about their leader's group identity, then the party associated with the numerically dominant caste group in a political constituency enjoys an electoral advantage. Its candidates will win even if they are less competent than their rivals. This implies that in constituencies with a numerically dominant caste group, an increase in caste-group identity will result in a decline in leadership quality if the elected candidate belongs to the party associated with that caste group. Conversely, if the elected candidate belongs to a party associated with the minority group in that constituency, he will be of relatively high quality.

Banerjee and Pande test this prediction with data from one hundred constituencies in Uttar Pradesh at two points in time, 1980 and 1996. Caste identity is believed to have increased substantially in Uttar Pradesh over

this period with the emergence of parties such as the SP, favoring the backward castes, and the BSP, favoring the scheduled castes. This implies that in a constituency dominated by low castes, elected politicians representing upper-caste parties should have relatively high quality in 1996. Conversely, in constituencies dominated by upper castes, elected politicians from the same upper-caste parties should be of relatively low quality. Politician quality is based on measures of corruption collected retrospectively from key informants (journalists and former politicians) in each constituency. The assumption here is presumably that more corrupt politicians capture more resources for themselves at the expense of the electorate.

The first result is that the probability that a constituency dominated by the low castes has a leader from an upper-caste party, that is, the Congress or the BJP, fell by 38 percent from 1980–96. The second result is that those leaders are substantially less likely to be corrupt in 1996 versus 1980, whereas the converse is true in constituencies dominated by the upper castes (where they are more likely to be corrupt in 1996). An additional prediction of Banerjee and Pande's model is that there should be no change in the quality of elected politicians over time in competitive constituencies where no caste group is dominant. This implication, which is especially useful in ruling out alternative explanations for the differential change in leadership quality across constituencies, is also borne out in the data.[28]

7.2 *Affirmative Action*

The change in Indian party politics over time would have served the interest of particular caste groups that succeeded in mobilizing politically. However, other caste groups could have been left worse off. Apart from these distributional effects via the competitive democratic process, the reservation for SCs and STs in state and national legislatures could also have shifted public resources toward these historically disadvantaged caste groups. In parallel with affirmative action in higher education and central government jobs, the Indian constitution also requires that reservation for SCs and STs in the national legislature reflect their share of the population in each state. If political parties can credibly commit to particular policies regardless of the preferences of their elected representatives, as assumed in the previous discussion on clientelist politics, then reservation of seats in specific constituencies (13 percent on average for the SCs and 7 percent on average for the STs) will have no effect. However, as Pande (2003) notes, if political parties cannot exercise complete control over their representatives once they are elected and if these representatives can influence the allocation of resources by the legislature, then reservation will improve outcomes for the beneficiaries of affirmative action.

Pande examines the effect of caste reservation on access to public resources using state-level data from 1960–92. Political reservation is measured by the fraction of jurisdictions reserved (separately) for SCs and STs in a state in a given year. Targeted public resources are measured by (i) the fraction of state-government spending devoted to SC and ST welfare programs (this is 6–8 percent of total spending on average), and (ii) the fraction of state government jobs reserved for SC/STs. One potential concern when estimating the relationship between targeted government spending and reservation is that these expenditures could have been made in the absence of reservation. A state with a greater share of reserved seats is, after all, a state with a greater share of SC/STs in the population, and they would naturally have a say in resource allocation.

---

[28] Banerjee and Pande's model also has implications for the difference in quality of the winning candidate and the runner-up that are borne out in the data.

Pande's solution to this identification problem, which is similar to Bagde, Epple, and Taylor's (2016) research design, takes advantage of the fact that the underlying variable upon which the affirmative action is based is not perfectly correlated with the fraction of SC/STs in the population at a given point in time. In Bagde, Epple, and Taylor's application, the effective rank in the list of engineering college applicants is positively (but not perfectly) correlated with the absolute rank, that is, the score on the college entrance exam. The latter variable can thus be used to control for individual ability when estimating the effect of reservation, measured by the effective rank, on college performance. In Pande's application, the slippage between the population share of the SC/STs and the fraction of reserved seats is generated by the time lag in changing quotas. These quotas only change in an election year and are based on the preceding census, whereas the population share changes continuously over time.[29] The main finding of the analysis is that conditional on the SC/ST population share, an increase in SC reservation has a strong positive effect on job quotas, but not on welfare programs, whereas an increase in ST reservation results in a substantial increase in ST welfare programs (without a significant effect on jobs or SC welfare programs).

In a related study, Banerjee and Somanathan (2007) examine the location of public goods between 1971–91 in approximately five hundred parliamentary constituencies in rural India to assess the changing relationship between the caste composition of the constituency and access to public resources. They consider fifteen public goods, covering the areas of education, health, water, electricity, and communication, with access to each good measured by the fraction of villages in the constituency where it is available. The 1971 data indicate that an increase in the share of SC/STs in the constituency is associated with a deficit in almost all public goods. Once the 1971 and 1991 data are pooled, however, the evidence indicates that there has been substantial convergence in access to public goods over time.[30] When this convergence is decomposed by caste group, it turns out that the gains are mostly in the SC-dominated areas, while the STs appear to have been left behind. The convergence documented by Banerjee and Somanathan can be explained, in part, by the political reservation policy and its redistributive effect that Pande finds. The reconfiguration of electoral politics on caste lines over time, coupled with the broad antipoverty programs that were pushed by the Indian government, could also have played a role in this convergence. The main message from the preceding analyses, complementing the findings of Hnatkovskaya, Lahiri, and Paul (2012) in the private sphere, is that there has been substantial equalization in access to public resources over time.

Did the targeting of public resources to specific caste groups have negative consequences for overall public good provision? Although the counterfactual experiment, without affirmative action, is unavailable, the evidence indicates that overall public good provision has, if anything, increased over time. Banerjee and Somanathan document a dramatic expansion in rural infrastructure between 1971–91. By 1991, primary schools had been established in 75 percent

[29] The population share, which is not directly measured, is derived by interpolating between census rounds. Because the preceding census is used for this interpolation, the reservation variable, which is constructed from the preceding census alone, provides no additional information about the actual population share.

[30] This convergence does not arise mechanically because access to public goods was close to complete in advantaged constituencies in 1971. Half the public goods that Banerjee and Somanathan consider were available in less than 5 percent of Indian villages in 1971 and in less than 10 percent of villages in 1991.

of all villages, and the share of villages with electricity increased from 18 percent to 70 percent. There was a ninefold increase in access to piped water over this period. Although there are many contributory factors, part of this dramatic increase can be directly attributed to targeted spending through affirmative action. Pande (2003) shows that ST (but not SC) reservation significantly increased per capita state government spending, and Chin and Prakash (2011), using a similar identification strategy, show that ST (but not SC) reservation significantly reduced poverty between 1960–2000. There are evidently important differences in the effect of SC and ST reservation, and these two population groups have fared very differently in the postindependence period. Further investigation of these differences would have great value for the design of development policy in India.

