Review of Economic Studies (2011) **78**, 1069–1101 doi: 10.1093/restud/rdq029

Advance access publication 4 March 2011

Strength in Numbers: Networks as a Solution to Occupational Traps

KAIVAN MUNSHI

Brown University and NBER

First version received September 2009*; final version accepted August* 2010 (*Eds.*)

The "new classical" theory states that families in low-skill occupations with low levels of human capital can stay poor from one generation to the next, while families in high-skill occupations with correspondingly high levels of human capital stay wealthy, despite being endowed with the same level of ability on average. This paper proposes an informal institutional mechanism—the community-based network—through which families belonging to the same neighbourhood or kinship group can bootstrap their way out of such low-skill occupational traps. The insight from the dynamic model that is developed is that once they form, new networks providing mutual support to their members and substituting for inherited parental human capital and wealth will strengthen most rapidly in historically disadvantaged communities, generating a correspondingly high level of intergenerational mobility. These predictions are successfully tested using unique data from India. The analysis in this paper, coupled with an emerging empirical literature on networks and migration, provides a new perspective on mobility in developing countries, with restrictive traditional networks decaying even as new networks supporting collective mobility form and strengthen over time.

Key words: Intergenerational mobility, Persistent inequality, Occupational traps, Network formation, Network dynamics, Entrepreneurship, Business networks

JEL Codes: D85, J62, L14, L22

1. INTRODUCTION

The relationship between inequality and growth has received much attention in recent years. The starting point for the "new classical" theory is the assumptions that either credit markets are imperfect and there are fixed costs to investing in human capital or that inherited (parental) human capital is occupation-specific. Under these assumptions, families in low-skill occupations with low levels of human capital stay poor from one generation to the next, while families in high-skill occupations with correspondingly high levels of human capital stay wealthy, despite being endowed with the same level of ability on average (Galor and Zeira, 1993; Banerjee and Newman, 1993; Maoz and Moav, 1999; Mookherjee and Ray, 2003). Persistent inequality of this sort, with its accompanying mismatch between individual ability and occupational choice, has negative consequences for growth.

Complementing the new classical theory, there have been a number of attempts to enrich the model of educational (and occupational) choice described above by introducing community effects. The general insight from this parallel literature is that if there are complementarities in human capital investments due to peer effects or wealth effects or alternatively if there are

complementarities in labour market outcomes due to discrimination, then families will stratify into homogeneous communities. This stratification will only serve to accentuate the family-level persistence in inequality implied by new classical models (Loury, 1977; Bénabou, 1996; Fernandez and Rogerson, 1996; Durlauf, 1996; Bowles, Loury and Sethi, 2009).

Focusing instead on the positive role that can be played by the community, this paper proposes an informal institutional mechanism—the community-based network—through which families belonging to the same neighbourhood or kinship group can bootstrap their way out of low-skill occupational traps. The model of occupational choice developed in Section 2 of the paper is set-up to initially deliver the familiar result that families or dynasties get locked into specific occupations, even though intrinsic individual ability and, hence, the match with other occupations might vary across generations. Introducing a new relatively remunerative occupation does little to help historically disadvantaged families since only families engaged in related occupations will have the resources to shift into it. Prospects for mobility increase, however, once we allow community-based networks to form in the new occupation. Community members from the same generation provide mutual support to each other, effectively substituting for inherited occupation-specific human capital or wealth and allowing individuals without an advantageous family background to move into the new occupation. The key insight from the dynamic model is that once they form, *new* networks should strengthen most rapidly in communities with the *weakest* outside options, and that intergenerational occupational mobility should be correspondingly greater in those communities.

The empirical analysis tests these predictions with unique data from the Indian diamond industry. India does not produce rough diamonds. The rough diamonds are imported, for the most part from Antwerp, then cut and polished in domestic factories, before being sold in the Mumbai (formerly Bombay) market to foreign buyers or shipped directly abroad.[1] Two distinct supply shocks serve as the point of departure for the empirical analysis: the first shock, in the mid-1960's, allowed Indian firms to use scarce foreign exchange to import rough diamonds. As described in Section 3, two traditional business communities—the *Marwaris* and the *Palanpuris*—dominated the business end of the industry from that point onwards, leaving the cutting and polishing to a community of lower-caste agricultural labourers known as the *Kathiawaris*. The second shock, in the late 1970's, coincided with a huge increase in the world supply of rough diamonds with the opening of Australia's Argyle mines. This allowed Kathiawari firms to enter the business, and today all three communities account for a substantial share of the industry.

The episodes that triggered the entry of different communities into the industry are distant enough to trace the growth of their underlying networks together with changes in the family background of entering firms over a thirty-year period, yet are recent enough to provide data that are amenable to statistical analysis. The empirical analysis in this paper uses two independent sources of data: (i) a comprehensive database covering all active firms over the 1995–2004 period, with their export performance in each year, maintained by an administrative agency, the Gem and Jewelry Export Promotion Council (GJEPC), and (ii) a survey of nearly 800 diamond export firms, with offices in the Mumbai market, conducted in 2004–2005. The survey collected information on the senior partner's personal and family background, the firm's history, and changes in the organization of its business over time.

1. The diamond industry accounts for roughly 14% of India's total merchandize exports and has competed with textiles, and more recently with computer software, as the country's top export industry over the past three decades. It is estimated that approximately 1000 Indian diamond export firms employ over a million workers and that this industry accounts for as much as 85% (65% by value) of the rough diamonds cut and polished worldwide (Gem and Jewelry Export Promotion Council, 1998; Purani, 2000).

Because of the difficulty in enforcing legal contracts, the diamond industry is associated with a high degree of community networking throughout the world. Ultra-orthodox Jews historically dominated the Antwerp market and continue to dominate the New York market (Coleman, 1988; Richman, 2006). Over 96% of the firms in our survey are drawn from just three communities, suggesting that community networks must be active here as well. Taking advantage of the co-existence of communities with very different histories in the diamond industry, our objective is to compare the growth of these networks, which substitute for a family background in business and the wealth, connections, and skills that such a background provides, with accompanying change in the composition of entrants. Based on information from the survey, we see in Section 4 that while there is a mild weakening in the inherited business background of the Marwaris and Palanpuris over time, there is a particularly steep decline in the background of the entering Kathiawaris from the late 1970's onwards. Although 70% of the Kathiawaris who entered the industry in 1970, before the second supply shock, reported that their father was a businessman, this statistic declines steadily and drops below 20% by 2000.

The explanation provided by the model for the relatively high intergenerational mobility observed in the Kathiawari community is that its rapidly strengthening industry-specific network was able to support increasingly disadvantaged entrants over time. The marriage institution is key to the maintenance of a strong network. The basic marriage rule in Hindu society is that no individual can match outside the subcaste or *jati*, which typically has a population of a few hundred thousand. The dense web of marriage ties that consequently forms over the course of many generations improves information flows and reduces commitment problems and, not surprisingly, networks serving different functions have historically been organized at the level of the subcaste. Given the particularly severe commitment problems associated with (risky) business activity, marriage alliances within specific industries are, in addition, commonly observed in business communities (Hazlehurst, 1966). In our sample communities, 35% of the entrepreneurs and 57% of their children married within their subcaste and within the diamond industry.

Providing direct support for the hypothesis that the Kathiawari network strengthened relatively rapidly, the frequency of intra-industry marriages increases steeply in that community over time. Almost none of the early Kathiawari entrants who established their firms before 1975 married within the industry. By 2004, however, 50% of the entrants were marrying within the industry, surpassing the corresponding marriage rates for the Marwaris and Palanpuris, which remained roughly constant over time.[2] Firm-level export data, available over a 10-year period, indicate that these investments in the network translated into superior firm performance. The Kathiawaris keep pace with their more established rivals despite the decline in the inherited business background of entrants from this community. Indeed, once this compositional change is accounted for with firm fixed effects, the Kathiawari export trajectory is significantly steeper than the corresponding trajectory for the Marwaris and Palanpuris, precisely as predicted by the model.

Previous theoretical work on networks has taken the position that this institution can restrict mobility during the process of development (Greif, 1994; Kranton, 1996; Rauch, 2001). Consistent with this view, Munshi and Rosenzweig (2006, 2009) show that pre-existing subcaste networks restrict the occupational and spatial mobility of their members. At the same time, the important role played by community networks in the migration process and in supporting mobility has also been extensively documented. In India, the setting for this paper, community networks supported urban migration and the movement of entire subcastes into the new occupations

2. We will also see that Kathiawari firms are more likely to organize production in ways that leave them more reliant on their network and that these community differences in organizational structure have widened over time.

that became available under colonial rule in the 19th and early 20th centuries (Chandavarkar, 1994; Rudner, 1994). More recently, a new business class that is active across diverse industries has emerged in post-colonial India, drawn from a select group of agricultural subcastes and from subcastes that had historically dominated the bureaucracy and various white-collar professions (Damodaran, 2008).

The Indian experience is in many respects similar to the historical and contemporary U.S. experience. For example, Gordon, Edwards and Reich (1982) describe how migrant European communities clustered in particular occupations and locations in the 19th century. While British coal miners may have been imported to work in American coal mines and German bakers and confectioners may have maintained their traditional occupations, their analysis indicates that the arriving migrants typically found niches in new occupations. These patterns of spatial and occupational mobility, supported by underlying networks, continue to this day, as evidenced by the rapidly growing literature connecting community networks and international migration to the U.S. (*e.g.* Kotkin, 1992; Fairlie and Meyer , 1996; Munshi, 2003; Patel and Vella, 2007; McKenzie and Rapoport, forthcoming).

The analysis in this paper, coupled with the emerging empirical literature on networks and migration, provides a new perspective on mobility in developing countries, with restrictive traditional networks decaying even as new networks supporting collective mobility form and strengthen over time. Although these new networks may facilitate mobility, the important caveat is that such mobility will be limited to those communities in which networks have formed. Whether collective mobility occurs or not will depend on whether specific pre-conditions (derived in our analysis) that allow new networks to form are satisfied. Section 5 concludes by identifying policies that could be implemented to encourage such mobility in other settings.

2. THE MODEL

The model developed in this section is set-up to initially deliver the familiar result that occupational traps emerge when networks are absent. Community networks open up the possibility of intergenerational mobility and the main result of the model is that these networks will strengthen most rapidly in communities with the weakest outside options. The model makes a number of simplifying assumptions for analytical convenience. Simulations reported in the appendix verify that the results are robust to some of these assumptions. Other assumptions will be dealt with in Section 4.

2.1. *Individual endowments*

Each individual i belonging to community j is endowed with ability ω_i^j. There are two types of occupations in this economy: business (B) and non-business (NB). The returns to ability vary by the type of occupation. Let the return to ability in the NB occupation be r_{NB} and let the corresponding return in the B occupation be $r_{\mathrm{B}} > r_{\mathrm{NB}}$. We could alternatively classify occupations as skilled and unskilled without changing any of the analysis that follows. The advantage of the context-specific business vs. non-business distinction is that it will allow us to assume that the B occupation and the new D occupation, introduced below, require interchangeable skills.

In addition to his intrinsic ability, each individual inherits an occupation-specific human capital endowment from his father, as in Galor and Tsiddon (1997) or Hassler and Rodríguez Mora (2000). The son of a businessman receives a utility pay-off U_{B}, in addition to the returns to his ability, if he chooses the same occupation as his father, and no inheritance if he chooses any other occupation. Similarly, a father in the NB occupation passes on $U_{\mathrm{NB}}^j < U_{\mathrm{B}}$ if his son remains in the same occupation, and nothing otherwise. The human capital endowment

U^j_{NB} reflects the historical opportunities that were available to communities and is the only dimension along which communities differ in the basic version of the model.[3]

The assumption that the human capital endowment is occupation-specific is standard in the new classical literature and is based on the idea that success in any occupation is associated with a specific set of actions or skills. For example, the son of a businessman in a developing economy will learn how to make connections to buyers and sellers, how to bribe government officials, and more generally to exploit arbitrage opportunities from his father. Such complex and multidimensional knowledge is non-tradeable and can only imbibed over a long period of time in the parental household (Hassler and Rodríguez Mora make an argument along the same lines).

