A Theory of Community Formation and Social Hierarchy

Susan Athey and Saumitra Jha*

Preliminary and Incomplete

November 27, 2006

Abstract

It is well-known that trust can be sustained among small groups with repeated interactions. But with rapid urbanisation in the developing world, or in online communities where people interact without physically meeting, the set of individuals available for transactions is seldom small, and some partners may be ideal matches at some times but not at others. The lack of effective systems of law and enforcement mean that self-enforcing institutions face the brunt of the challenge of supporting exchange.

In this paper, we explore how informal institutions may sustain cooperation among members of large populations when the legal ordering is weak and outside options are strong. We delineate equilibria where cooperative exchange is made possible through the formation of distinct groups or "communities." We contrast a number of social conventions that may support such group formation, including the use of identity-specific investments and the creation of social hierarchies. We highlight a type of social hierarchy whereby agents lower in the hierarchy change groups more often, and thus cannot be trusted as much as senior agents. The value of remaining loyal to a group increases over time without resorting to institutions that "artificially" restrict trade with trustworthy agents; instead, at all points in the game agents are trusted up to the point where they would "cheat and run." We also explore the robustness of cooperation to the entry of institution-free alternative groups.

*Department of Economics, Stanford University. We would like to thank Avner Greif, Pedro Miranda, Jeffrey Wu and participants in the Stanford theory lunch, the Stanford law and economics lunch, and CDDRL student lunch for helpful comments. Athey thanks the Toulouse Network for Information Technology for support. All errors and omissions are our own.

1 Introduction

In this paper, we explore the different institutions that can sustain cooperation among members of large populations when the legal ordering is weak and outside options are strong. We are motivated by settings such as online communities and informal marketplaces in the developing world. In such settings, informal, self-enforcing institutions often emerge that help sustain trust. We focus on institutions that do not rely on third party observability of individual behavior.[1] Instead we focus on the potential for coordination among individuals to form distinct groups or "communities." If members maintain their loyalty, repeated interaction among a small group provides incentives for cooperation within the group. If the incentive for loyalty is strong enough, they will also be deterred from cheating a member of the community and then starting over in a new community. The question becomes, what creates incentives for loyalty, and how do we balance the benefits of loyalty–sustaining cooperation–against the costs of inducing members to sacrifice other trading opportunities that may be more convenient.

Formally, we model communities, or "markets," as being potential loci for interaction located on a circle. In every period, each individual realizes a new location on the circle and then chooses which market to attend, where travelling to a market is costly. Then individuals observe who else has chosen to attend that market and engage in "trust games" with partners within the market. In our baseline model, matching within a community is random, capturing the idea that selecting a community inherently commits players to interacting with all of its members. However, we also consider an extension where players can select partners within a community. The principal in the trust game chooses how much to trust the agent (the scale of the trade), and the agent chooses whether to work or shirk. We are particularly interested in the role of endogenous "institutions," by which we mean strategies where agents differentiate among trading partners based on factors other than their individual trading history. In our model, individuals may also differentiate based on a partner's record of continuous attendance, or, closely related, an agent's "seniority level," where advancement to the next level may be determined as the outcome of public randomization. We assume that each agent can observe the records/seniority levels of all other agents in his market upon arrival there, but not any features of trades other than his own.

[1] The degree of observability among traders in real-world markets, which undoubtedly varies widely across contexts, has been the subject of some debate. For example, Esfahani and Salehi-Isfahani [13] and Banerjee and Newman [6] emphasise the augmented ability of business partners to monitor one another, often through third party channels, in the informal sector. Ahmad [1], on the other hand, describes contexts where the ability to monitor is extremely constrained, often to those within a family, and it is limited observability and lack of collateral that denies informal sector entrepreneurs the credit they need to expand.

Loyalty arises when it is less attractive to start fresh in a new community than to return to the old one. A number of alternative institutions can induce loyalty. Observable community-specific monetary investments are one alternative; we propose other alternatives where investments of time substitute for money, and new arrivals trade less than more senior members. Thus, social hierarchies based on seniority engender loyalty.

We consider a series of questions in this context. First, which institutions do well in an environment where players have the option to travel to new markets in each period, and where efficiency demands the maximum possible mobility without disturbing trust? Do social hierarchies help us understand the greater mobility of the young? Second, how do institutions based on observable investments compare in terms of efficiency to those based on social hierarchies? Third, how does a social structure with more than two agents interacting improve upon what can be accomplished with pairwise relationships, absent any technological reason for group production? Fourth, can we find equilibria that have certain desirable robustness properties? For example, can the equilibria be characterized by behavior where all individuals are trusted up to the point where their incentive constraints bind? And can the equilibria withstand entry by an alternative group that does not have barriers to entry?