7.3 *Local Politics*

At the state and national level, the literature has devoted much attention to the role that caste plays in party politics. At the local level, parties are less relevant. In fact, many Indian states prohibit candidates in local elections from contesting on party lines. Caste matters now through the caste affiliation of the elected representative, and the focus of much of the literature on local governments in India has been to understand the relationship between the characteristics of the leader (including caste and gender) and the supply of public resources.

The seventy-third amendment of the Indian Constitution, passed in 1992, established a three-tier system of local governments or panchayats, at the village, block, and district level, with all seats to be filled by direct election. The village panchayats, which often cover multiple villages, were divided into 10–15 wards. Panchayats were given substantial power and resources, and

regular elections for the position of panchayat president and for each ward representative have been held every five years in most states. At the panchayat level, seats are reserved in each election for SCs, STs, and OBCs, but now in proportion to their share of the population in the district. Seats are, in principle, reserved randomly across wards and, for the position of the president, randomly across panchayats, from one election to the next in each district. The only restriction is that no seat can be reserved for the same group across consecutive elections within a constituency. In practice, however, seats have historically been reserved in earlier election rounds in constituencies with a larger share of SCs, STs, and OBCs, respectively (Dunning and Nilekani 2013).

Without the discipline of a party affiliation, and with limited possibility of reelection (given the frequent exogenous turnover due to randomized reservation), candidates in panchayat elections cannot credibly commit to exert the effort that is needed to ensure an adequate supply of public goods once in office. In addition, local political representatives are poorly compensated for their efforts; panchayat presidents are paid 50–60 dollars per month (less than the minimum wage), while ward representatives earn even less. In the absence of the standard political incentive mechanisms, Munshi and Rosenzweig (2018) posit that social ties within castes could be exploited to elicit effort by elected representatives. As discussed above, these ties have been used to support clientelist arrangements between political parties and *jatis* since independence. They have also been used to facilitate private economic activity in an economy where many markets function imperfectly. In the Indian local governments, they could be used to ensure that the representative puts in the requisite level of effort, even when he is elected for a single term, while at the same time ensuring that he is compensated (informally) for his efforts.

While caste networks may facilitate economic activity, their involvement in the economy generates particular inefficiencies, as discussed above. Many of these inefficiencies can be traced to the fact that cooperation and trade cannot cross caste lines. The use of the caste to solve political incentive problems is similarly associated with its own limitations. In particular, the competence of the candidate put forward by a caste and the effort he exerts, conditional on being elected, will be optimal from the perspective of the caste rather than the constituency as a whole. Because each caste accounts for only a share of the local population, there will be a suboptimal level of leadership competence and effort, with an accompanying undersupply of public goods.

To derive an empirical test of caste-specific cooperation and to quantify the resulting undersupply of public goods, Munshi and Rosenzweig develop a model in which N individuals belonging to K castes reside in a local constituency. Each caste k consists of N_k individuals, with the k subscript sorting castes by size; $N_{k-1} < N_k$. An individual from among the constituency's residents is elected to represent it in each election term. The representative's task is to supply a non-excludable public good to the residents of the constituency. Representatives receive no monetary compensation and are elected for a single term. A (long-term) reciprocal relationship between the representative and the entire constituency is thus infeasible. However, the representative's caste could credibly commit to compensate him ex post for his efforts through informal side transfers, even if he is elected for a single term, because they are connected to each other in many ways outside the political system. Once the side transfers are in place, each caste will select a candidate whose ability, and the effort he subsequently exerts if elected to represent the constituency, are optimal from its perspective. We first derive the effort, a,

that is optimal from caste k's perspective, taking as given the candidate's ability, ω. The effort is chosen to maximize

$$(13) \qquad N_k a^\beta - \frac{a}{\omega},$$

where a^β is the level of the public good, and the assumption is that the level of the public good maps one for one into the utility derived from its consumption by each resident of the constituency. The second term in the expression above measures the effort cost of the caste's chosen candidate, conditional on being elected. We make the usual assumption that the unit cost is decreasing in ability. Based on the solution to the maximization problem, the optimal level of effort from the caste's perspective, $a(\omega, N_k)$, is an increasing function of the candidate's ability, ω, and the size of the caste in the constituency, N_k.

The next step is to determine which individual will be selected by the caste as its candidate. Caste k will put forward as its candidate the individual with ability ω that maximizes

$$(14) \quad N_k \left[a(\omega, N_k) \right]^\beta - \frac{a(\omega, N_k)}{\omega} - \alpha\omega,$$

where the first term in the expression above measures the utility derived from the public good by the N_k members of the caste. The second term measures the candidate's effort cost, conditional on being elected, and the third term the opportunity cost to the candidate of holding public office. Although the representative could extract personal rents for himself and advance his political career, the private benefit of holding local public office is limited. The α parameter measures the difference in the returns to ability in the private sector and public office, which we thus assume is positive. Substituting the expression for $a(\omega, N_k)$ derived from the first maximization problem and then maximizing

with respect to ability, both ω and a are ultimately derived as increasing functions of N_k. Intuitively, because more individuals benefit from the non-excludable public good in a large caste, it will be in its collective interest to select a more competent candidate who will, in turn, exert greater effort. Cooperation within the caste results in a higher level of public good provision relative to the benchmark model where the representative only cares about his own utility (N_k would be replaced by one in the maximization problems above). However, public good provision is less than first best (where N_k would be replaced by N). The difference between the size of the representative's caste and the population of the constituency, N, determines the extent of the undersupply of the public good. This model of public good provision shares an important feature of previous models of insurance, migration, and occupational choice, which is that outcomes depend on the size of the relevant group, the caste network in previous applications, and the size of the caste in the constituency in the current model.

Munshi and Rosenzweig (2018) test the hypothesis that cooperation is restricted to the representative's caste and then quantify the resulting undersupply of public goods using data from the 2006 REDS, which covers a large sample of wards—the most local level of government—across the major Indian states over three election terms. The data include the caste and education of the elected representative, the supply of public goods at the street level, which can be mapped to the ward level, and the receipt of welfare transfers by specific households in each election term. In addition, the data provide the caste of every household in each ward. Providing direct support for caste-specific cooperation, there is a positive and significant relationship between the size of the elected representative's caste in the ward and both the representative's education, which proxies for his competence, and

the supply of public goods. In contrast, the size of the representative's caste in the village or the number of ward residents that belong to the representative's broader caste grouping, that is, SC, ST, or OBC, has no bearing on the supply of public goods.