2.2. *Occupational choice without networks*

Individual i belonging to community j will choose a career on the basis of the utility he receives from the alternative occupations available to him. If his father was in the B occupation, he will certainly choose this occupation since

$$U_{\mathrm{B}} + r_{\mathrm{B}}\omega^j_i > r_{\mathrm{NB}}\omega^j_i.$$

If the individual's father was in the NB occupation instead, he will continue to choose the parental occupation despite the fact that $r_{\mathrm{B}} > r_{\mathrm{NB}}$ if the human capital inheritance U^j_{NB} is sufficiently large

$$U^j_{\mathrm{NB}} + r_{\mathrm{NB}}\omega^j_i > r_{\mathrm{B}}\omega^j_i.$$

As in Galor and Tsiddon, we make the distributional assumption $\omega^j_i \sim U[0,1]$, regardless of the occupation of the father (B, NB) or the community that the individual is born into. Moreover, ability is uncorrelated across generations. It is then easy to verify from the expression above that the children of fathers in the NB occupation will always choose that occupation, regardless of their ability if

Condition 1. $U^j_{\mathrm{NB}} \geq r_{\mathrm{B}} - r_{\mathrm{NB}}$.

The occupational persistence derived above, with its associated inefficiency, would continue to be obtained even if human capital were not occupation-specific as long as (the now generalized) human capital requirements vary across a discrete set of occupations. With missing or incomplete credit markets, capable individuals would be unable to invest in the higher levels of human capital needed to move out of the parental occupation into more remunerative activities with greater returns to ability (Galor and Zeira, 1993; Maoz and Moav, 1999; Mookherjee and Ray, 2003). Introducing a new D occupation, that is perfectly substitutable with the B occupation and provides a higher utility pay-off for the children of businessmen $U_D > U_{\mathrm{B}}$ will have little impact on inequality in this case. All the children of businessmen would switch to the D occupation, which in this context refers to business in the diamond industry, but differences in utility within and between communities would remain qualitatively unchanged. An intervention that exogenously shifts a small group of individuals without a business background into the D

3. U^j_{NB} could alternatively measure the pay-off that individuals receive from their community's network in its traditional occupation. Although U^j_{NB} stays constant in the model, we could allow this parameter to decline over time as a new network formed in the D occupation and the traditional network decayed, without changing the results that follow. The main result of the model is that new networks strengthen most rapidly in communities with the weakest outside options (lowest U^j_{NB}). The correspondingly steep decline in U^j_{NB} in those communities as their traditional networks decayed would only reinforce this result.

occupation would reduce inequality, but without the support of a new network as described below, there would be no multiplier and the benefits would be restricted to individuals that receive the intervention.

2.3. *Occupational choice with networks*

Prospects for intergenerational mobility out of the NB occupation and a better match between ability and occupational choice improve once we allow networks to form in the new D occupation. These networks will substitute for the parental endowment U_D that the children of businessmen receive, allowing capable individuals without a business background to enter the new occupation where returns to ability are relatively high. In a business-related occupation, the network will provide connections to buyers and sellers, as well as credit, all of which would be provided by a businessman to his son. In other occupations, the network provides information and strategies for career advancement that would otherwise be provided by the father.

The son of a businessman can either operate independently in the D occupation, taking advantage of the human capital he inherits from his father or participate in his community network (in which case, he foregoes his inheritance). Support for this assumption in the context of the diamond industry will be provided in Section 3. Given this assumption, we will see below that it is optimal for the sons of businessmen to operate independently. The children of individuals in the NB occupation, in contrast, must participate in the network if they are to choose the D occupation. We will focus on the choice between the NB and the D occupation for those children, in the first generation that has access to the new occupation, in the discussion that follows.

Let each individual work for M periods and allow for $N < M$ cohorts in each generation. A fresh cohort enters the workforce in each period and each member of that cohort must make an irreversible career decision at that point in time. For the child of an individual in the NB occupation, this choice will depend on the utility from the NB occupation, as described above, and the utility from the D occupation, which will change across cohorts as the community network strengthens over time. An individual who enters the workforce in period t and chooses to participate in the network will receive referrals (or support) from all members of his community who preceded him and chose to participate in the network. Each of those benefactors provides the individual with a referral when he enters occupation D, which subsequently increases his utility by a constant amount with constant probability in each period of his working life. Let the measure of individuals in each community cohort be equal to one and let a fraction λ of each cohort consist of children of businessmen. We will see momentarily that $(1-\lambda)\Delta\omega_\tau^j$ is the measure of individuals from community j who entered the workforce in period τ and choose to participate in the network. The utility from occupation D for individual i from community j who must make his career choice in period t can then be described by the expression

$$h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j + r_B\omega_i^j,$$

where $W_t^j \equiv (1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j$ measures the stock of individuals in his network or its strength, h maps that stock into the utility that the individual derives from the network over his working life and r_B is the return to ability in the D occupation (and the B occupation). The additive separability in the preceding specification implies that individual ability can substitute for the network. This will be reasonable in many contexts since the services that the network provides, such as information and connections, can be obtained by individuals operating independently if

they are capable. The assumption that the individual only benefits from preceding cohorts and the assumption that network size maps linearly into the individual's utility are made for analytical convenience. These assumptions are relaxed in the simulations reported in the appendix without qualitatively changing the key result of the model. We also implicitly assume that the individual can credibly commit to supporting other members of his network and to repaying his obligations. A precise characterization of the role played by the network and the mechanism through which cooperative behaviour is supported in the diamond industry is postponed to Section 3. That discussion will generate additional context-specific predictions, which we test in Section 4.

Given the utility derived from different occupations, the son of an individual in the NB occupation who enters the workforce in period t will choose the D occupation if

$$h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j + r_B\omega_i^j \geq U_{NB}^j + r_{NB}\omega_i^j.$$

Since $r_B > r_{NB}$, there exists an ability threshold $\underline{\omega}_t^j$ above which all members of community j without a business background who enter the workforce in period t enter the D occupation, where

$$\underline{\omega}_t^j = \frac{U_{NB}^j - h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j}{r_B - r_{NB}}.$$

Individuals with $\omega_i^j \in [0, \underline{\omega}_t^j)$ remain in the traditional NB occupation, while individuals with $\omega_i^j \in [\underline{\omega}_t^j, 1]$ move into the D occupation.

2.4. *Network dynamics*

Given that ability is distributed uniformly over the unit interval, the range of ability over which individuals enter the network in period t is $\Delta\omega_t^j \equiv 1 - \underline{\omega}_t^j$. Substituting the expression for $\underline{\omega}_t^j$ derived above,

$$\Delta\omega_t^j = \frac{h(1-\lambda)}{r_B - r_{NB}}\Sigma_{\tau=0}^{t-1}\Delta\omega_\tau^j - \left(\frac{U_{NB}^j}{r_B - r_{NB}} - 1\right) \equiv \beta\Sigma_{\tau=0}^{t-1}\Delta\omega_\tau^j - \alpha^j. \tag{2.1}$$

Note that $\alpha^j \geq 0$ from Condition 1. Since $r_B > r_{NB}$, it follows that $\beta > 0$ and $d\alpha^j/dU_{NB}^j > 0$. We will consequently use α^j interchangeably with U_{NB}^j in the discussion that follows. Equation (2.1) also allows us to verify that the choice by children of businessmen to operate independently, as we have assumed, is indeed optimal: $h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j < U_D$. Since ability is distributed uniformly on the unit interval, $\Delta\omega_t^j \leq 1$. It therefore follows from equation (2.1) that $h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j \leq U_{NB}^j$, which is less than U_B and, hence, U_D (by assumption).

To initiate the network dynamics, a measure $(1-\lambda)\Delta\omega_0$ of individuals without a business background are moved exogenously into the D occupation in Period 0.[4] Starting with Period 1 and moving forward in time, equation (2.1) can be solved recursively to derive a simple expression relating $\Delta\omega_t$ to $\Delta\omega_0$:

$$\Delta\omega_t^j = (\beta\Delta\omega_0 - \alpha^j)(1+\beta)^{t-1}. \tag{2.2}$$

4. If the range of ability over which individuals enter the network in period t is $\Delta\omega_t^j \equiv 1 - \underline{\omega}_t^j$, then $(1-\lambda)\Delta\omega_t^j$ is the measure of individuals that enters the network since individuals with a business background operate independently. It is thus convenient to specify the initial condition as $(1-\lambda)\Delta\omega_0$ in all communities.

Note that the network will only set off on a positive trajectory if there is a sufficiently large influx of initial entrants:

Condition 2. $\Delta\omega_0 \geq \alpha^j/\beta$.

This requirement is a standard feature of "tipping" models and if it is satisfied, the main result of the model follows directly from equation (2.2).

Proposition 1. *If Condition* 2 *is satisfied then (a) the strength of the network increases over time, more steeply in communities with weak outside options (low* U^j_{NB}*), (b) the share of entrants with a family background in business* S^j_t *decreases over time and, moreover, will decrease more steeply in communities with weak outside options when* $S^j_t > 0{\cdot}5$.

From equation (2.1), the strength of the network in period t,

$$W^j_t \equiv (1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega^j_\tau = \frac{(1-\lambda)}{\beta}[\Delta\omega^j_t + \alpha^j].$$

$$\frac{dW^j_t}{dt} = \frac{(1-\lambda)}{\beta}\frac{d\Delta\omega^j_t}{dt}, \quad \frac{d^2W^j_t}{d\alpha^j dt} = \frac{(1-\lambda)}{\beta}\frac{d^2\Delta\omega^j_t}{d\alpha^j dt}.$$

It follows that changes in network strength over time and across communities are driven entirely (up to a positive constant) by changes in $\Delta\omega^j_t$. The first part of the proposition then follows directly from equation (2.2):

$$\frac{d\Delta\omega_t}{dt} = (\beta\Delta\omega_0 - \alpha^j)(1+\beta)^{t-1}\ln(1+\beta), \tag{2.3}$$

$$\frac{d^2\Delta\omega_t}{d\alpha^j dt} = -(1+\beta)^{t-1}\ln(1+\beta). \tag{2.4}$$

If Condition 2 is satisfied, $d\Delta\omega_t/dt > 0$ and $d^2\Delta\omega_t/d\alpha^j dt < 0$. The strength of the network is thus increasing over time, more steeply in communities with weak outside options (low U^j_{NB}). Although the negative relationship between conditions in the origin and the propensity to migrate is well understood, our model of endogenous network formation at the destination generates the stronger prediction that the gap in network strength between communities with different outside options will widen over time. This result is obtained with a linear network technology and simulations reported in the appendix indicate that it holds up when the network technology is concave.

To prove the second part of the proposition, the share of entrants in period t with a business background can be expressed as

$$S^j_t = \frac{\lambda}{\lambda + (1-\lambda)\Delta\omega^j_t}.$$

$$\frac{dS^j_t}{dt} = \frac{-\lambda(1-\lambda)}{[\lambda+(1-\lambda)\Delta\omega^j_t]^2}\frac{d\Delta\omega^j_t}{dt},$$

$$\frac{d^2S^j_t}{d\alpha^j dt} = \frac{2\lambda(1-\lambda)^2}{[\lambda+(1-\lambda)\Delta\omega^j_t]^3}\frac{d\Delta\omega^j_t}{dt}\frac{d\Delta\omega^j_t}{d\alpha^j} - \frac{\lambda(1-\lambda)}{[\lambda+(1-\lambda)\Delta\omega^j_t]^2}\frac{d^2\Delta\omega^j_t}{d\alpha^j dt}.$$

Substituting from equations (2.2–2.4), and noting that $d\Delta\omega_t^j/d\alpha^j = -(1+\beta)^{t-1}$ from equation (2.2),

$$\frac{d^2S_t^j}{d\alpha^j dt} = \frac{\lambda(1-\lambda)(1+\beta)^{t-1}\ln(1+\beta)[\lambda-(1-\lambda)\Delta\omega_t^j]}{[\lambda+(1-\lambda)\Delta\omega_t^j]^3}.$$

Although the share of entrants with a business background is unambiguously declining over time, $dS_t^j/dt < 0$ from equation (2.3), this decline will only be steeper in communities with weak outside options, $d^2S_t^j/d\alpha^j dt > 0$, if the share of entrants with a business background exceeds 0·5: $\lambda > (1-\lambda)\Delta\omega_t^j$. This last condition will be satisfied, at least in the early stages of the transition, if the initial share of entrants in Period 0 satisfies this condition $S_0^j > 0{\cdot}5$. We will verify empirically that the initial entrants are overwhelmingly from business backgrounds in all communities in Section 4.

Networks allow capable individuals without a business background to enter the new D occupation where the returns to ability are greater. While this occupational migration will generally increase efficiency, the first-best allocation of individuals to jobs is unlikely to be obtained. To see why this is the case, consider an economy with a fixed number of positions in each occupation. Returns to ability will adjust across occupations to clear the market and ideally we would like to see the most capable individuals in the economy assigned to the D occupation. Although new networks increase efficiency by expanding the pool of individuals that can potentially enter this occupation, this pool will nevertheless be restricted to those communities in which networks have formed.