We begin our formal analysis by briefly describing two benchmarks. In a *peddler's equilibrium*, individuals trade in the most convenient markets. With large populations, this requires high degrees of patience to sustain cooperation. The *loyalty equilibrium* allows individuals to coordinate into communities without supporting institutions. Here, the ability to start over in a new group makes cooperation break down altogether, since an individual can potentially avoid all future contact with a partner he has cheated.

We then consider equilibria where individuals make *specific investments* in a shared group identity. Such investments can take a number of forms, depending on the exogenous dimensions through which groups may assert distinction between one another and conformity within themselves. The dimensions could be as diverse as language, "appreciation capital" for a particular form of music or food, choosing where to live, getting a tattoo, choosing clothes, even choosing a religion or ethical code (Becker and Stigler [7].) We show that for this class of equilibria, for some parameter values (when the entry barriers are high) the group size for which it is easiest to sustain cooperation is two: partners make specific investments that commit them to one another, and there is no gain to being part of a larger group. However, for other parameter values, partners are especially tempted to "cheat and leave," and the presence of additional groups makes it cheaper for an individual to avoid a market where he has cheated someone in the past. If the population is fixed, increasing the number of groups decreases the group size. But the size of the group plays no direct role in enabling cooperation.

If no dimensions exist that can be used to form the basis of group distinctiveness, then attendance within a group– and hence not elsewhere– may act as a specific investment. In a *purgatorial equilibrium*, agents do not trade with new members of a group for a pre-specified number of periods, or more generally, agents engage in reduced trade with new members. An interesting question concerns the optimal group size in a purgatorial equilibrium. Answering this question helps us to understand why groups can be more efficient than partnerships, and thus why social institutions can play an important role. When there are only two players in each group, the equilibrium has the flavor of a partnership, or a marriage. However, whenever one partner dies or leaves (due to unusually high costs of travel to meet the group), the remaining partner must reduce his trade temporarily until the new partner finishes purgatory. On the other hand, the incentives to cheat a partner within a group of two are quite low, since there will be no benefit to returning after cheating. Thus, a player will be most tempted to "cheat and leave," behavior that is dissuaded by the desire to avoid going through purgatory in a new group. We show that larger groups can actually increase efficiency, by increasing the return to seniority. Holding features of purgatory fixed, with larger groups, senior players trade more often with other senior players and less often with junior players. Junior players trade more often with other junior players and less often with senior players. This increases the gap between the value of being junior and that of being senior. Then, the incentive to "cheat and leave" is reduced, and purgatory can be made less severe. Although the larger group reduces the sanction to cheating and *staying* within the group, for many parameter values and for moderate group sizes this constraint is not the binding one.

Although the equilibria based on specific investments as well as those based on purgatory seem to be consistent with some real-world examples, they also share two somewhat undesirable features. First, newly arrived agents or agents who have not made group-specific investments are actually trustworthy (they would be willing to cooperate today if they expected it would lead to cooperation in the future when rematched with today's partner), yet trust is withheld. In the case of purgatory, the period before purgatory ends, an individual's continuation value is the same as that of senior agents, but yet equilibrium strategies call for reduced trade. Thus, if an individual is matched with a newly arrived agent in the purgatorial equilibrium, the pair would be better off agreeing to trade at a higher level.

A second undesirable feature is that a group as a whole would often be better off if it abandoned an institution of requiring specific investments or withholding trade from newcomers (so long as all other groups maintained their institutions), because these requirements serve only to induce loyalty in other groups.

Motivated by these concerns, we construct an alternative class of equilibria. A *hier-*

archical equilibrium is closely related to a purgatorial equilibrium. It also has the feature that, on average, agents who have longer attendance records are given more trust. However, instead of using a fixed length of time where an agent is at the "lower level," there may be multiple levels, and further each level of the hierarchy has a fixed size and advancement from one level to the next occurs only if agents at the higher level leave or die. This feature plays two roles. First, it simplifies the analysis by keeping the seniority distribution stable over time. Second, and more substantive, when individuals at the next-to-highest level do not deterministically advance to the highest level, their continuation values are not the same as those at the highest level. Thus, the individuals at the next-to-highest level have less to lose from a cheat-and-leave strategy, and the potential exists to find equilibria where (i) trust increases at each level of the hierarchy, but (ii) individuals at each level of the hierarchy are trusted to the point where their incentive constraints bind.