Munshi and Rosenzweig quantify the undersupply of public goods due to caste-specific cooperation by estimating the structural parameters of the model and then conducting counterfactual simulations. Their analysis examines the supply of six major public goods falling within the panchayat's jurisdiction measured by whether there were expenditures on either new construction or maintenance for each of those goods on each street in the ward in a given election term.[31] Across the major Indian states covered by the REDS data, the average fraction of the ward population receiving a given public good across the six goods they consider is less than one-third in 40 percent of wards. The first counterfactual result is that if the ward representative internalized the benefit derived from the public goods by all residents rather than just his caste, then the entire population would receive all public goods in all wards.

The counterfactual simulations can also be used to assess the impact of existing policies on the supply of public goods. Ward representatives are responsible for the administration of welfare programs, in addition to their traditional role of delivering public infrastructure. Following Besley et al. (2004), Munshi and Rosenzweig measure welfare transfers by the receipt of Below the Poverty Line (BPL) cards that give recipients access to subsidized food and most government welfare programs. BPL cards are meant to be received by economically disadvantaged

---

[31] The six public goods—drinking water, sanitation, roads, electricity, street lights, and public telephones—account for 67.5 percent of the panchayat's discretionary spending. Education and health are not included in the list of public goods because they fall outside the purview of the panchayats (Bardhan and Mookherjee 2006).

households, but it is well known and well documented, for example, by Besley, Pande, and Rao (2012), that ineligible households that are politically connected can also benefit from them. If the total supply of the welfare transfers is fixed and the representative first targets his own caste, then outsiders will be worse off with respect to the welfare transfers when a larger caste is in power (as documented empirically by Munshi and Rosenzweig). This makes larger castes relatively unpopular with the electorate, reducing the likelihood that their candidates will be elected and, thereby, lowering the supply of public goods. Counterfactual simulations of the estimated structural model indicate that the effect of decoupling the delivery of public infrastructure and the administration of welfare programs on the supply of public goods would be modest. However, another policy—reservation for historically disadvantaged castes and tribes in panchayat councils—would appear to have more impact. Reservation will mechanically reduce the size of the representative's caste in the ward by restricting the set of castes that are eligible to stand for election. The disadvantaged minority groups will also account for a relatively small share of the ward population on average. Based on the counterfactual simulations, a policy that combines the decoupling of welfare transfers and public good provision with de-reservation would reduce from 40 percent to 20 percent the fraction of wards in which less than one-third of the population received each public good, averaged over the six goods.

A key assumption in Munshi and Rosenzweig's analysis is that the public goods they consider are non-excludable at the local—ward—level, whereas the welfare transfers can be targeted at specific households. Exploiting changes in the caste of the elected representative over election terms within each ward, they are able to verify this assumption. If the household belongs to the caste of the elected representative, it is significantly more likely to hold a BPL card compared to when the representative belongs to another caste. However, there is no evidence that public goods are targeted to streets on which the representative's caste is disproportionately represented. As with the results on the supply of public goods, the fact that a household belongs to the broad caste group of the elected representative, that is, SC, ST, or OBC, does not affect the probability that it will receive a BPL card conditional on whether or not the household belongs to the same caste. Previous research on caste-based targeting in Indian local governments has documented that SC/ST households are more likely to receive publicly provided private goods (welfare transfers) when the panchayat president's position is reserved for SC/STs; see Besley et al. (2004) and Bardhan, Mookherjee, and Torrado (2010). While it may make sense from a policy perspective to assess whether political reservation for particular caste groups benefits the members of those caste groups on average, Munshi and Rosenzweig's results indicate that targeting, at least at the ward level, is occurring at the finer *jati* level.[32]

Political clientelism is characterized by the transfer of targeted public goods, jobs, or services to groups of voters in return for their political support (Stokes 2009). The caste or *jati*, rather than caste groupings such as SC, ST, or OBC, is the social unit around which these reciprocal arrangements are organized at the local level (Brass 1994). This argument runs as a common thread through the preceding sections of the paper and is consistent with Munshi and Rosenzweig's

[32] Although most studies on targeting in panchayats focus on welfare transfers, Gille (2016) uses the 2006 REDS to show that individuals are more likely to apply for reserved public sector jobs when the panchayat president belongs to the same caste. Her interpretation of this result is that the president uses his political connections at the state level to help the applicant from his caste get a job.

finding that welfare transfers are targeted to the representative's *jati*. Their analysis goes beyond the standard clientelist model by allowing transfers to simultaneously flow in the opposite direction, from the caste to the representative. This informal compensation ensures that the competence of the individual selected to represent the caste, and the effort he exerts conditional on being elected, are optimal from the perspective of the caste. Anderson, Francois, and Kotwal (2015) extend the standard clientelist model in a different direction by allowing the elite group in power to transfer private rather than public resources to the electorate, in the form of insurance against income shocks, in return for their votes.

Anderson, Francois, and Kotwal's analysis is set in rural Maharashtra. An unusual feature of Maharashtrian society is that a single caste, the Marathas, accounts for a large fraction of the total population. Moreover, while Marathas are often the dominant landowning caste in the village, many Marathas do not own land and supply labor to Maratha and non-Maratha landlords. The starting point for Anderson, Francois, and Kotwal's model is the assumption that landlords do not value public programs because they raise wages and lower labor compliance. Political power allows the landlords to reduce these programs, but this is only possible if the electorate votes them to power. The landlords achieve this objective by offering private insurance to the landless in return for their votes. For this reciprocal arrangement to be an equilibrium, each side must credibly commit to honoring its obligation. An important assumption in the model is that the Marathas are more cohesive than other castes; in particular, only Maratha landlords can cooperate with each other. Given that the Marathas are the only caste on both sides of the clientelist arrangement, the social cohesion within that caste can also be exploited to sustain cooperation between Maratha landlords and workers. The additional advantage enjoyed by the Marathas is that they control local agricultural input and product trading networks. Maratha patrons can offer access to these networks to their non-Maratha clients in return for their votes. The more workers a Maratha landlord employs, the greater is the gain from suppressing wages. It follows that if the Marathas are the dominant landowning caste in the village, then it is in the interest of the Maratha landlords to establish a clientelist political arrangement.

To test this prediction, Anderson, Francois, and Kotwal estimate the following equation:

$$(15) \quad Y_{ik} = \beta_0 + \beta_1 MLD_k + \beta_2 MPROP_k$$

$$+ \beta_3 MLD_k \cdot MPROP_k$$

$$+ \psi Z_{ik} + \gamma Z_k + \epsilon_{ik},$$

where Y_{ik} is an outcome for household i residing in village k, MLD_k indicates whether village k is Maratha land dominated, $MPROP_k$ is the proportion of Marathas in the village, Z_{ik}, Z_k are household- and village-level controls, and ϵ_{ik} is a mean-zero disturbance term. The same specification is used for village-level outcomes, except that the individual, i, subscript is suppressed.