2.5. *Extensions to the model*

The advantage of assuming that ability is uniformly distributed, that the network technology is linear, that individuals make their entry decision at a fixed age, and that they only benefit from the cohorts that have preceded them is that $\Delta\omega_t^j$ can be expressed as a linear function of the stock of past participants in the network $(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau$. This expression can then be solved recursively to derive a simple relationship between $\Delta\omega_t^j$ and the exogenous initial condition $(1-\lambda)\Delta\omega_0$, which delivers the main result of the model that the new network strengthens more rapidly in communities with weak outside options (low U_{NB}^j). Simulations reported in the appendix relax each of these assumptions in turn to verify the robustness of the main analytical result. I also show in the appendix that this result holds when communities are allowed to vary on a second dimension: the share of individuals with a business background (λ).

Two extensions to the model are particularly useful for the empirical analysis that follows. The analysis in the appendix indicates that the divergence in network strength over time and across communities continues to be obtained when the individual receives support from his own cohort and the following cohort, not just the cohorts that preceded him. We would expect the divergence across communities to be maintained as we allowed additional cohorts to provide support, ultimately including all cohorts that were active in the industry at a given point in time. What is essentially a statement about changes in network strength across cohorts in Proposition 1 could then be reformulated as a statement about changes in network strength over time. In the basic version of the model, individuals are locked into a cohort-specific network with a fixed level of utility over their working lives. In contrast, each individual's utility will be increasing over time if he benefits from subsequent cohorts as well. This distinction will be important in the analysis of firm performance in Section 4, where we exploit variation in utility (or profits) within firms over time to identify underlying changes in the strength of community networks.

The model assumes that individuals make an irreversible choice between the new D occupation and the traditional occupation when they enter the workforce. Without entry costs, there is

no reason to delay entry into the D occupation strategically. However, an individual belonging to community j who enters the workforce in period t and has ability below $\underline{\omega}_t^j$ could still choose to enter later once the network strengthens and the threshold has dropped sufficiently. To allow for this possibility, I solved the model numerically with individuals below the threshold when their cohort entered the workforce entering occupation D with a lag if their ability exceeded the threshold in the subsequent period. Proposition 1 continues to hold in the simulation (reported in the appendix), although the network strengthens more rapidly over time in all communities when individuals can delay entry and establish their firms at different ages. This extension to the model will allow us to control for age effects in Section 4, while studying changes in the pattern of entry into the diamond industry across communities over time.

3. INSTITUTIONAL SETTING

3.1. *The survey*

Although aggregate diamond statistics are available over many years, detailed firm-level information could only be obtained by conducting a survey of the industry. The population of firms upon which the survey is based is obtained from a computerized database maintained by the GJEPC of all its members. The import of rough diamonds was banned after independence in 1947 to preserve scarce foreign exchange. This policy was reversed in the mid-1960's with the introduction of the Multi-Rate Import Replenishment Scheme, which allowed firms to import rough diamonds against their previous exports of polished diamonds. The GJEPC verified the export figures for its members under this scheme and then forwarded them to the Government of India. Almost all active exporters availed of this useful service, and so the GJEPC database provides us with a comprehensive list of firms that exported polished diamonds each year. I was able to gain access to this database, covering the 1995–2003 period, at the beginning of 2004. The GJEPC database includes the name of the firm, its address and telephone numbers, the name of a contact individual (typically the senior partner), and the firm's export figures, each year from 1995 onwards.[5] The population of firms for the survey was restricted to firms with offices in the Mumbai market, listed in the GJEPC database as exporting in any year over the 2001–2003 period. Seven hundred and seventy seven firms were ultimately interviewed, with an overall response rate of 84·3%.[6]

Diamond firms are notoriously secretive and so every effort was made to establish connections within the industry before the survey commenced. Assisted by a few close personal contacts, I gradually built up a small network of influential diamond exporters over a two-year period, which in turn helped the survey team penetrate each of the communities in the industry. A computerized referral system was set-up, and each individual in my personal network provided a list of exporters that he was tied closely with. We would speak on behalf of these individuals when arranging interviews with the firms; in many cases, this was sufficient for the firm to agree to be interviewed, but in other cases, the firms did contact the individual who had provided the referral to verify its authenticity. These firms, in turn, provided additional referrals, and the process continued until *all* the names on our list had been covered. While the order of interviews may not have been random, the referral-based approach that we took did not result in a non-representative sample since all firms were ultimately contacted. When providing referrals, our respondents were simply asked to list firms that they were closely tied with, without any

5. The export figures were provided in 2005 after the survey was completed and so the export data cover a 10-year period from 1995 to 2004.

6. The response rate across communities was 85·7% for the Kathiawaris, 89·3% for the Marwaris, and 81·9% for the Palanpuris. See Munshi (2007) for further details of the survey.

TABLE 1
Referral pattern

Source of referrals	Number of individuals that provided referrals (1)	Total number of referrals provided (2)	Percent of referrals for Kathiawaris (3)	Percent of referrals for Marwaris (4)	Percent of referrals for Palanpuris (5)
Kathiawari exporters	60	212	74·06	2·83	20·28
Marwari exporters	24	206	12·62	42·72	37·86
Palanpuri exporters	128	707	9·19	9·05	78·64
Brokers	47	239	31·38	14·23	51·05
Other	36	109	18·35	21·10	49·54

Notes: Other sources of referrals include personal connections of the survey team and firms belonging to other communities.
A total of 295 individuals provided referrals in Column 1.
These individuals provided a total of 1473 referrals in Column 2.
Columns 3–5 sum to approximately 95% because some referrals are also made to exporters from other communities.

prompting from our side. It is worth noting that only 5·7% of the sampled firms do not appear in the GJEPC database, consistent with the assumption that this database effectively covers the population of active exporters.

Although it is well known that community networks play an important role in this industry, the survey respondents were generally reluctant to report the support they received from members of their community or from other close connections. The pattern of referrals that was received evidently had research value since it could be used to provide direct evidence on the importance of community ties and so the survey team was instructed to continue to fill the data fields, which recorded the identity of up to five individuals who had provided referrals, even after a firm had been interviewed. Table 1 lists the major sources of referrals, the number of referrals that they provided, and the community-wise breakdown of firms that received these referrals. A total of 295 individuals provided 1473 referrals. Looking across Columns 3–5, it is apparent that exporters from each community disproportionately provide referrals to members of their own group. The Kathiawaris make up just 29% of the firms in the sample, yet 74% of the referrals from Kathiawari exporters are to members of their community. The Marwaris who make up 17% of the sample and the Palanpuris who account for another 54% also favour members of their own community, but not as conspicuously as the Kathiawaris. The Marwaris in particular make a substantial number of cross-community referrals and we will later see that the Marwaris concentrate on the polished side of the market where community affiliation is less important.

3.2. *Exit from the industry*

The surveyed firms are all currently active. Much of the analysis will use the establishment year of these firms to describe historical changes in the industry across communities. However, firms that exited the industry will be missing, which could bias our interpretation of these changes if exit rates varied by community. A particularly useful feature of the administrative data made available by the GJEPC is that they cover all active firms over the 1995–2003 period. I assume that a firm exits in a given year if it was exporting in that year but fails to show up in the database thereafter. It seems reasonable to assume that a firm which fails to show up continuously for three years or longer has permanently exited, allowing exit rates to be computed each year from 1995 to 2000. Restricting attention to firms with offices in the Mumbai market, we see in Table 2 that

TABLE 2
Exit statisiics

Sample	All firms (1)	Kathiawaris (2)	Marwaris (3)	Palanpuris (4)
Exit rates	1·53	1·79	1·14	1·55
Exit rates by age				
One-year old firms	2·43	1·94	1·01	3·57
Two-year old firms	2·73	1·23	1·18	4·72
Three-year old firms	2·73	5·56	3·23	0·96
Four-year old firms	2·05	2·94	0·00	2·70

Notes: Exit rates, in percent, are computed over the 1995–2000 period.
A firm is assumed to have exited if it was active in a given year and then drops out of the database for three years or longer.
A firm is assumed to be one year old in a given year if it first shows up in the database in that year.
Exit rates for one-year olds can thus be computed over the 1996–2000 period, for two-year olds over the 1996–1999 period, for three-year olds over the 1996–1998 period, and for four-year olds over the 1996–1997 period.

annual exit rates over this period are extremely low—just around 1·5%—and do not vary by community.[7] Although not reported, there is no time trend in these statistics.

Diamonds are bought and sold on credit. Faced with unexpected delays in receiving payment from their polished buyers, it is often a challenge for exporters to turn their working capital around on schedule. Based on my conversations with diamond exporters, it appears that firms are most likely to exit the industry when they face a large demand or supply shock and subsequently fall behind on their payments. Although there is evidently a random aspect to these shocks, we might expect firms that have recently begun to export, with smaller cash reserves, to be especially vulnerable. Exit rates by the age of the firm are thus presented next in Table 2. I assume that a firm enters the industry in a given year if it first appears in the database in that year, allowing me to compute exit rates for one- to four-year old firms over the 1996–2000 period.[8] Exit rates remain extremely low even for these young firms and while the small number of observations in each age-community cell generates some noise in these statistics, one-way analysis of variance tests do not reject the hypothesis that exit rates are equal across communities in each age category.

3.3. *The communities*

The history of the industry described earlier would suggest that entrepreneurs from the three communities should come from very different backgrounds. The Palanpuris entered the diamond business in the 1880's and later expanded into the pearl trade, with their business activity extending as far as Antwerp by the 1930's (Chhotalal, 1990). The Marwaris are a community of traders and moneylenders who made the transition into industry around 1914 and subsequently expanded their trading and industrial activity throughout the country (Lamb, 1955). The Marwari

7. The contact names included in the GJEPC database, together with a detailed knowledge of firms in the industry, allowed the exporters that I knew and their employees to assign a community affiliation to each firm in the database with an office in the Mumbai market. Names are a good indicator of community affiliation, and comparing this assignment to the actual affiliation, obtained from the survey, just 6·3% of the sampled firms were miscoded. We would expect similarly low rates of miscoding for firms out of sample as well.

8. Exit rates can be computed for one-year olds over the 1996–2000 period, for two-year olds over the 1996–1999 period, for three-year olds over the 1996–1998 period, and for four-year olds over the 1996–1997 period.

TABLE 3
Characteristics of enterprencurs

Community	Kathiawari (1)	Marwari (2)	Palanpuri (3)
Panel A: individual characteristics			
Age	42·46	46·13	49·05
	(11·37)	(10·32)	(10·48)
Years of schooling	10·84	14·41	12·87
	(3·77)	(2·09)	(2·51)
Percent that grew up in Mumbai	22·02	26·40	49·38
	(41·53)	(44·26)	(50·06)
Panel B: family background			
Father's occupation (%)			
Farming	53·02	2·46	2·54
White-collar professional	5·58	13·93	15·52
Other business/store-owner/sales	11·16	27·05	27·23
Other jewelry business	5·12	29·51	11·96
Diamond cutting & polishing	7·44	1·64	6·62
Diamond broker/trader	2·79	3·28	9·92
Diamond exporter	14·88	22·13	26·21
Any business	34·56	82·40	75·81
	(47·67)	(38·24)	(42·88)
Number of firms	218	125	405

Notes: Standard deviations in parentheses.
Any business includes other business/store-owner/sales, other jewelry business, diamond broker/trader, and diamond exporter.

network is more diversified, both spatially and by business activity, than any other community network in the country (Timberg, 1978). The Kathiawaris, in contrast, are a lower caste of agricultural labourers and sharecroppers who made the transition to industrial labour in the 1960's when the Indian diamond industry started to grow with the introduction of the Multi-Rate Import Replenishment Scheme (Engelshoven, 2002). The cutting and polishing of the small stones that the Indian diamond industry specializes in does not require great skill or ability, and this historically disadvantaged community only moved into business in the late 1970's when a supply shock hit the world diamond industry.