We allow that agents lower in the hierarchy actually leave in equilibrium when they receive adverse location shocks. The less likely an agent is to leave, the greater his incentive to cooperate today (that is, the greater are the losses from a cheat and leave strategy that entails starting anew in a new community), and the more trustworthy is the agent. Thus, mobility in equilibrium reinforces the effects described above, creating a second reason that junior agents are less trustworthy. We construct some numerical examples of these equilibria.

Although the hierarchical equilibrium potentially has one of our desired robustness properties, it still shares the feature that the institutions of one group are critical for supporting trust in other groups. A single community could potentially adopt the convention that all individuals (including newcomers) are treated the same as the most senior level of other communities. Individuals would be dissuaded from leaving, since starting over in any other group would mean going to the bottom of the hierarchy. Thus, cooperation could be sustained in the single community without institutions.

In addition, the hierarchical equilibrium (like the purgatorial equilibrium) has some inefficiencies relative to an equilibrium based on specific investments. Since senior agents are randomly matched with (untrustworthy) junior agents within a community, they suffer losses in potential trade even though the senior agents are trustworthy. This decreases the value of being a senior agent, reducing the incentive for loyalty and necessitating still lower levels of trade at the bottom of the hierarchy. In contrast, in an equilibrium based on specific investments, once investments have been made, trade occurs efficiently.

We then consider an extension of the model where agents can choose their partners within a community–a *choice-based hierarchy*. With this extension, we can avoid the inefficiency just described. In an appropriately constructed hierarchy, all individuals prefer to trust senior members, since they are more trustworthy (less tempted to use a cheat-and-leave strategy) and the scale of trade can be greater. Thus, senior agents will receive more trades. This type of equilibrium is more efficient–if senior agents are fully trustworthy, first-best trade may be attained, though travel costs will still be inefficiently high, since individuals return to the same community rather than attend the closest community. It also has a further robustness advantage, however. In a choice-based hierarchy, senior members get higher payoffs than they would in an institution-free, full-trust community. So, it is possible that this type of equilibrium is robust to the entry of an institution-free community.

The rest of the paper proceeds as follows. In Section 2, we discuss some motivational examples that have features we wish to capture in our model. In Section 3, we present our main model of group formation and social hierarchy. Section 4 looks at benchmark equilibria that do not rely on "institutions" to support cooperation or identities. Section 5 is devoted characterizing equilibria based on specific investments, the purgatorial equilibrium, and the hierarchical equilibrium. Section 6 discusses the extension of the model where players can select their partners within a market. In Section 7, we review related work. Section 8 concludes.

2 Motivational Examples

We are motivated by two types of examples. First, we discuss examples from online communities, such as open source software (OSS) projects and internet discussion groups. Second, we describe examples from the developing world. In each example, certain elements are common. Market participants may have insufficient incentives to deliver promised goods and services, or to contribute to a public good. There are multiple possible groups individuals may affiliate with. Trades and interactions are repeated, but are also spread among different participants rather than being exclusive. Neither the legal environment nor community sanctions seem to play an important role in resolving disputes or sanctioning free riders.

2.1 Online Communities

OSS communities and internet discussion groups are a major phenomenon (see, e.g. , Raymond [35], Lerner and Tirole [30], or Shah [36]), and as of early 2006 there were over 100,000 open source projects with over a million registered users. Individuals also provide information and support for commercial products, especially software products, in internet discussion groups. These discussion groups substitute for costly technical support provided by companies. Many features of online communities at first seem surprising from the perspective of standard incentive theory. Individuals regularly contribute to a public good despite the lack of monetary incentives, the large sizes of the groups, and the relative anonymity of the interactions. A closer look at these communities reveals that it is not just public-spiritedness that motivates users. Many groups have formal or informal hierarchies. In OSS communities, there is often a group of senior members known as "committers" who have the authority to incorporate their own code into the project, as well as the code of others. In addition, there is often an informal hierarchy, where more senior members will help answer questions for other senior members, but not junior members. In other internet discussion groups, it is common to show the date a user joined beside their posts. As Microsoft's Office Forum web page indicates, "the information about your activity in the Community (such as how many posts you have contributed...) gives others a sense of your trustworthiness and a way to gauge how valuable your comments might be." Other types of designations (such as top reviewers for Amazon.com, or Most Valuable Professionals on Microsoft Forums) also help distinguish the senior members of a community. Our analysis helps to understand these phenomena.