Based on the model, clientelism is more likely to be observed when $MLD_k = 1$. Thus, we expect public programs and wages to be decreasing in MLD_k, while access to private insurance (through their landlords) by the landless and landlord profits should be increasing in MLD_k. The sign of the coefficient on the interaction term, β_3, is ambiguous. If intra-caste ties between Maratha workers and landlords are more (less) important than the access to trading networks by non-Maratha workers through their Maratha landlords, then clientelism is increasing (decreasing) in the interaction term. In contrast, $\beta_2 = 0$ because if $MLD_k = 0$, then there is no clientelism,

irrespective of the caste composition of the village population. Using data from 320 villages in rural Maharashtra, Anderson, Francois, and Kotwal verify each of these predictions, with the sign of the β_3 coefficient indicating that access to trading networks dominates intra-caste cohesion as a determinant of clientelism.

Caste reservation is implicitly based on the assumption that clientelism is present, that is, that political representatives will target public resources to their own group. The available evidence indicates that disadvantaged minorities (SCs and STs) are more likely to benefit from welfare programs when the political representative belongs to their own caste. While reservation for disadvantaged minorities thus has distributional benefits, Munshi and Rosenzweig's (2017) analysis indicates that there can be negative efficiency consequences because reservation reduces, on average, the size of the representative's caste. The other obvious reason why the supply of public goods could decline with reservation is that the pool of potential candidates is weaker in historically disadvantaged minorities. However, new research by Anderson and Francois (2017) indicates that the effect could go in the opposite direction under certain circumstances, with reservation increasing the supply of public goods. Their analysis builds on Padró i Miquel's (2007) model of clientelism in which ethnic groups will tolerate the corruption of their co-ethnic political representative and continue to vote him into power when the alternative with another group's representative in power is much worse. Reservation generates exogenous turnover in the political leadership, reducing the ability of incumbent representatives to extract rents. Using the same data from a sample of Maharashtrian villages as Anderson, Francois, and Kotwal (2015), Anderson and Francois (2017) estimate the following equation:

$$(16) \quad Y_{ik} = \beta_0 + \beta_1 RESERVED_k$$
$$+ \psi_k X_{ik} + \gamma_k Z_k + \epsilon_{ik},$$

where Y_{ik} is the level of public resources received by household i residing in village k, $RESERVED_k$ indicates whether the panchayat election was reserved for lower castes (SC, ST, OBC) in village k, X_{ik} and Z_k are household and village-level controls, and ϵ_{ik} is a mean-zero disturbance term. The same specification is used for village-level outcomes, except that the individual, i, subscript is suppressed.

Based on the preceding discussion, the prediction of Anderson and Francois' model is that $\beta_1 > 0$. This is because in reserved elections, the next term will certainly be unreserved, which implies that the ability of the incumbent elected representative to extract rents is reduced. However, this argument does not apply if the *jati* in power accounts for a majority of the population, in which case its representative would remain in power, with or without reservation. It would also not apply if the representative's caste accounts for a small minority of the population, in which case he would only be elected by chance. Anderson and Francois thus focus on panchayat elections in which the caste in power comprises 25–50 percent of the population. In those elections, the β_1 coefficient is positive and significant. However, this coefficient is statistically indistinguishable from zero in elections where the population share of the representative's caste lies outside this range (this result is robust to the specific range that is specified). One limitation of this analysis, which also applies to Anderson, Francois, and Kotwal (2015), is that panel data are unavailable. The potential for omitted variable bias is always greater with a cross-sectional analysis. A second limitation of the analysis is that the electoral process and, hence, the caste in power, is not endogenized. Nevertheless, these results,

as in Munshi and Rosenzweig (2017), indicate that caste reservation interacts with the caste composition of the village in complex ways that have an important bearing on the supply of public resources in Indian local governments.

In parallel with the literature on caste reservation in local politics, there has also been substantial work on the reservation of seats in the panchayat council for women. Previous research on caste networks has found that men and women benefit differentially from these networks, with resulting differences in participation by gender (Munshi and Rosenzweig 2006, Luke and Munshi 2011). We would similarly expect caste and gender to interact in Indian local governments, and this is the subject of a new paper by Cassan and Vandewalle (2017). In a well-known paper, Chattopadhyay and Duflo (2004) examine the impact of gender reservation for the panchayat president's position on the choice of public goods. The main finding is that public goods that are preferred by women, based on the frequency of complaints lodged by men versus women in public meetings, are more likely to be supplied in election terms where the president's position is randomly reserved for women. Cassan and Vandewalle add to this research by adding a caste component to the gender reservation. In particular, they hypothesize that lower-caste women will be disproportionately more likely to be elected than upper-caste women in reserved elections. This is not because those elections are also more likely to be reserved for lower castes. Within each caste-reservation category and in open elections, one-third of the seats in each election term is randomly reserved for women. However, as discussed in section 6.4, lower-caste women have historically had higher rates of labor force participation and enjoyed greater autonomy than upper-caste women. It is entirely possible that they would also be more likely to participate in local governments. Using data from the 2006 REDS,

Cassan and Vandewalle find that lower-caste women are more likely to be elected to the president's position than upper-caste women in reserved elections (women are almost never elected in unreserved elections). They then proceed to show that goods preferred by women and, independently, goods preferred by lower castes, based on preferences obtained from the 2006 REDS, are more likely to be supplied in elections reserved for women. Research on caste that adds a gender component consistently uncovers interesting and novel findings. Further work in this area will likely have high payoff.

8. *Conclusion*

Caste continues to play an important role in the Indian economy. Networks organized at the level of the caste or *jati* provide insurance, jobs, and credit for their members in an economy where market institutions are inefficient. Affirmative action for large groups of historically disadvantaged castes in higher education and India's representative democracy has, if anything, made caste more salient in society and in the public discourse.

Newly available evidence with nationally representative data indicates that there has been convergence in education, income, occupations, and consumption across caste groups over time. If this trend continues, then there will no longer be a rationale for affirmative action. The available evidence indicates that caste discrimination, at least in urban labor markets, is statistical, that is, based on differences in socioeconomic characteristics between upper and lower castes. This type of discrimination will disappear once there is convergence between these broad caste groups. It is important to realize, however, that these changes will not occur in the near future. The current gap in educational attainment between SC/STs and upper castes is very wide, over two years of schooling. Given the strong intergenerational

persistence in human capital, the key variable driving convergence, it will be many generations before income and consumption are equalized across caste groups.