The descriptive statistics in Table 3, Panel A, based on data collected from the senior partner in each firm, are conditional on entry into the diamond industry. They are, nevertheless, broadly consistent with the historical differences across communities. The entrepreneur's age is (mechanically) negatively correlated with the year that the firm was established. Given the long history of the Palanpuris in the diamond business and the late entry of the Kathiawaris into this occupation, it is not surprising that the Kathiawari respondents are younger than the Marwari respondents, who in turn are younger than the Palanpuri respondents in our sample. The Kathiawaris also have significantly lower educational attainment, measured by years of schooling, than the entrepreneurs from the historically advantaged communities. They are also less likely to have grown up in Mumbai (as compared with the Palanpuris). A relatively low proportion of Marwaris also report having grown up in Mumbai, but this simply reflects the wide scope of their commercial activities; although not reported, many of them grew up in urban centers elsewhere in the country and this will become apparent in a moment when I describe the occupations of their fathers.

Table 3, Panel B describes the entrepreneur's father's occupation, which is aggregated into seven categories: farming, white-collar professional, other business, other jewelry business, diamond cutting and polishing, diamond broker or trader, and diamond exporting. The Kathiawaris are significantly less likely to belong to a business family than the other two communities: 35% of the Kathiawaris vs. 82% of the Marwaris, and 76% of the Palanpuris report that their father was engaged in any type of business. Although these patterns are consistent with differential selection across communities into the industry, as predicted by the model, they could also reflect the fact that a greater share of the population has a business background in the Marwari and Palanpuri communities (in the context of the model, this is saying that λ is greater in those communities). The most striking feature of these statistics, however, is the differential access to occupations outside business with 53% of the Kathiawaris, but just over 2% of the Marwaris and Palanpuris, reporting that their fathers were farmers.

Although the Palanpuris, and especially the Marwaris, have many opportunities outside the diamond business, the next best option for a Kathiawari entrepreneur is farming or working as a labour contractor, neither of which is particularly remunerative. In terms of the model, the historical background and the descriptive statistics in Table 3 indicate that the Kathiawaris have the worst options outside business among the three communities. Although the differences between the Palanpuris and the Marwaris are not as large, non-business (and business) activity in the Palanpuri community is largely restricted to the diamond industry, whereas the Marwaris are well diversified and can choose from a wider range of non-business (and business) activities.

3.4. *Organization of production*

Most diamond exporters visit Antwerp once every month or every other month for a few days to acquire rough diamonds, have these diamonds cut and polished in domestic factories, and then sell the polished diamonds on the Mumbai market or directly to foreign buyers. "Much of the diamond industry revolves around the issue of getting a regular supply of good quality [rough] diamonds" (Engelshoven, 1999, p. 371). Rough suppliers in Antwerp and the largest exporters receive parcels directly from the Diamond Trading Corporation (DTC), the trading arm of DeBeers, or from other primary suppliers of rough diamonds. These parcels will typically comprise stones of various grades and sizes. Individual exporters, however, will tend to specialize in stones of a particular size.[9] This implies that they would like to buy from suppliers in Antwerp who happen to be well stocked with the type of stones they specialize in on any given trip, with the set of preferred suppliers changing from one trip to the next. The rough stones are received on credit, without a written contract stipulating the principal, interest rate, and time of repayment, giving rise to a potentially substantial commitment problem.

Based on my conversations with numerous exporters, two solutions are available to avoid this commitment problem. The first solution takes advantage of the community network. Although an exporter could establish long-term bilateral relations with a small number of suppliers in Antwerp, these relations would not be sufficient to satisfy his demand for rough diamonds, given the variation in the type of stones received by suppliers from one month to the next. What the network does is to diversify the supply of rough diamonds, with exporters that have established long-term relations with particular suppliers providing referrals for other members of

9. Diamonds are classified by size and shape. In the questionnaire, we defined eight categories—seven sizes and a separate category for "fancy shapes"—and asked the entrepreneurs to report the proportion of their output (by value) in each category. Despite this fine classification of stones, a substantial fraction of the firm's output is devoted to a single—most popular—category: 52% for the Kathiawaris, 42% for the Marwaris, and 48% for the Palanpuris. The Marwaris are significantly less specialized, in large part because their business is centred on the polished side of the market, where flexibility is less costly.

their community. The set of exporters providing referrals will vary from one period to the next depending on the mix of stones received by the suppliers in Antwerp. Exporters thus draw upon different members of their community to provide referrals over time, expanding the set of suppliers that is available to them. Exporters providing referrals have long-term relationships at stake and so will ensure that members of their community receiving the rough stones do not renege on their credit obligations. Exporters receiving referrals will not cheat, even if they do not expect to be helped by the same exporter in the future, if the threat of community-based sanctions is sufficiently severe. Numerous accounts (as described below) of the serious economic and social punishments faced by exporters who reneged on their obligations, on the few occasions when such transgressions did occur, would tend to support this characterization of the cooperative equilibrium.[10]

Most exporters follow the strategy described above, building long-term relationships with a few suppliers in Antwerp, while using the community network to expand their access to rough diamonds. Other exporters operate independently, by setting up branches in Antwerp. Exporters who are based permanently in Antwerp also function as rough suppliers and so will interact frequently with other suppliers in the Antwerp market. These interactions and their permanent presence in Antwerp allow them to establish a reputation in the market, which serves as a commitment device and gives them access to rough diamonds from numerous suppliers without the support of a community network.[11] Based on the model in Section 2, we would expect the sons of businessmen, who have the wealth, connections, and skill needed to operate independently, to be most likely to establish branches in Antwerp.[12]

Table 4, Column 1 describes transactions on the rough side of the market. Firms have 11 suppliers per year on average and 71% of the firms have a dominant supplier who provides more than 30% of their rough diamonds. Different firms will have different dominant suppliers, allowing for the cross-referrals across firms that are needed for the network to function effectively. Much of the rough supply (70%) comes from Antwerp. The other major alternative source of roughs is the Mumbai secondary market, where the price is substantially higher but the commitment problem less severe since all firms have a permanent presence in the city.[13] Despite the high value of the rough diamonds and the potential for default, much of the rough supply is obtained on credit and rarely involves a written contract.

In contrast with rough diamond transactions, where referrals are critical and firms tend to do business with a limited number of suppliers, the polished side of the industry, described in Column 2, appears to operate very much like a spot market. Firms have as many as 34 buyers per

10. Greif's (1993) description of the Maghribi traders' coalition is very similar to this characterization of the diamond exporters' network. Although deviations from cooperative behaviour are associated with severe sanctions in the diamond industry, increasing the level of commitment that can be sustained, Greif shows formally that the termination of future network services can be sufficient to support a cooperative equilibrium.

11. Once reputation is introduced, the market has many features of the network such as a collective punishment mechanism. The advantage of the community network is that it allows exporters without a permanent presence in the Antwerp market to receive rough diamonds on credit.

12. We assume in the model that the son of a businessman can either operate independently, taking advantage of the human capital inherited from his father, or participate in the network. This assumption is reasonable in the context of the diamond industry. Once a branch has been established in Antwerp, the network's ability to punish the individual is limited, and it follows that he will no longer have access to its services. In addition, the wealth and connections that the individual inherits from his father have limited value if he chooses to participate in the network since much of his rough supply will be received on credit (without collateral) and since he will rely for the most part on his community connections.

13. The very largest firms, known as *sightholders*, receive rough diamonds directly from the DTC. A relatively small number of firms also buy rough diamonds from Israel. However, Antwerp is the dominant source of rough diamonds from abroad.

TABLE 4
Organization of production

Business activity	Buying roughs (1)	Selling polished (2)
Number of suppliers/buyers per year	10·60	34·18
	(16·85)	(76·08)
Percent of firms with a single dominant supplier/buyer	71·18	58·64
	(45·34)	(49·28)
Percent of stones (by value) bought/sold directly abroad	70·26	63·10
	(33·80)	(37·74)
Percent of stones (by value) bought/sold on credit	77·10	82·08
	(29·58)	(25·59)
Average repayment period (days)	101·48	110·33
	(25·54)	(37·02)
Percent of transactions involving a written agreement	5·79	4·89
	(2·34)	(2·16)

Notes: Standard deviations in parentheses.
Dominant supplier is defined as a supplier who provides more than 30% of the firm's roughs.
Dominant buyer is defined as a buyer who accounts for more than 20% of the firm's polished.
Percent of stones bought directly abroad in Column 1 is computed using stones sourced from Antwerp alone.
Merchant exporters, who restrict their activity to the polished side of the market, are excluded from Column 1.

year on average, and fewer (less than 60%) of the firms have a single dominant buyer, despite the fact that a dominant buyer is now defined to account for just 20% of the firm's product. A substantial fraction of the polished diamonds are also sold on the Mumbai market, either to *merchant exporters*, who restrict their activity to buying and selling polished diamonds and brokerage or to visiting foreign buyers.

Polished diamonds are largely sold on credit and these transactions rarely involve a written contract in Column 2, so commitment problems could potentially arise on this side of the market as well, with buyers reneging on their obligations. Because firms specialize in particular stone sizes, they tend to build long-term relationships with a few foreign buyers, channeling the rest of their output abroad through numerous merchant exporters, who buy the product directly or serve as commission agents for foreign buyers with whom they have established long-term relationships. All export firms, including the merchant exporters, have a permanent presence in Mumbai and so can build a personal reputation in the polished market. As with the exporters with branches in Antwerp, this market reputation serves as a commitment device, expanding business relations across community lines. The analysis that follows will consequently assume that the primary role of the community network is to increase the amount of rough diamonds that its members can procure, with larger networks providing more referrals.[14]

3.5. *Network formation*

Recall from Condition 2 in Section 2 that a sufficiently large influx of initial entrants is needed to set the network on a positive trajectory. The entry of the Kathiawaris as a group into business

14. We do not expect community networks to play an active role at the cutting and polishing stage of the production process either. Entrepreneurs can always establish long-term bilateral relationships with their manufacturing contractors to avoid the commitment problems, associated with the swapping of roughs, that arise at this stage. Consistent with this view, the respondents in the survey reported an average relationship of 16 years with their manufacturing contractors.

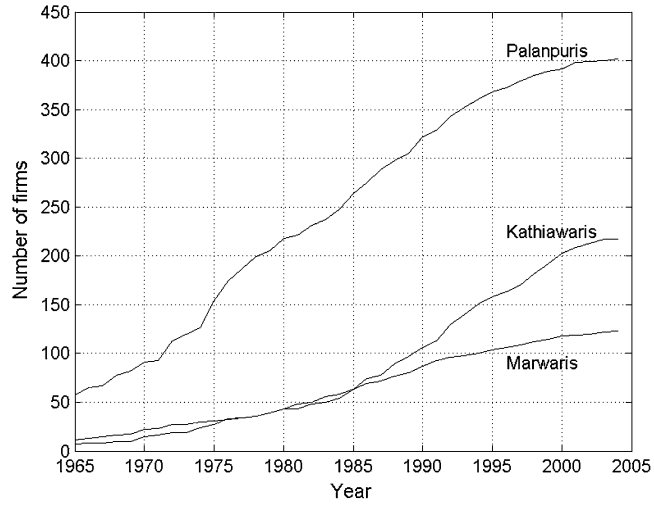

FIGURE 1
Number of Firms

can be traced to the discovery of massive diamond deposits in Australia's Argyle mines in 1979. Although the Argyle mines account for as much as one third of the global production of natural diamonds, these diamonds tend to be small and low quality. The Indian industry with its low labour costs was particularly suited for the cutting and polishing of these diamonds, resulting in the entry of new firms at this time and substantial expansion in the decades that followed.

Figure 1 plots the number of firms by community over the 1965–2004 period, based on the establishment year of the firms in the sample. The Palanpuris are always the dominant group, growing at a constant rate over the entire 40-year period. The Marwaris and the Kathiawaris track together at a slower rate initially, but while the Marwaris continue to grow at that rate, the Kathiawaris shift to a steeper trajectory around 1980. Although this is not visually discernable, if we were to divide the 1965–2004 period into two equal halves, the Kathiawaris would grow significantly faster than the Palanpuris in the latter period.

Contrary to the popular perception in the industry that the Kathiawaris suddenly entered business in the late 1970's, we see that a few Kathiawari firms were active as far back as the 1960's. Our theoretical framework assumes that individuals with a family background in business can always move into a new business-related activity. For hundreds of Kathiwaris to enter as in Figure 1, however, the support of a network is needed since the bulk of that community lacks a business background. The supply shock in 1979 may have jump-started the Kathiawari network, but how did the *first* members of the network succeed in business? The story told in the industry is that the Palanpuris supported the early Kathiawaris.[15] Based on our characterization of the organization of production, such support could only have been provided if the Palanpuris had started to establish branches in Antwerp by that time. An entrepreneur with a branch in Antwerp functions simultaneously as an exporter and a rough diamond supplier. He could thus use his

15. The fact that the early Kathiawari entrants were supported by Palanpuris is not disputed in the diamond industry, although individual entrepreneurs are reluctant to admit that they were personally assisted in this way. Statements such as the following are often heard: "'Kathiawadis are here because of the Palanpuris' admits a Kathiawadi diamond merchant. The Palanpuris, who were the market leaders brought the Kathiawadis into the trade. Help came not only in the form of finance but as initiation into the import-export sector". (Diamond World, 1999, pp. 52–53).