Consider how OSS projects fit into the model outlined in the introduction. Even though in principle individual contributions may be publicly observable in many online communities, in practice it requires a substantial investment to learn about the quality of their work. For example, one might need to read the code a programmer has written carefully. In addition, some aspects of support and collaboration in OSS projects take place in private correspondence. It may thus be difficult for outsiders to ascertain the quality of interactions others are having in the community, and one may only learn about an individual's behavior by trying to adopt the code or otherwise having a close interaction with someone.

Our model applies to a setting where programmers select whether to become involved in a particular project. Some projects may be more closely related to the programmer's skills or needs at particular points in time. Programmers may write code for different purposes. Individuals associated with the project have needs arise which may be met using code that others have written, and they may choose whether or not to invest in reading and trying to use the code (as opposed to writing from scratch or finding other sources). They may then have support questions. The author of the code then chooses whether to support the user, answer questions, etc. Some of the author's choices may be made ex ante, such as documentation and code quality, modularity/adaptability, etc. The "scale" of trust on the part of the user can be interpreted in several ways: for example, the user might make use of only the basic features of the code, or the user might invest rewriting portions himself. The user observes whether the ex ante quality was high, and also observes whether support was provided. The author gets higher utility the more the user relies on the author's original code (rather than rewriting it), because duplicate/competing versions of the same code are avoided.

2.2 Informal marketplaces in Lima, 1985

Hernando de Soto's [39] account of street vendors in the slums of Lima provides a detailed look at the emergence of new informal property rights in a period of rapid urbanisation. In 1985, Lima's informal marketplaces outnumbered her official commercial districts by a ratio of five to one. These employed approximately four hundred and thirty-nine thousand people, responsible for annual commerce of around $40.9 million.

De Soto charts the course followed by newcomers seeking to establish themselves in the informal sector. Many begin as peddlers who move around in search of the most favourable opportunity. These sell products of easily observable quality and are largely denied access to credit. This can change if they establish a predictable route, or alternatively set up in stationary location- a "pitch" or marketplace, usually by public roads that carry large amounts of traffic.

When a marketplace is young, resident vendors encourage newcomers to set up as well. However, as a marketplace becomes more "established," incumbents begin to resist further expansion, sometimes violently. Within the group, incumbents begin to trade and sell property rights - for example, time shares on a stall become available, allowing food vendors to set up specifically during meal times and then be replaced by other tradespeople later in the day.

The particular elements of interest in this example are that informal marketplaces are established as the result of coordination by individual vendors in certain locations. There are limits to group size. Newcomers face challenges in becoming established, while incumbents enjoy greater trust from customers and other vendors. Thus, this is an example of a hierarchical system, where returning to a regular location creates opportunities for repeated interaction, and seniority provides incentives to stay.

2.3 A commercial revolution in Kenya, 1940-50

Jean Ensminger's [12] study of the expansion of trade by the Galole Orma in Kenya provides a telling portrait of competing communities and institutions. The Orma people were pastoralists near Lamu on the Swahili Coast of Kenya. Orma were governed by a seniority based system, whereby the elders ruled the tribe, and possessed most of the wealth. They were hostile to trade.

The town of Lamu was an independent Muslim trading town. Unusually, it was an African republic with a seniority-based hierarchy ruled by elders. Tradition dictated that elders made investments in ownership "stone houses." In the 1930s, Shambaro, who was an Orman orphan and dispossessed of family wealth, became the first Orma trader to set up permanent trading post in Lamu. He converted to Islam and became a successful agent for cattle trading.

By 1940, almost the whole Orma population had converted to Islam. This lead to a collapse of the seniority system among Orma, as the young men converted first, and the "elders were disrespected, their cattle sold without permission." The culture became characterized by competitive religiosity: people outdoing each other to show how Islamic they were. Initially, they made pilgrimages to Mecca at a cost of $200. Traders invested in observables specific to Islam- e.g. mosques, buying pilgrimages for others. Anthropologist Last [29] explicitly calls it a "membership fee" for doing trade. By assuming their new identity, the Orma joined the larger seniority system of the Lamu traders.

This example highlights the (more or less successful) use of institutions such as seniority and group-specific investments, and it illustrates that the ability to "cheat and leave" can undermine cooperation.