The caste-based economic networks that currently serve many functions will also disappear once markets begin to function efficiently. These networks continue to be active in the globalizing Indian economy because information and commitment problems are exacerbated during a period of economic change. In the long run, however, the markets will settle into place and the caste networks will lose their purpose. This has certainly been the experience in many developed countries. In the United States, for example, ethnic networks based on a European country (region) of origin supported their members through the nineteenth century into the middle of the twentieth century. Ultimately, however, these networks no longer served a useful role and today, outside of a few pockets, European ethnic identity in the United States is largely symbolic. We might expect caste to similarly lose its salience as India develops into a modern market economy, and there is some evidence that this process may have already begun.

Munshi and Rosenzweig (2006) found that lower-caste boys in Bombay were held back by their working class networks, whereas lower-caste girls (without that historical baggage) were apparently willing and able to take advantage of the new opportunities that became available with economic liberalization in the 1990s. This is reflected not only in their schooling choices, which will presumably map into their labor market outcomes, but also in their marriage decisions; specifically, whether they married within their caste. Figure 3 reports the marriage choices of female family members of the students who were sampled by Munshi and Rosenzweig. The marriage choices are reported in five-year intervals over a time period (1970–2000) that substantially overlaps with figure 1, which documented

corresponding marriage patterns in rural India in ten-year intervals from 1950–2000. While outmarriage remained extremely stable at 5 percent over the entire fifty-year period in rural India, outmarriage increases over time in the urban data. It is even lower than 5 percent, somewhere between 2–4 percent, prior to economic liberalization, but then starts to increase steeply in the 1990s and is as high as 12 percent in the 1995–2000 period. Munshi and Rosenzweig's study is set in Bombay city, which is at the vanguard of economic development in India. These results are thus almost certainly unrepresentative, but they do tell us what we might expect to see in the future.[33] Marriage within the caste is key to the efficient functioning of caste-based networks and the formation of a caste identity. The changes in marriage patterns uncovered in Bombay will likely be replicated elsewhere in the future as the economy develops, and this will signal the end of the caste system as we know it.

These social changes, just like the end of affirmative action, will take many decades (even generations) to unfold. In the interim, both research and policy would benefit from a recognition of the informal caste-based institutions that are active in the economy. A counterfactual policy experiment described in Munshi and Rosenzweig (2016) highlights the potential cost of ignoring these institutions. Recall that Munshi and Rosenzweig's explanation for India's low mobility is based on a combination of well-functioning (caste-based) rural insurance networks and the absence of formal insurance, which

---

[33] Similar results are obtained by Banerjee et al. (2013), who examine the marriage advertisements that were placed in an English-language newspaper in Calcutta and the matches that subsequently took place. They find that 72 percent of the prospective grooms and 68 percent of the prospective brides who had married after a year found a partner within their caste. These statistics are substantially lower than the standard figure of 95 percent obtained from nationally representative surveys.
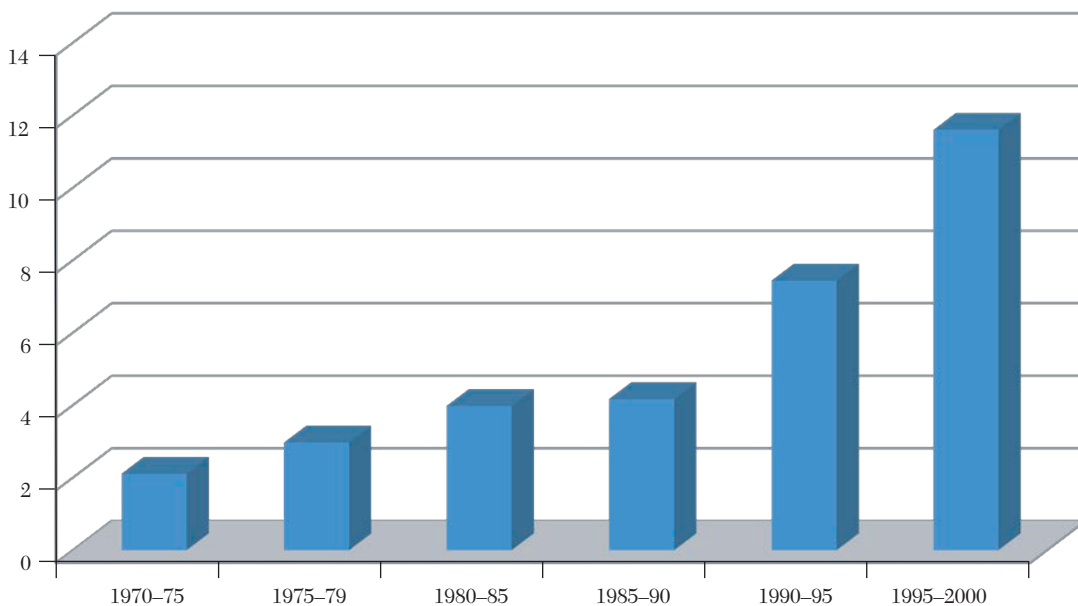
Figure 3. Marriage outside the Caste in Urban India

Note: Percentage of marriages by women in the student's family (sisters and mothers) outside their caste.

Source: Survey of school children and their families in Bombay, 2001.

includes government safety nets and private credit. Having empirically validated this explanation and having estimated the model of insurance and migration, they proceed to implement a counterfactual experiment in which private credit is made available to wealthy households (with the resources to put up collateral). We see that these households subsidize less wealthy members of their caste and thus are more likely to leave in the current equilibrium. While the absence of alternative consumption-smoothing mechanisms keeps their migration currently at a low level, this level increases dramatically in the counterfactual experiment when they get access to credit and are no longer reliant on the rural insurance network.

The increased migration by wealthy households leaves less-wealthy households in their caste worse off in two ways. First,

the size and scope of the insurance network is reduced by the exit of the wealthy households, making it function less effectively and increasing the consumption variance of the households that remain. Second, the level of redistribution will decline as a way of getting some wealthy households to stay in the network, reducing the consumption level of the poor households. A well-intentioned policy that succeeded in its objective of increasing migration could thus have unintended distributional consequences that would go unnoticed if researchers evaluating the program and the policy makers who designed it were unaware of the underlying links within the caste. There is an increasing awareness that spillovers must be accounted for when designing and evaluating development programs. In the Indian context, these spillovers will often extend beyond the village to

cover the entire caste, spread over a wide area. This fact seems to have gone unnoticed by development economists evaluating government programs and even explicitly analyzing networks in India.

REFERENCES

Akerlof, George. 1976. "The Economics of Caste and of the Rat Race and Other Woeful Tales." *Quarterly Journal of Economics* 90 (4): 599–617.