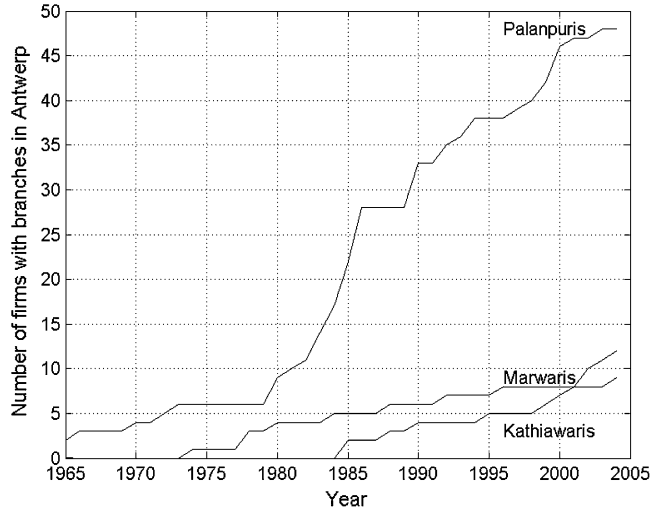

FIGURE 2
Number of Firms with Branches in Antwerp

market reputation to provide referrals to many other suppliers (as opposed to just a few if he was part of the community network) and independently support the new entrant. At the same time, he would have an incentive to bring his trusted Kathiawari labour contractor into business, as a way of unloading his stock of rough diamonds.[16] The survey collected information on the year in which each firm established each of its foreign branches (where relevant) and Figure 2 uses this information to plot the number of firms with branches in Antwerp. As predicted, we see a sharp increase in the number of Palanpuri firms with branches in Antwerp just around 1979. It was this change in the organizational structure of firms belonging to a community that was already established in the industry that allowed the Kathiawaris to enter, emphasizing the fortuitous confluence of circumstances that was needed to initiate the large-scale occupational migration that followed.

While the Palanpuris may have been able and willing to support the entering Kathiawaris in Antwerp, how did they sell their polished diamonds in the Mumbai market? Figure 3 plots the growth in the number of merchant exporters across communities and over time. We see that these firms, drawn predominantly from the Marwari and Palanpuri communities, emerged as early as the mid-1970's, allowing the initial Kathiawari entrants to sell their polished diamonds without established foreign buyers. We saw above that it is optimal for the children of businessmen to operate independently in the new *D* occupation. Figures 2 and 3 are broadly consistent with this claim since firms with branches in Antwerp and merchant exporters operate independently of their networks and the Marwaris and Palanpuris are much more likely to be children of businessmen than Kathiawaris. Formal statistical support for the claim that children of businessmen are less likely to participate in the network, with its corollary that Marwaris and Palanpuris are more likely to operate independently will be provided below.[17]

16. Community ties are not needed to maintain long-term bilateral relations. Many Palanpuri businessmen would have established such long-term relations with their Kathiawari contractors by the late 1970s.

17. Although few Marwari firms establish branches in Antwerp in Figure 2, entrepreneurs from this community are disproportionately represented among the merchant exporters who also operate independently of their networks.

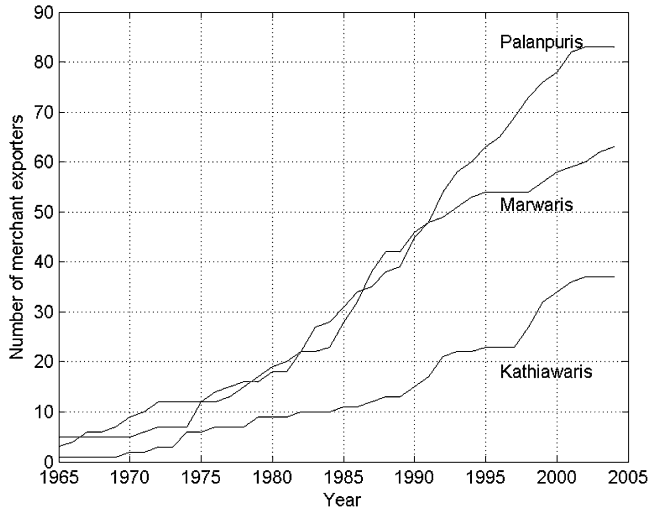

FIGURE 3
Number of Merchant Exporters

4. TESTING THE MODEL

We are particularly interested in the transition into business in the historically disadvantaged Kathiawari community. Having described the formation of a new network in this community in the late 1970's, the empirical analysis focusses on the comparison between the Kathiawaris and the more established communities in the decades that followed when testing Proposition 1. We will take account of the difference in outside options between the Marwaris and Palanpuris at the end of this section to provide additional support for the predictions of the model.

4.1. *Characteristics of entrants*

Proposition 1(b) predicts that the share of entrants with a business background should be declining over time, more steeply in communities with weak outside options. We test this prediction at the individual level by estimating the linear probability regression

$$\Pr(f_i^j = 1) = \sum_j [\gamma_j EY_i^j \cdot \delta^j + \lambda_j \delta^j], \tag{4.1}$$

where $f_i^j = 1$ if the entrepreneur's father was a businessman, $f_i^j = 0$ if he was not, EY_i^j is the year in which his firm was established, and δ^j is a vector of community dummies. We subtract 1965 from the establishment year in these regressions, so the estimated community dummies can be interpreted as the mean share of entrants with a business background around the time of the first shock to the industry. Treating the Palanpuris as the reference category, the model predicts that the coefficient on the uninteracted establishment year variable should be negative. Once we allow for changing conditions in the diamond industry over time, the sign on this coefficient is ambiguous. However, as long as these changes independently affect all three communities, we continue to expect the coefficient on the interaction of the establishment year variable with the Kathiawari dummy to be negative.

The dependent variable in Table 5, Column 1 takes the value one if the entrepreneur's father was a businessman, and zero if he was engaged in farming, a white-collar professional occupation or diamond cutting and polishing. Column 2 expands the set of occupations that are assumed

TABLE 5
Organization of production

Dependent variable	Father business (1)	Father non-agric. (2)	Schooling (3)	Father business (4)	Father non-agric. (5)	Schooling (6)
Establishment	−0·011	−0·008	−0·017	−0·016	−0·016	−0·065
year–Kathiawari	(0·004)	(0·004)	(0·024)	(0·004)	(0·004)	(0·025)
Establishment	−0·003	−0·00004	0·025	−0·003	0·0001	0·031
year–Marwari	(0·002)	(0·001)	(0·017)	(0·003)	(0·001)	(0·018)
Establishment year	−0·002	−0·001	−0·022	−0·002	−0·001	−0·030
	(0·002)	(0·001)	(0·007)	(0·002)	(0·001)	(0·007)
Kathiawari	−0·134	−0·290	−1·404	0·333	0·448	2·187
	(0·105)	(0·096)	(0·667)	(0·214)	(0·213)	(1·434)
Marwari	0·135	0·006	1·153	0·224	−0·029	−0·532
	(0·047)	(0·007)	(0·367)	(0·182)	(0·059)	(1·182)
Constant	0·778	0·988	13·176	0·859	1·093	15·423
	(0·028)	(0·008)	(0·144)	(0·135)	(0·068)	(0·649)
Age terms	No	No	No	Yes	Yes	Yes
Number of observations	737	737	737	737	737	737

Notes: Standard errors in parentheses clustered by establishment year. Establishment year is subtracted by 1965 so intercepts are interpreted as average levels in that year.
Entrepreneur's age is included, uninteracted and interacted with Kathiawari and Marwari dummies, in Columns 4–6.
Business occupations include other business/store-owner/sales, other jewelry business, diamond broker/trader, and diamond exporter.
Schooling is measured as years of educational attainment.

to be substitutable with the diamond business, setting the dependent variable to one if the entrepreneur's father was engaged in any non-agricultural occupation and zero if he was a farmer. Column 3 replaces the father's occupation with the entrepreneur's years of schooling as the dependent variable. The model assumes that the fraction of individuals with fathers in business and the ability distribution do not vary across communities and over time. This implies that average ability among the entrants into the *D* occupation must be declining over time as the network strengthens, more steeply in communities with weak outside options. Although ability is not observed directly, I assume as do Munshi and Rosenzweig (2006) in the same setting (Mumbai city) that ability maps directly into years of schooling completed. The additional prediction of the model is that schooling levels among the entrants should decline over time, more steeply among the Kathiawaris.

The coefficient on the establishment year variable is negative in the first three columns of Table 5 but only significant in Column 3. As noted, this coefficient cannot be interpreted once we allow for changing conditions within the diamond industry. However, the coefficient on the interaction of this variable with the Kathiawari dummy continues to be interpretable and, as predicted by the model is negative and significant (except with schooling as the dependent variable). The Marwari-establishment year coefficient, in contrast, is small in magnitude and imprecisely estimated.[18]

The explanation put forward in this paper for the widening gap in the business background and education of the Kathiawaris and the more established communities is that a rapidly

18. The results in Table 5 are unaffected when the establishment year is replaced by the year in which the firm started exporting.

strengthening Kathiawari network was able to support an increasing share of disadvantaged entrants. An alternative explanation is based on change in the pool of potential entrants. For example, the share of individuals with a business background (λ in the model) and the level of education could vary differentially by community across age cohorts. This could be because business opportunities were changing in the parental generation or because access to education was changing (across age cohorts) in the current generation. Because the entrepreneur's age and the establishment year of the firm are correlated, changes across cohorts could give rise to changes in the pool of potential entrants over time.

Suppose that the share of potential entrants with a business background is declining more steeply across age cohorts among the Kathiawaris. The results in Columns 1 and 2 could then be obtained even if entrepreneurs were drawn randomly from the pool of potential entrants. Alternatively, suppose that access to education exogenously increased more rapidly over time in the established communities. The relative decline in educational attainment for the Kathiawaris in Column 3 could then simply reflect the differential mapping from ability to education across communities over time.

To disentangle differential selection into the industry due to changes in the underlying networks from exogenous variation in population characteristics or access to education across age cohorts, we take advantage of the fact that entrepreneurs establish their firms at different ages. Recall from Section 2 that Proposition 1 follows through even when individuals can delay entry and establish their firms at different ages. The specifications in Columns 4–6 consequently include the entrepreneur's age and age–community interactions as additional regressors. Conditional on the age variables, the establishment year effect can then be attributed to differential selection into the industry in response to contemporaneous changes in the underlying community networks. The Kathiawari–establishment year coefficient becomes even more negative once the age terms are included and is now significant even with schooling as the dependent variable.[19] While the augmented specification in Columns 4–6 accounts for variation across community-cohorts, if children of businessmen and more educated individuals establish their firms at a younger age, then the pool of potential entrants could still weaken over time (net age effects) without requiring networks to be active. Our communities consist of thousands of individuals, and the pool of potential entrants are continually refreshed with new cohorts, so we do not expect this alternative explanation to be relevant in practice. It will nevertheless be discussed in Section 4.4 and ruled out using the results on firm performance.

To illustrate the magnitude of the community differences reported in Table 5, non-parametric estimates of the relationship between business background and the firm's establishment year are presented in Figure 4. Subsequently, this relationship is estimated net age effects or estimated net of age effects in Figure 5.[20] Ninety percent of the Marwaris and Palanpuris that started their

19. Although the age coefficients are not reported in Table 5, the coefficient on the Kathiawari–age interaction term is negative and significant in all columns. This indicates that schooling levels and business background in the population are converging across these communities, which is not surprising since the Kathiawaris started at such a low level.

20. The non-parametric kernel estimates in Figures 5–6 are constructed in two stages: (1) Separately regress father's occupation and the entrepreneur's age non-parametrically on the establishment year. Regress the residual from the first regression on the residual from the second regression to obtain a consistent estimate of the age coefficient. (2) Compute mean age by community and subtract this from each entrepreneur's age. This allows the intercepts in the second stage non-parametric regression (described below) to reflect differences in average age across communities. Subtract the differenced variable, multiplied by the estimated age coefficient, from the father's occupation. This generates a measure of father's occupation net age effects. Then non-parametrically regress this measure on the firm's establishment year, separately by community. The Epanechnikov kernel function is used in Figures 4–6. The fact that the predicted probability exceeds one for the established communities in Figure 6 is a consequence of the two-stage procedure that we employ.