Alexander, K. C. 1989. "Caste Mobilization and Class Consciousness: The Emergence of Agrarian Movements in Kerala and Tamil Nadu." In *Dominance and State Power in Modern India: Decline of a Social Order*, Vol. 1, edited by Francine R. Frankel and M. S. A. Rao, 362–413. Oxford: Oxford University Press.

Ambrus, Attila, Arun G. Chandrasekhar, and Matt Elliot. 2014. "Social Investments, Informal Risk Sharing, and Inequality." Economic Research Initiatives at Duke Working Paper 179.

Ambrus, Attila, Wayne Y. Gao, and Pau Milán. 2017. "Informal Risk Sharing with Local Information." http://www.waynegao.com/uploads/8/¼/6/81465138/agm170226.pdf.

Anderson, Siwan. 2011. "Caste as an Impediment to Trade." *American Economic Journal: Applied* 3 (1): 239–63.

Anderson, Siwan, and Patrick Francois. 2017. "Reservations and the Politics of Fear." http://faculty.arts.ubc.ca/fpatrick/documents/reservations_Bread_may17.pdf.

Anderson, Siwan, Patrick Francois, and Ashok Kotwal. 2015. "Clientelism in Indian Villages." *American Economic Review* 105 (6): 1780–816.

Angelucci, Manuela, Giacomo De Giorgi, and Imran Rasul. 2015. "Resource Pooling within Family Networks: Insurance and Investment." http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.722.3002&rep=rep1&type=pdf.

Appadurai, Arjun. 1993. "Number in the Colonial Imagination." In *Orientalism and the Postcolonial Predicament: Perspectives on South Asia*, edited by Carol A. Breckenridge and Peter van der Veer, 314–39. Philadelphia: University of Pennsylvania Press.

Bagde, Surendrakumar, Dennis Epple, and Lowell Taylor. 2016. "Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India." *American Economic Review* 106 (6): 1495–521.

Banerjee, Abhijit, Marianne Bertrand, Saugato Datta, and Sendhil Mullainathan. 2009. "Labor Market Discrimination in Delhi: Evidence from a Field Experiment." *Journal of Comparative Economics* 37 (1): 14–27.

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune. 2013. "Marry for What? Caste and Mate Selection in Modern India." *American Economic Journal: Microeconomics* 5 (2): 33–72.

Banerjee, Abhijit, and Kaivan Munshi. 2004. "How Efficiently Is Capital Allocated? Evidence from the Knitted Garment Industry in Tirupur." *Review of Economic Studies* 71 (1): 19–42.

Banerjee, Abhijit, and Rohini Pande. 2009. "Parochial Politics: Ethnic Preferences and Politician Corruption." https://epod.cid.harvard.edu/sites/default/files/2018-02/parochial_politics-_ethnic_preferences_and_politician_corruption.pdf.

Banerjee, Abhijit, and Rohini Somanathan. 2007. "The Political Economy of Public Goods: Some Evidence from India." *Journal of Development Economics* 82 (2): 287–314.

Banerjee, Biswajit, and J. B. Knight. 1985. "Caste Discrimination in the Indian Urban Labour Market." *Journal of Development Economics* 17 (3): 277–307.

Bardhan, Pranab, and Dilip Mookherjee. 2006. "Pro-poor Targeting and Accountability of Local Governments in West Bengal." *Journal of Development Economics* 79 (2): 303–27.

Bardhan, Pranab, and Dilip Mookherjee. 2012. "Political Clientelism and Capture: Theory and Evidence from West Bengal, India." http://thred.devecon.org/papers/2012/2012-029_Bardhan_Political_Clientilism_and.pdf.

Bardhan, Pranab K., Dilip Mookherjee, and Monica Parra Torrado. 2010. "Impact of Political Reservations in West Bengal Local Governments on Antipoverty Targeting." *Journal of Globalization and Development* 1 (1): 1–36.

Bayly, Susan. 2001. *Caste, Society and Politics in India from the Eighteenth Century to the Modern Age*. Cambridge: Cambridge University Press.

Bertrand, Marianne, Rema Hanna, and Sendhil Mullainathan. 2010. "Affirmative Action in Education: Evidence from Engineering College Admissions." *Journal of Public Economics* 94 (1–2): 16–29.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991–1013.

Besley, Timothy, Rohini Pande, Lupin Rahman, and Vijayendra Rao. 2004. "The Politics of Public Good Provision: Evidence from Indian Local Governments." *Journal of the European Economic Association* 2 (2–3): 416–26.

Besley, Timothy, Rohini Pande, and Vijayendra Rao. 2007. "Political Economy of Panchayats in South India." *Economic and Political Weekly* 42 (8): 661–67.

Besley, Timothy, Rohini Pande, and Vijayendra Rao. 2012. "Just Rewards? Local Politics and Public Resource Allocation in South India." *World Bank Economic Review* 26 (2): 191–216.

Bidner, Chris, and Mukesh Eswaran. 2015. "A Gender-Based Theory of the Origin of the Caste System of India." *Journal of Development Economics* 114: 142–58.

Bisin, Alberto, and Thierry Verdier. 2000. "'Beyond the Melting Pot': Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits." *Quarterly Journal of Economics* 115 (3): 955–88.

Brass, Paul R. 1994. *The Politics of India Since Independence.* Cambridge: Cambridge University Press.

Burnett-Hurst, A. R. 1925. *Labour and Housing in Bombay: A Study in the Economic Conditions of the Wage-Earning Classes of Bombay*. London: P. S. King and Son.

Caldwell, John C., P. H. Reddy, and Pat Caldwell. 1986. "Periodic High Risk as a Cause of Fertility Decline in a Changing Rural Environment: Survival Strategies in the 1980–1983 South Indian Drought." *Economic Development and Cultural Change* 34 (4): 677–701.

Cassan, Guilhem. 2015. "Identity-Based Policies and Identity Manipulation: Evidence from Colonial Punjab." *American Economic Journal: Economic Policy* 7 (4): 103–31.

Cassan, Guilhem. 2019. "Affirmative Action, Education and Gender: Evidence from India." *Journal of Development Economics* 136: 51–70.

Cassan, Guilhem, and Lore Vandewalle. 2017. "Identities and Public Policies: Unintended Effects of Political Reservations for Women in India." Graduate Institute Geneva Working Paper HEIDWP18-2017.

Chakravarti, Uma. 1993. "Conceptualising Brahmanical Patriarchy in Early India: Gender, Caste, Class and State." *Economic and Political Weekly* 28 (14): 579–85.

Chandavarkar, Rajnarayan. 1994. *The Origins of Industrial Capitalism in India: Business Strategies and the Working Classes in Bombay, 1900–1940*. Cambridge: Cambridge University Press.

Chattopadhyay, Raghabendra, and Esther Duflo. 2004. "Women as Policy Makers: Evidence from a Randomized Policy Experiment in India." *Econometrica* 72 (5): 1409–43.

Chhibber, Pradeep K. 2002. *Democracy without Associations: Transformation of the Party System and Social Cleavages in India*. Ann Arbor: University of Michigan Press.

Chin, Aimee, and Nishith Prakash. 2011. "The Redistributive Effects of Political Reservation for Minorities: Evidence from India." *Journal of Development Economics* 96 (2): 265–77.

Cholia, Rasiklal Popatlal. 1941. *Dock Labourers in Bombay*. New York: Longmans, Green, and Co.

Choy, James P. 2018. "Social Divisions with Endogenous Hierarchy." *Economic Journal* 128 (615): 2711–42.

Coleman, James S. 1988. "Social Capital in the Creation of Human Capital." *American Journal of Sociology* 94: S95–S120.

Damodaran, Harish. 2008. *India's New Capitalists: Caste, Business, and Industry in Modern India*. New York: Palgrave Macmillan.

Dandekar, Hemalata C. 1986. *Men to Bombay, Women at Home: Urban Influence on Sugao Village, Deccan Maharashtra, India, 1943–1982*. Ann Arbor: University of Michigan Press.

Darity, William, Jr., Ashwini Deshpande, and Thomas Weisskopf. 2011. "Who Is Eligible? Should Affirmative Action be Group- or Class-Based?" *American Journal of Economics and Sociology* 70 (1): 238–68.

Das Gupta, Ranajit. 1981. "Structure of the Labour Market in Colonial India." *Economic and Political Weekly* 16 (44–46): 1781–806.

De Haan, Arjan. 1997. "Unsettled Settlers: Migrant Workers and Industrial Capitalism in Calcutta." *Modern Asian Studies* 31 (4): 919–49.

de Zwart, Frank. 2000. "The Logic of Affirmative Action: Caste, Class and Quotas in India." *Acta Sociologica* 43 (3): 235–49.

Deshpande, Ashwini, and Katherine Newman. 2007. "Where the Path Leads: The Role of Caste in Post-University Employment Expectations." *Economic and Political Weekly* 42 (41): 4133–40.

Deshpande, Ashwini, and Thomas E. Weisskopf. 2011. "Does Affirmative Action Affect Productivity in the Indian Railways?" https://pdfs.semanticscholar.org/0358/870b8a73d388a6547bf6ae36e3ef01a2c692.pdf.

Dhillon, Amrita, Vegard Iversen, and Gaute Torsvik. 2013. "Employee Referral, Social Proximity and Worker Discipline: Theory and Evidence from India." Center for Economic Studies ifo Institute Working Paper 4309.

Dirks, Nicholas B. 2001. *Castes of Mind: Colonialism and the Making of Modern India*. Princeton: Princeton University Press.

Drèze, Jean, and Geeta Gandhi Kingdon. 2001. "School Participation in Rural India." *Review of Development Economics* 5 (1): 1–24.

Dumont, Louis. 1970. *Homo Hierarchicus: The Caste System and Its Implications*. Chicago: University of Chicago Press.

Dunning, Thad, and Janhavi Nilekani. 2013. "Ethnic Quotas and Political Mobilization: Caste, Parties, and Distribution in Indian Village Councils." *American Political Science Review* 107 (1): 35–56.

Engelshoven, Mirandu. 2002. "Rural to Urban Migration and the Significance of Caste." In *Development*

and Deprivation in Gujarat, edited by Ghanshyam Shah, Mario Rutten, and Hein Streefkerk. Thousand Oaks: Sage Publications.

Fafchamps, Marcel, and Susan Lund. 2003. "Risk-Sharing Networks in Rural Philippines." *Journal of Development Economics* 71 (2): 261–87.

Ferraz, Claudio, and Frederico Finan. 2011. "Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance." https://eml.berkeley.edu/~ffinan/Finan_MPoliticians.pdf.

Field, Erica, Seema Jayachandran, and Rohini Pande. 2010. "Do Traditional Institutions Constrain Female Entrepreneurship? A Field Experiment on Business Training in India." *American Economic Review* 100 (2): 125–29.

Frisancho Robles, Verónica C., and Kala Krishna. 2012. "Affirmative Action in Higher Education in India: Targeting, Catch up, and Mismatch." NBER Working Paper 17727.

Gadgil, D. R. 1959. *Origins of the Modern Indian Business Class*. New York: Institute of Pacific Relations.

Gans, Herbert J. 1962. *The Urban Villagers: Group and Class in the Life of Italian-Americans*. New York: Free Press of Glencoe.

Ghurye, Govind S. 1932. *Caste and Race in India*. London: Routledge.

Gille, Véronique. 2016. "The Role of Local Politics and Caste Networks in the Take-up of Affirmative Action in India." Unpublished.

Gokhale, R. G. 1957. *The Bombay Cotton Mill Worker*. Bombay Millowners' Association.

Gore, M. S. 1970. *Immigrants and Neighbourhoods: Two Aspects of Life in a Metropolitan City*. Bombay: Tata Institute of Social Sciences.

Greif, Avner. 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition." *American Economic Review* 83 (3): 525–48.

Grimard, Franque. 1997. "Household Consumption Smoothing through Ethnic Ties: Evidence from Cote d'Ivoire." *Journal of Development Economics* 53 (2): 319–422.

Gupta, Bishnupriya. 2014. "Discrimination or Social Networks? Industrial Investment in Colonial India." *Journal of Economic History* 74 (1): 141–68.

Gupta, Dipankar. 2000. *Interrogating Caste: Understanding Hierarchy and Difference in Indian Society*. New Delhi: Penguin Books.

Hanna, Rema N., and Leigh L. Linden. 2012. "Discrimination in Grading." *American Economic Journal: Economic Policy* 4 (4): 146–68.

Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh Paul. 2012. "Castes and Labor Mobility." *American Economic Journal: Applied Economics* 4 (2): 274–307.

Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh B. Paul. 2013. "Breaking the Caste Barrier: Intergenerational Mobility in India." *Journal of Human Resources* 48 (2): 435–73.

Hoff, Karla, and Priyanka Pandey. 2006. "Discrimination, Social Identity, and Durable Inequalities." *American Economic Review* 96 (2): 206–11.

Jodhka, Surinder S., and Katherine Newman. 2007. "In the Name of Globalisation: Meritocracy, Productivity and the Hidden Language of Caste." *Economic and Political Weekly* 42 (41): 4125–32.

Kapadia, Karin. 1995. *Siva and Her Sisters: Gender, Caste, and Class in Rural South India*. Boulder: Westview Press.

Kijima, Yoko. 2006. "Caste and Tribe Inequality: Evidence from India, 1983–1999." *Economic Development and Cultural Change* 54 (2): 369–404.