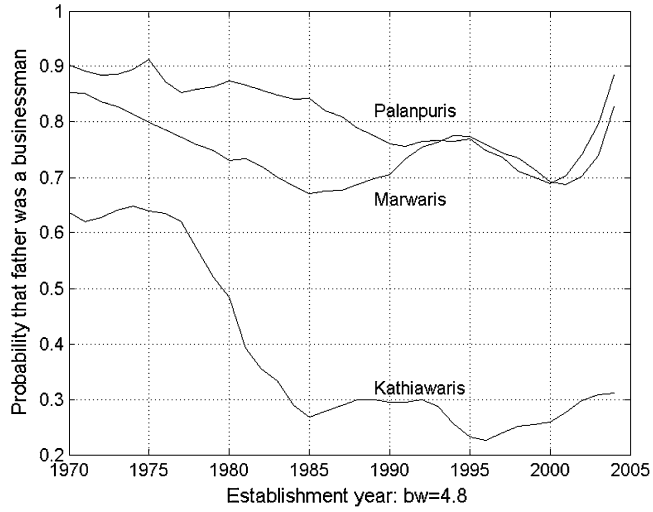

FIGURE 4
Family Background of Entering Entrepreneurs (Business)

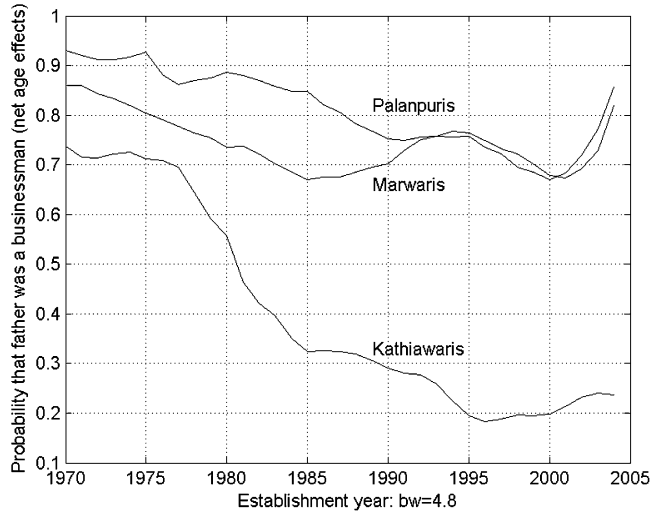

FIGURE 5
Family Background of Entering Entrepreneurs (Business)

firms in 1970 had fathers in business. This statistic drops to 70% for both communities by 2000. However, this decline is dwarfed by the corresponding decline for the entering Kathiawaris; from 70% in 1970 to 20% in 2000.[21] Figure 6 repeats this exercise, measuring business background by whether the entrepreneur's father was a farmer or not. Almost all entrants, regardless of their community, came from non-agricultural backgrounds in 1970. This pattern remains constant over time for the Marwaris and the Palanpuris since there are almost no farmers in those

21. Schooling levels match these trends in occupational background. The Marwari entrepreneurs maintain roughly 14 years of schooling, and the Palanpuris roughly 13 years of schooling, over the 1970-2000 period. The Kathiawaris start with 13 years of schooling in 1970 and fall below 11 years by 2000.

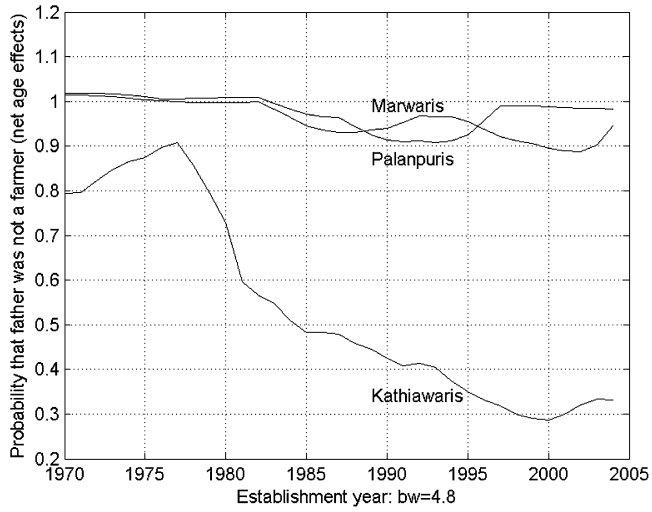

FIGURE 6
Family Background of Entering Entrepreneurs (Non-Agriculture)

communities. Starting from the late 1970's, the Kathiawari entrants, however, are increasingly likely to have fathers who were farmers and by 2000 over 70% of the Kathiawari entrants have farming backgrounds. It is evident that most of the change in the business background of entering Kathiawari entrepreneurs in Figures 4 and 5 was driven by the dramatic shift out of agriculture in this community over a single generation.

4.2. *Participation in the network*

The model implicitly assumes that participants in the network can credibly commit to cooperative behaviour. Previous research has identified the important role that the marriage institution plays in maintaining the integrity of caste-based networks (Munshi and Rosenzweig 2006, 2009) and the following episode exemplifies the role played by this institution in supporting cooperation and building strong networks in the diamond industry. In September 2003, a series of articles appeared in local Mumbai newspapers, including *The Times of India*, describing the high profile murder in Mumbai of a Palanpuri polished buyer who was based in Belgium. This individual owed an enormous sum of money to export firms in India, but since diamond transactions are not in writing, his creditors had no recourse to the legal system. Because he was based abroad, there was little they could do to punish him directly either. However, he had married a woman from a diamond family and at the insistence of his in-laws, it appears that he agreed to return to Mumbai to face his creditors. The situation subsequently took an unfortunate turn (off the equilibrium path), but what this episode makes clear is that punishments are severe in this industry and that intra-industry marriage can serve as a very powerful commitment device.

The preceding description of the marriage institution tells us that individuals that participate in the network should be more likely to marry within the industry. Our characterization of the organization of production in Section 3 tells us that merchant exporters and firms with branches in Antwerp operate independently, while all other exporters rely on their community network. These observations generate additional context-specific predictions that complement the general predictions of the model. Recall that the model assumes that individuals with a business background operate independently of the network. It follows directly from Proposition 1(b) that the share of entrants that chooses to participate or that the share of entrants that choose to participate

in the network should then be increasing across cohorts, more steeply in communities with weak outside options. In the diamond industry, this implies that the probability that (i) an entrant will organize his firm to be reliant on the network, and (ii) marry within the industry, will also be increasing over time, more steeply in communities with weak outside options. We proceed to test these predictions by estimating regressions analogous to equation (4.1),

$$\Pr(y_i^j = 1) = \sum_j [\tilde{\gamma}_j EY_i^j \cdot \delta^j + \tilde{\lambda}_j \delta^j], \tag{4.2}$$

with parental occupation replaced as the dependent variable by $y_i^j = 1$ if firm i in community j participates in the network, $y_i^j = 0$ if it operates independently.

As with equation (4.1), the coefficient on the uninteracted establishment year variable cannot be interpreted once we allow for changing conditions in the industry over time. However, as long as these changes affect all three communities, the coefficient on the community-establishment year interaction can still be interpreted. We saw in Table 5 that the Kathiawaris were relatively *less* likely to be drawn from a business background over time. Treating the Palanpuris as the reference category, we expect the Kathiawari–establishment year coefficient to be *positive* with either measure of network participation as the dependent variable.

Starting with the first measure of participation in Table 6, Column 1, we see that the organization of production of the Kathiawari and Palanpuri firms is statistically indistinguishable in 1965 (the base year) but that the Kathiawaris are significantly more likely to be organized as network firms over time. The Marwaris, in contrast, are much less likely to be dependent on their network than the Palanpuris to begin with, and there is no change in the cross-community gap over time. These estimates broadly match the cross-community differences in the organization of production presented earlier in Figures 2 and 3.

The dependent variable in Table 6, Column 2 takes the value one if the spouse's family was in the diamond industry prior to their marriage and zero otherwise. Both the Kathiawaris and the Marwaris are less likely to marry within the industry than the Palanpuris (the reference category) to begin with. While the Kathiawari-establishment year coefficient is positive and significant once again, the Marwari-establishment year coefficient is much smaller in magnitude and insignificant.[22]

One explanation for variation in the probability of intra-industry marriage across communities and over time is that these marriages respond to growth in the *stock* of firms in the industry, which expands the pool of prospective partners from within the industry and the community. The number of firms from the entrepreneur's own community that were already active when his firm was established, and the squared value of this variable, are consequently included as additional regressors in Table 6, Column 3. Reassuringly, the results reported in Column 2 are unchanged, with the Kathiawari–establishment year coefficient, in particular, continuing to be positive and significant.

22. Marriage within the community or subcaste would seem to be a pre-condition for marriage within the industry, based on our characterization of the industry-specific community network. As expected, while 92% of the entrepreneurs in the sample married within their community, the corresponding statistic for entrepreneurs who married within the industry is as high as 98%. Although caste networks have been historically very stable, recent evidence from urban India indicates that some of these traditional networks may be starting to decay, with an accompanying decline in intra-community marriage (Munshi and Rosenzweig, 2006). An alternative explanation for the cross-community variation in Table 6, Column 2 would then be that intra-community marriage has declined more rapidly over time among the urbanized Marwaris and Palanpuris. However, regressions (not reported) with our sample of entrepreneurs indicate that intra-community marriage actually *increased* over time among the Palanpuris (the reference category), while remaining roughly constant for the other two communities.

TABLE 6
Participation in the network

	Network firm	Married within the industry		Children married within industry	
Dependent variable	(1)	(2)	(3)	(4)	(5)
Establishment year–Kathiawari	0·009	0·009	0·011	0·014	0·017
	(0·003)	(0·004)	(0·005)	(0·004)	(0·005)
Establishment year–Marwari	0·002	0·003	−0·0002	0·008	0·010
	(0·004)	(0·003)	(0·001)	(0·008)	(0·008)
Establishment year	−0·005	−0·001	0·007	0·001	0·002
	(0·001)	(0·002)	(0·004)	(0·002)	(0·004)
Kathiawari	−0·050	−0·372	−0·598	−0·441	−0·509
	(0·081)	(0·089)	(0·164)	(0·094)	(0·111)
Marwari	−0·255	−0·331	−0·489	−0·509	−0·567
	(0·081)	(0·070)	(0·082)	(0·108)	(0·143)
Constant	0·735	0·463	0·671	0·709	0·790
	(0·030)	(0·034)	(0·105)	(0·041)	(0·092)
Number of active firms in the community	No	No	Yes	No	Yes
Number of observations	742	742	742	588	588

Notes: Standard errors in parentheses clustered by establishment year. Establishment year is subtracted by 1965 so intercepts are interpreted as average levels in that year.
Network firms are organized to be reliant on the network and exclude merchant exporters and vertically integrated firms. Columns 4–5 include a gender dummy as an additional regressor.
Number of active firms in the community is computed in the year that the firm was established (linear and quadratic terms are included as regressors).

Table 6, Column 4 repeats the regression that we ran for the entrepreneur in Column 2, with intra-industry marriage for the children as the dependent variable. The child's gender is now included as an additional regressor but the specification from Column 2 is otherwise unchanged. Once again, the Kathiawari–establishment year coefficient is positive and statistically significant. Table 6, Column 5 includes the number of firms from the entrepreneur's community that were already active in the industry when his firm was established (linear and quadratic terms) as additional regressors, without changing the results.