Kingdon, Geeta Gandhi. 1998. "Education of Females in India: Determinants and Consequences: A Case Study of Urban Uttar Pradesh." Unpublished.

Kochar, Anjini. 2008. "Can Schooling Policies Affect Schooling Inequality? An Empirical Evaluation of School Location Policies in India." *India Policy Forum* 4 (1): 53–99.

Kochar, Anjini. 2010. "Affirmative Action through Quotas: The Effect on Learning in India." Stanford Center for International Development Working Paper 426.

Kohli, Atul. 1990. *Democracy and Discontent: India's Growing Crisis of Governability*. Cambridge: Cambridge University Press.

Kooiman, Dick. 1989. *Conversion and Social Equality in India: The London Missionary Society in South Travancore in the 19th Century*. Delhi: Manohar Press.

Kornblum, William. 1974. *Blue Collar Community*. Chicago and London: University of Chicago Press.

Kothari, Rajni. 1970. *Politics in India*. Boston: Little, Brown, and Company.

Lamb, Helen B. 1955. "The Indian Business Communities and the Evolution of an Industrialist Class." *Pacific Affairs* 28 (2): 101–16.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69 (1): 209–44.

Luke, Nancy, and Kaivan Munshi. 2011. "Women as Agents of Change: Female Income and Mobility in India." *Journal of Development Economics* 94 (1): 1–17.

Madheswaran, S., and Paul Attewell. 2007. "Caste Discrimination in the Indian Urban Labour Market: Evidence from the National Sample Survey." *Economic and Political Weekly* 42 (41): 4146–53.

Mazzocco, Maurizio, and Shiv Saini. 2012. "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences." *American Economic Review* 102 (1): 428–68.

Mehta, Makrand. 2001. "Gujarati Business Communities in East African Diaspora: Major Historical Trends." *Economic and Political Weekly* 36 (20): 1738–47.

Min, Pyong Gap. 2001. "Changes in Korean Immigrants' Gender Role and Social Status, and Their Marital Conflicts." *Sociological Forum* 16 (2): 301–20.

Mobarak, A. Mushfiq, and Mark Rosenzweig. 2012. "Selling Formal Insurance to the Informally Insured." https://web.stanford.edu/group/SITE/archive/SITE_2012/2012_segment_½012_SITE_Segment_1_papers/rosenzweig.pdf.

Moorjani, Priya, et al. 2013. "Genetic Evidence for Recent Population Mixture in India." *American Journal of Human Genetics* 93 (3): 422–38.

Morris, Morris David. 1965. *The Emergence of an Industrial Labor Force in India: A Study of the Bombay Cotton Mills, 1854–1947*. Berkeley: University of California Press.

Morten, Melanie. 2019. "Temporary Migration and Endogenous Risk Sharing in Village India." *Journal of Political Economy* 127 (1): 1–46.

Munshi, Kaivan. 2011. "Strength in Numbers: Networks as a Solution to Occupational Traps." *Review of Economic Studies* 78 (3): 1069–101.

Munshi, Kaivan, and Mark Rosenzweig. 2006. "Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy." *American Economic Review* 96 (4): 1225–52.

Munshi, Kaivan, and Mark Rosenzweig. 2015. "Insiders and Outsiders: Local Ethnic Politics and Public Goods Provision." NBER Working Paper 21720.

Munshi, Kaivan, and Mark Rosenzweig. 2016. "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap." *American Economic Review* 106 (1): 46–98.

Munshi, Kaivan, and Mark Rosenzweig. 2018. "Ethnic Diversity and the Under-supply of Local Public Goods." https://www.histecon.magd.cam.ac.uk/km/panchayat79.pdf.

Nafziger, E. Wayne. 1978. "Class, Caste, and Entrepreneurship: A Study of Indian Industrialists." Honolulu: University Press of Hawaii.

Padró i Miquel, Gerard. 2007. "The Control of Politicians in Divided Societies: The Politics of Fear." *Review of Economic Studies* 74 (4): 1259–74.

Pande, Rohini. 2003. "Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India." *American Economic Review* 93 (4): 1132–51.

Rawal, Shenila, and Geeta Kingdon. 2010. "Akin to My Teacher: Does Caste, Religious, or Gender Distance between Student and Teacher Matter? Some Evidence from India." UCL Institute of Education Department of Quantitative Social Science Working Paper 10-18.

Rees, Albert. 1966. "Information Networks in Labor Markets." *American Economic Review* 56 (1–2): 559–66.

Richman, Barak D. 2006. "How Community Institutions Create Economic Advantage: Jewish Diamond Merchants in New York." *Law and Social Inquiry* 31 (2): 383–420.

Rudner, David West. 1994. *Caste and Capitalism in Colonial India: The Nattukottai Chettiars*. Berkeley: University of California Press.

Rudolph, Lloyd I., and Susanne Hoeber Rudolph. 1960. "The Political Role of India's Caste Associations." *Pacific Affairs* 33 (1): 5–22.

Schuler, Sidney Ruth, Syed M. Hashemi, Ann P. Riley, and Shireen Akhter. 1996. "Credit Programs, Patriarchy and Men's Violence against Women in Rural Bangladesh." *Social Science and Medicine* 43 (12): 1729–42.

Singh, K. S., ed. 1992. *An Introduction*. Vol. 36, *People of India*. Calcutta: Anthropological Survey of India.

Srinivas, Mysore N. 1951. "The Social Structure of a Mysore Village." *Economic Weekly*. October 30, 1051–56.

Stokes, Susan C. 2009. "Political Clientelism." In *Oxford Handbook of Comparative Politics*, edited by Carles Boix and Susan C. Stokes. Oxford: Oxford University Press.

Swaminathan, Padmini, and J. Jeyaranjan. 1994. "The Knitwear Cluster in Tiruppur: An Indian Industrial District in the Making?" Unpublished.

Tabellini, Guido. 2008. "The Scope of Cooperation: Values and Incentives." *Quarterly Journal of Economics* 123 (3): 905–50.

Thorat, Sukhadeo, and Paul Attewell. 2007. "The Legacy of Social Exclusion: A Correspondence Study of Job Discrimination in India." *Economic and Political Weekly* 42 (41): 4141–45.

Timberg, Thomas A. 1978. *The Marwaris: From Traders to Industrialists*. New Delhi: Vikas.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62 (3): 171–84.

Weber, Max. 1958. *The Protestant Ethic and the Spirit of Capitalism*. New York: Charles Scribner's Sons.

Yadav, Yogendra. 1999. "Electoral Politics in the Time of Change: India's Third Electoral System, 1989–99." *Economic and Political Weekly* 34 (34–35): 2393–99.