On average, 16% of the Marwaris and 45% of the Palanpuris married within the industry, with little change in these statistics over time as indicated by the estimated coefficients in Columns 2 and 3. The corresponding statistic for the Kathiawaris is 28%, with a substantial increase in intra-industry marriages over time, matching the positive and significant establishment-year coefficient in Columns 2 and 3. While a negligible fraction of the Kathiawaris who established their firms in 1970 married within the industry, this fraction reaches 50% by 2004. The marriage patterns for the children broadly match the patterns reported for their parents but change even more steeply over time, with 65% of the Kathiawaris marrying within the industry by 2004.[23]

4.3. *Network strength*

Proposition 1(a) predicts that networks will strengthen over time, more steeply in communities with weak outside options. Although network strength cannot be observed directly, firm

23. Ninety percent of the children from all three communities married within their subcaste, highlighting their continued ties to the broader community networks.

performance data can be used to infer changes in underlying networks. Let the entrepreneur's utility be determined by his firm's profits. Based on the model, average profits for firms belonging to community j in period t can then be expressed as

$$\pi_t^j = \frac{\lambda t\left[U_D + \frac{r_B}{2}\right] + (1-\lambda)\sum_{\tau=1}^{t}\Delta\omega_\tau^j\left[h(1-\lambda)\sum_{\tau=0}^{t-1}\Delta\omega_\tau^j + r_B\frac{\sum_{\tau=1}^{t}\Delta\omega_\tau^j\left(\frac{1+\underline{\omega}_\tau^j}{2}\right)}{\sum_{\tau=1}^{t}\Delta\omega_\tau^j}\right]}{\lambda t + (1-\lambda)\sum_{\tau=1}^{t}\Delta\omega_\tau^j}. \quad (4.3)$$

The first square bracket in the preceding equation measures the average profit of entrepreneurs with a business background, which remains the same across cohorts (and over time). The first term in the second square bracket measures network strength. Recall from Section 2 that the predictions of the model hold when the individual is allowed to benefit from his own cohort and the subsequent cohort. We expect these results to hold as additional cohorts provide support, ultimately including all cohorts that are active in the industry at a given point in time. The expression above thus specifies that all individuals without a business background benefit equally from the network in period t, regardless of their cohort. The second term in the second square bracket measures the effect of ability on profits for entrepreneurs without a business background. Although each entrepreneur's ability is fixed, the average ability of entrepreneurs without a business background declines over time as the ability threshold $\underline{\omega}_\tau^j$ shifts down over successive cohorts. Finally, the denominator in equation (4.3) measures the stock of firms from community j that are active in the industry in period t, with average profits for firms with and without a business background (the terms in square brackets) weighted by the corresponding stock of firms in each category in the numerator of that equation.

We know from equations (2.3) and (2.4) that $\underline{\omega}_\tau^j = 1 - \Delta\omega_\tau^j$ is decreasing over time, more steeply in communities with weaker outside options. The weighting of average ability among the entrepreneurs without a business background across cohorts in equation (4.3) will reinforce this effect and so compositional change in entrepreneurial ability will counteract changes in network strength predicted by Proposition 1(a) in the second square bracket. An unconditional comparison of changes in average firm performance across communities is consequently uninformative about variation in underlying network strength.

Prospects for identifying changes in network strength across communities and over time improve, however, if firm-level panel data are available. Firm fixed effects account for the first square bracket and the second term in the second bracket in equation (4.3). This last term is essentially the average of a changing set of firm fixed effects and once it is accounted for, changes across communities and over time reflect changes in the strength of underlying networks alone. Because the ability of entrepreneurs without a business background is declining relatively steeply in communities with weak outside options, inclusion of firm fixed effects should have a correspondingly large (positive) effect on the performance trajectory in those communities. Once firm fixed effects are included, we also expect the firm performance trajectory to be unambiguously steeper in communities with weaker outside options as predicted by the model. The fact that the second square bracket is weighted more heavily in those communities in equation (4.3) will only reinforce these predictions.

Although the preceding discussion generates predictions for conditional profits across communities and over time, the firm-level data provided by the GJEPC consist of annual exports for 95% of the surveyed firms over the 1995–2004 period (or as long as the firm was exporting if it commenced after 1995). There are three stages in the diamond manufacturing process—splitting, cutting, and polishing—with a worker assigned to each stage. Each diamond must go through this process, so firms do not benefit from economies of scale in production. Polished

TABLE 7
Firm performance

Dependent variable:	Exports					
Sample:	All firms		father non-business		All firms	
	(1)	(2)	(3)	(4)	(5)	(6)
Year–Kathiawari	1·874	7·419	10·076	16·752	2·744	8·266
	(1·511)	(2·223)	(4·758)	(5·242)	(1·626)	(2·362)
Year–Marwari	−7·514	−6·626	−8·018	−9·374	−8·214	−7·583
	(1·452)	(2·153)	(2·130)	(2·432)	(1·776)	(2·408)
Year	12·940	14·272	7·941	9·784	17·592	20·585
	(2·169)	(1·906)	(1·658)	(2·137)	(3·565)	(3·287)
Kathiawari	−22·282	−48·921	0·652	−31·167	−15·579	−44·070
	(5·454)	(11·357)	(16·602)	(13·913)	(5·833)	(11·568)
Marwari	−3·789	−8·163	−53·953	−47·865	−2·755	−4·287
	(4·773)	(13·964)	(8·815)	(9·201)	(5·639)	(14·051)
Constant	97·953	92·103	72·199	63·788	119·121	89·101
	(10·805)	(6·777)	(6·657)	(8·984)	(14·200)	(6·770)
Year-proportion small stones	—	—	—	—	−0·100	−0·123
					(0·035)	(0·031)
Firm fixed effects	No	Yes	No	Yes	No	Yes
Number of observations	6114	6114	2034	2034	5965	5965

Notes: Standard errors in parentheses clustered by year.
Exports are measured in millions of 1994 rupees. The exchange rate was 31 rupees to the dollar in that year.
Proportion small stones measures the proportion of the firm's output that is accounted for by −2, stars, and mele.
Proportion small stones is included as an additional regressor in Column 5.
The intercept and community dummies in the regressions with firm fixed effects are computed *ex post* from the estimated fixed effects.

diamonds produced in India are almost exclusively exported and firms are price takers on both the rough and polished side of the market. This implies that there is a linear mapping from exports to profits, with this mapping depending on the polished diamond yield, the labour required per unit of rough diamonds, and the price of rough and polished diamonds, all of which may vary across size classes of diamonds but are approximately equal across firms within a size class.[24]

Ignoring variation across size classes for a moment, Table 7, Column 1 regresses exports on a time trend, the interaction of the time trend with Kathiawari and Marwari dummies, and a full set of community dummies.[25] We subtract 1995 from the year variable so that the intercept and the community coefficients can be interpreted as average exports in the first year of our panel. The coefficient on the Kathiawari–year interaction term is positive but insignificant in Column 1; Kathiawari exports do not lag behind Palanpuri exports (the reference category) despite the fact that entrepreneurs from this community with relatively weak business backgrounds

24. Let the price of rough and polished diamonds be P_R and P_P, respectively, let $\theta < 1$ be the yield of polished diamonds per unit of rough diamonds, and let L be the labour requirement (at a wage rate w) per unit of rough diamonds. The firm's profit π can then be expressed as a linear function of exports, X: $\pi = [P_P - (P_R + wL)/\theta]X$.

25. To mask firm-specific figures, the firms in the database were sorted by export level and then divided into 100 groups in each year by the GJEPC. The average export level in a group was then assigned to all firms in that group. While this procedure generates noise in the export data, it does not bias the estimated community coefficients in the export regressions.

were entering the industry over time. The community-year effects in Column 1 reflect relative changes in the strength of the network and the composition of firms over time. Controlling for compositional change with firm fixed effects in Column 2, the Kathiawari–year interaction coefficient increases in size and is now positive and significant at the 5% level as predicted.[26] The increase in the steepness of the export trajectory from Columns 1 to 2 is also substantially larger for the Kathiawaris than for the Marwaris or Palanpuris, as predicted.

The export trajectories estimated with firm fixed effects can be interpreted as the average effect of the network on all firms in the community, including firms operating independently that are disproportionately owned by children of businessmen (as assumed theoretically and shown empirically in Tables 5 and 6).[27] Because the parental occupation is predetermined rather than a choice variable, the export regression can be estimated with a restricted sample consisting of entrepreneurs without a business background, thus recovering the effect of the network on those firms that benefit from it. The pattern of coefficients with the reduced sample of firms in Table 7, Columns 3and 4 broadly matches what we obtained with the full sample in Columns 1 and 2. The notable differences are that the Kathiawari–year coefficient is now large and significant even without firm fixed effects and that this coefficient subsequently increases even more with the fixed effects than it did earlier with the full sample.[28] The estimated coefficients in the fixed-effects regressions with the full sample indicate that Palanpuri exports, incorporating growth in the network and secular changes in the industry, were increasing by 450,000 dollars per year on average. The Kathiawari network increased exports for its members by 240,000 dollars over and above this benchmark. To get a sense of the importance of this differential network effect, average annual exports for Kathiawari firms were roughly 4·7 million dollars per year over the 1995–2004 period.[29] Restricting attention to firms without a business background who directly benefit from the network, the differential network effects are even larger, with the Kathiawari network increasing exports for its members, over the Palanpuri benchmark, by over 10% of average annual exports.

Although we have focussed on the comparison between the Kathiawaris and the established communities in the empirical analysis, the distinction between the Marwaris and the Palanpuris is also of some interest. Based on the discussion in Section 3, we expect the Marwaris to have superior outside options than the Palanpuris, although the gap between these communities will be small relative to the gap between them and the historically disadvantaged Kathiawaris. We do not observe a statistically significant difference in the changing background of entering entrepreneurs between these two communities in Table 5, but this may simply reflect the fact that the decline in the business background of the entering Palanpuris was so shallow. Since business background is correlated with network participation, it is not a surprise that changes over time in participation were also statistically indistinguishable between the Marwaris and Palanpuris in

26. The intercept and the community coefficients in Column 2, and in the fixed-effects regressions that follow, are computed as the average of the estimated firm fixed effects within the relevant groups, with the community coefficients interpreted as deviations from the reference (Palanpuri) category.

27. The share of entrants with a business background declines more steeply among the Kathiawaris in Table 5. Participation in the network, measured by intra-industry marriage and the organization of production, increases more steeply for the Kathiawaris in Table 6. These results, taken together, are consistent with the assumption that a family background in business and participation in the network are negatively correlated.

28. Regressions (not reported) that restrict the sample to (i) firms in which the senior partner or his children married within the industry, and (ii) firms that are organized to be reliant on the network yield estimates that are qualitatively similar to the results in Columns 3 and 4.

29. Exports are measured in millions of 1994 rupees in Table 7 and the exchange rate was 31 rupees to the dollar in that year. Average exports over the 1995–2004 period (in millions of dollars per year) are 3·8 for the Marwaris and 5·0 for the Palanpuris.

Table 6. Recall, however, that intra-industry marriage and the choice of organizational structures that left them reliant on the network were both especially infrequent for the Marwari firms. The export regressions indicate, moreover, that the Marwari network was strengthening significantly more slowly than even the Palanpuri network, providing additional support for the negative relationship between outside options and the growth in new networks.

4.4. *Alternative explanations*

The empirical analysis provides us with two stylized facts. First, the proportion of Kathiawari entrants with a business background declines relatively rapidly over time, with an accompanying decline in their educational attainment. Second, Kathiawari exports increase as fast as exports in the established communities. There is a relatively large increase in the steepness of the Kathiawari export trajectory once fixed effects are included, with Kathiawari exports now increasing more rapidly over time than exports in the established communities. My interpretation of these stylized facts is that a rapidly strengthening Kathiawari network was able to support increasingly disadvantaged entrants over time. Alternative explanations for these facts will be discussed and ruled out below.

4.4.1. Ability is mismeasured. Once we relax the assumption that each firm consists of a single entrepreneur, the senior partner could compensate for the absence of a business background or low ability by matching with capable partners, hiring well-qualified employees or gaining personal experience in the industry prior to establishing the firm. The substantial increase in the Kathiawari export trajectory once fixed effects are included, however, suggests that the relatively steep decline in the business background and education of entering entrepreneurs from that community in Table 5 and Figures 4–6 did indeed have negative consequences for firm performance.

4.4.2. Changing conditions outside the industry. Suppose that outside options, U_{NB}^{j} are declining over time, more steeply in historically disadvantaged communities. It is straightforward to verify from the model that the share of entrants without a business background will then increase over time, more steeply in historically disadvantaged communities, without requiring networks to be active. As discussed in Section 4.1, if children of businessmen and more educated individuals establish their firms at a younger age, then the pool of potential entrants will weaken over time, changing the characteristics of the entrants without requiring networks to be active. Previously, we used the entrepreneur's age to control for differential changes in parental occupations and education across community-cohorts. Changes over time in U_{NB}^{j} and the pool of potential entrants are more difficult to control for empirically because they are not observed and because they occur contemporaneously with entry decisions. Although we have shown that exit rates are low and do not vary by community, suppose that the entering Kathiawaris always consisted of a high proportion of first-generation businessmen, but that those individuals had a greater propensity to exit, perhaps because they were more willing to exploit opportunities outside the industry. Such selective exit would also generate the observed compositional change across communities. Changing conditions outside the industry, however, cannot explain the differences in firm performance across communities and over time in Table 7. Without a role for underlying community networks, there is no obvious reason why Kathiawari exports should track with Palanpuri exports despite the observed decline in education and the business background of entrants from that community over time or why the Kathiawari–year coefficient should be positive and significant once firm fixed effects are included.

4.4.3. Differences in initial conditions. If the Kathiawaris, who are more likely to be first-generation businessmen, start with a lower level of exports, and firms that start smaller grow faster, then this would explain why the export trajectory for Kathiawari firms is relatively steep. To begin with, note that differences in the export trajectory across communities are not driven by differences in the proportion of first-generation firms. Indeed, differences in the export trajectory are amplified in Table 7 when the sample is restricted to first-generation firms. Moreover, the export data allow us to compare initial exports for firms that entered over the 1995–2004 period. Initial exports are statistically indistinguishable across communities (not reported).

Although initial exports may not vary across communities, we must also account for the fact that Kathiawaris entered the diamond business later than the established communities. The model and the empirical analysis implicitly ignore individual experience effects by assuming that ability is constant over time. Introducing linear experience effects that do not vary by community creates no problems for our analysis since these effects will be subsumed in the uninteracted year variable. However, the observed patterns in Table 7, Columns 1–4 could also be obtained, without requiring networks to be active, if experience effects are concave since the Kathiawari firms were established more recently. Adding an experience-squared term actually increased the Kathiawari–year coefficient in the regression with fixed effects (not reported), with the positive and significant coefficient on this term indicating that experience effects were actually convex.

4.4.4. Changing conditions inside the industry. Changes in the new D occupation (diamond industry) that *equally* affect all communities can easily be incorporated in the empirical analysis. As noted, the uninteracted establishment year coefficient in Tables 5 and 6 and the uninteracted year coefficient in Table 7 can no longer be interpreted. However, the interaction of these variables with the Kathiawari dummy will continue to identify differential changes across communities in the composition of entrants, participation in the network, and network strength. The discussion that follows goes a step further by allowing for *differential* changes in firm profits across communities in the diamond industry.

While the mapping from exports to profits may be the same for firms specializing in a particular size class of diamonds, it could potentially vary across size classes. Kathiawari firms tend to specialize in small stones; these stones account for 57% of their output by value vs. 44% and 49% for the Marwaris and the Palanpuris, respectively.[30] Given the constraints on entry into this industry, Kathiawari exports and profit margins would have grown relatively rapidly if there was an increase in the supply of small rough stones or the demand for small polished stones, matching the patterns in Table 7, Columns 1–4. If we continue to assume that the children of businessmen shift immediately to the new industry, while individuals without a business background enter selectively as the returns in the new industry grow, then this alternative explanation for the superior performance of the Kathiawari firms would also explain the observed compositional change across communities without requiring networks to be active. To account for this possibility, I include the proportion of small stones in the firm's output interacted with time as an additional variable in the export regression (the uninteracted proportion is also included in the specification without fixed effects) in Table 7, Columns 5 and 6. The coefficient on the interaction term is actually *negative* and significant in both columns, while the Kathiawari–year coefficient continues to be positive and significant once fixed effects are included.

30. We classified stones into seven sizes in the survey: −2, stars, mele, +11, pointers, stones, and larger stones. Small stones are defined to include −2, stars, and mele.

5. Conclusion

This paper provides theoretical and empirical support for the role played by newly established community networks in facilitating occupational mobility. The key insight from the analysis is that such networks will strengthen most rapidly in communities with the weakest outside options, providing a useful counterpoint to the existing literature in which community effects are seen to reinforce persistent differences in occupational choice and wealth across families or communities.

What factors restrict such community-based mobility? One potential constraint is that historically disadvantaged communities will lack the social capital that is needed to coordinate and form new networks and to subsequently enforce cooperative behaviour (Bloch, Genicot and Ray, 2007). In the Indian context, however, caste networks have retained their importance, and social capital measured by within-caste marriage remains high, across the social hierarchy. A second potential constraint follows from the analysis in this paper, which indicates that a sufficiently large influx of initial entrants is needed to set the network on a positive trajectory. In the diamond industry, a large supply shock, together with a change in the organizational structure of established firms, allowed the Kathiawaris to enter business. Such a confluence of favourable circumstances may not occur very frequently.

Could interventions be designed in that case to move groups of individuals out of occupational traps by jump-starting new networks? To be successful, such interventions must move a sufficiently large number of individuals from the same community *and* ensure that those individuals choose activities that leave them reliant on the network in the new occupation. For example, most firms in the diamond industry rely on supplier credit to procure rough stones, which effectively prevents individuals without inherited wealth and prior connections from operating on their own. If an outsider with business skills but limited wealth received a permanent supply of bank credit, he could enter the industry by paying cash for the rough diamonds. If his credit line was sufficiently generous, he could even set-up a branch in Antwerp. The drawback of these solutions is that they do not bring other capable individuals, who do not have access to bank credit, into the business. A more cost-effective intervention would, instead, provide a limited amount of credit on a short-term basis to a sufficiently large number of individuals, facilitating the formation and subsequent expansion of new community networks. Credit infusions that exploit network externalities in this way may be especially effective in settings where these informal institutions continue to be active.

Acknowledgements. The survey upon which much of the empirical analysis is based was conducted under the leadership of Sam Taraporevala, department of sociology, St. Xavier's College, Mumbai. He was assisted by a team of interviewers supervised by Noor Aboobakar and Delnaz Chinwalla. Among the many individuals in the diamond industry that I wish to thank I am especially indebted to Suhrud Jhaveri and Rajul Jhaveri for their constant support. GJEPC graciously gave me access to their database of diamond export firms. I am grateful to Andrew Foster, Ashley Lester, and, especially, Oded Galor for helpful discussions that clarified my thinking on the theoretical framework. Abhijit Banerjee, Asim Khwaja, Nancy Luke, Bruce Meyer, Mark Rosenzweig, three anonymous referees, and numerous seminar and conference participants provided helpful comments. Katherine Casey, Scott Geller, Karim Ladha, Jonathan Stricks, and Sailesh Tiwari provided excellent research assistance and Ashish Khosla wrote the computer programs that translated the GJEPC data into a format that was suitable for analysis. Research support from the National Science Foundation through grants SES-0431827 and SES-0617847 is gratefully acknowledged. I am responsible for any errors that may remain.

REFERENCES

BANERJEE, A. and NEWMAN, A. (1993), "Occupational Choice and the Process of Development", *Journal of Political Economy*, **101**, 274–298.

BÉNABOU, R. (1996), "Equity and Efficiency in Human Capital Investment: The Local Connection", *The Review of Economic Studies*, **63**, 237–264.

BLOCH, F., GENICOT, G. and RAY, D. (2007), "Reciprocity in Groups and the Limits to Social Capital", *American Economic Review Papers and Proceedings*, **97**, 65–69.
BOWLES, S., LOURY, G. and SETHI, R. (2009), "Group Inequality" (Typescript, Columbia University).
CHANDAVARKAR, R. (1994), *The Origins of Industrial Capitalism in India: Business Strategies and the Working Classes in Bombay, 1900–1940* (Cambridge: Cambridge University Press).
CHHOTALAL, K. (1990), *Diamonds: From Mines to Markets* (Bombay: The Gem and Jewelry Export Promotion Council.
COLEMAN, J. S. (1988), "Social Capital in the Creation of Human Capital", *American Journal of Sociology*, **94**,: S95–S120.
DAMODARAN, H. (2008), *India's New Capitalists: Caste, Business, and Industry in a Modern Nation* (Ranikhet, Uttarakhand: Permanent Black).
DIAMOND WORLD (1999). *Diamond World* (Jaipur: Journal House).
DURLAUF, S. (1996), "A Theory of Persistent Income Inequality", *Journal of Economic Growth*, **1**, 75–93.
ENGELSHOVEN, M. (1999), "Diamonds and Patels: a Report on the Diamond Industry of Surat", *Contributions to Indian Sociology*, **33**, 353–377.
ENGELSHOVEN, M. (2002), "Rural to Urban Migration and the Significance of Caste: The Case of the Saurashtra Patels of Surat", in Shah, G., Rutten, M. and Streefkerk, H. (eds) *Development and Deprivation in Gujarat* (New Delhi: Sage Publications) 294–314.
FAIRLIE, R. W. and MEYER, B. D. (1996), "Ethnic and Racial Self-Employment Differences and Possible Explanations", *Journal of Human Resources*, **31**, 757–793.
FERNANDEZ, R. and ROGERSON, R. (1996), "Income Distribution, Communities, and the Quality of Public Education", *Quarterly Journal of Economics*, **111**, 35–164.
GALOR, O. and TSIDDON, D. (1997), "Technological Progress, Mobility, and Economic Growth", *American Economic Review*, **87**, 363–382.
GALOR, O. and ZEIRA, J. (1993), "Income Distribution and Macroeconomics", *Review of Economic Studies*, **60**, 35–52.
GEM AND JEWELRY EXPORT PROMOTION COUNCIL (1998), *Indian Gems and Jewelry Statistical Booklet* (Bombay: Gem and Jewelry Export Promotion Council).
GORDON, D. M., EDWARDS, R. and REICH, M. (1982), *Segmented Work, Divided Workers: The Historical Transformation of Labor in the United States* (Cambridge: Cambridge University Press).
GREIF, A. (1993), Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition", *American Economic Review* **83**, 525–548.
GREIF, A. (1994), "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies", *Journal of Political Economy*, **102**, 912–950.
HASSLER, J. and RODRÍGUEZ MORA, J. (2000), "Intelligence, Social Mobility, and Growth", *American Economic Review*, **90**, 888–908.
HAZLEHURST, L. W. (1966), *Entrepreneurship and the Merchant Castes in a Punjabi City* (Durham, NC: Duke University Press).
KOTKIN, J. (1992), *Tribes: How Race, Religion, and Identity Determine Success in the New Global Economy* (New York: Random House).
KRANTON, R. (1996), "Reciprocal Exchange: A Self-Sustaining System", *American Economic Review*, **86**, 830–851.
LAMB, H. B. (1955), "The Indian Business Communities and the Evolution of an Industrialist Class", *Pacific Affairs*, **28**, 101–116.
LOURY, G. (1977), "A Dynamic Theory of Racial Income Differences", in Le Mond, A. (ed.) *Women, Minorities and Employment Discrimination* (Lexington, MA: Lexington Books) 153–186.
MAOZ, Y. D. and MOAV, O. (1999), "Intergenerational Mobility and the Process of Development", *Economic Journal*, **109**, 677–697.
MCKENZIE, D. and RAPOPORT, H. (forthcoming), "Self-Selection Patterns in Mexico-U.S. Migration: The Role of Migration Networks", *Review of Economics and Statistics*.
MOOKHERJEE, D. and RAY, D. (2003), "Persistent Inequality", *Review of Economic Studies*, **70**, 369–393.
MUNSHI, K. (2003), "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market", *Quarterly Journal of Economics*, **118**, 549–597.
MUNSHI, K. (2007), "From Farming to International Business: The Social Auspices of Entrepreneurship in a Growing Economy" (Working Paper No. 13068, NBER).
MUNSHI, K. and ROSENZWEIG, M. (2006), "Traditional Institutions Meet the Modern World: Caste, Gender and Schooling Choice in a Globalizing Economy", *American Economic Review*, **96**, 1225–1252.
MUNSHI, K. and ROSENZWEIG, M. (2009), "Why is Mobility in India so Low? Social Insurance, Inequality and Growth", (Working Paper No. 14850, NBER).
PATEL, K. and VELLA, F. (2007), "Immigrant Networks and Their Implications for Occupational Choice and Wages" (Discussion No. 3217, IZA).
PURANI, K. (2000), "Gujarat Model of Entrepreneurial Innovation: A Study of Surat Diamond Industry" (Working Paper No. SEAKE Center, University of Brighton).
RAUCH, J. E. (2001), "Business and Social Networks in International Trade", *Journal of Economic Literature*, **39**, 1177–1203.

RICHMAN, B. (2006), "How Community Institutions Create Economic Advantage: Jewish Diamond Merchants in New York", *Law and Social Inquiry*, **31**, 383–420.
RUDNER, D. (1994) *Caste and Capitalism in Colonial India: The Nattukottai Chettiars* (Berkeley: University of Calfornia Press).
TIMBERG, T. (1978), *The Marwaris: From Traders to Industrialists* (New Delhi: Vikas Publishing House Pvt. Ltd).