2.4 Supplier relationships in Hanoi, 1995-97

McMillan and Woodruff [32], [33] provide a case study of new entrepreneurs in Vietnam, describing how the deregulation of Vietnam's economy resulted in the sudden flowering of small-scale commercial activity. They document a number of important stylized facts about supplier relationships. First, supplier firms offer 15% less credit to customers that have the possibility of going to alternative suppliers. Second, the duration of firms' relationship improves credit. Third, those that lived further away were less likely to be given credit. Fourth, 12% engaged in cafe chats with other members at least once a month, sharing information, but the vast majority did not. Yet, even among those that did engage in chats, almost all said they would trade with someone who had been accused of cheating by another member. Finally, itinerant temptations to cheat seemed to play an important role, as the temptation to cheat in bad times was particularly severe.

This example illustrates the importance of long-term relationships in enforcing cooperation, and the lack of reliance on group sanctions, even in the presence of organized groups that share information. Outside options were important in creating the incentive to cheat. Seniority seemed to play an important role in sustaining trust.

3 The Model

The game takes place over an infinite horizon, with periods indexed by $t = 1, .., \infty$. Consider the following environment:

A 1. *There is a stationary population of N individuals, of whom δ survive every period. N is even.*

Thus to maintain a stationary population, $(1-\delta)N$ are born every period. Let I be the set of individuals alive at a given point in time, and when a player dies, another player inherits his index with a null history.

A 2. *All exchange takes place in one of $J \leq N/2$ markets evenly spaced over a circular city, circumference $2\bar{c}$. N/J, $\delta N/J$ and $(1-\delta)N/J$ are integers.*

In equilibrium, trade may or may not take place at a particular market. The markets may be thought of as actual geographical marketplaces, as with the concentrations of street-vending in Peru, different online communities or OSS projects, or alternatively as communities that are differentiated in some other way.

A 3. *In each period t, individual i begins the period at one of the J markets, denoted θ_{it}. We assume that there is no aggregate uncertainty about the number of individuals born in each location, but an individual's location is random: each population center will have N/J individuals, and each individual views each population center as equally likely.*

If an individual's trading opportunity is θ_{it} and a market j is located at θ_j, the cost of attending market j is

$$c_{ij} = \min ||\theta_i - \theta_j||. \tag{1}$$

Thus a particular location θ_i implies a J-dimensional vector of costs $\mathbf{c_i} = (c_{i1} \ldots c_{iJ})^T$.

To simplify the calculations, we make the following assumption about the timing of death:

A 4. *In each period, individuals die after selecting a market but before they make any decisions. Newly born agents replace them starting at the same market. After arriving at a market, an agent observes the identities of all present in their chosen market, but not in other markets.*[2]

Next, we consider the interaction of players in a market.

[2]This assumption, which makes the total number of deaths in each market deterministic, has a few awkward features. For example, players might be "surprised" by seeing a change in the market size, if equilibrium calls for market size to be constant. To deal with this, we can assume that there is a small probability that the population grows (in proportion to J) in each period. If equilibrium requires that market size only changes when this event happens, then when it happens, agents will attribute the change to an exogenous shock. Below, we also discuss an alternative approach to ruling out unexpected changes in market size: markets may have an upper bound on the number of agents who may enter, and new agents are born in each market until the maximum size is reached.

A 5. *Individuals engage in transactions in pairs. Each pair has a principal and an agent, and each individual can play each role exactly once in each period. After being matched to an agent, the principal chooses how much to trust the agent, offering a trade of scale* $\lambda \in [0,1]$ *to the agent. Then, if the agent works* (W)*, payoffs are* $(\lambda R, \lambda w)$ *to the principal and agent, respectively. If the agent shirks* (S)*, payoffs are* $(-\lambda r, \lambda s)$*. We assume* $w < s$*,* $r > 0$ *and* $R + w > s - r$*.*

A 6. *Inside a market with* n_j *players, each individual acts as principal to one randomly assigned agent, and the randomization ensures that each individual is an agent to no more than one principal.*

Notice that a market will therefore have n_j trust games taking place, and that an individual will typically be a principal to one individual and an agent to a different individual. Later, we will consider an extension where principals can select their agents from individuals in the same market.

A 7. *Players observe personal bilateral trade histories from birth to the current period, but not trades of other players.*

An individual's trading history with a given partner has two components, one where the individual acted as a principal and one where the individual acted as an agent. When the individual is the principal, the possibilities are $H = \{\{(\lambda, W), (\lambda, S), \emptyset\} : \lambda \in [0,1]\}$: the principal trusted the agent at scale λ and the partner, as agent, worked (W) or shirked (S); or the pair was not matched in this way $(\emptyset)$. The same possibilities can occur when treating the individual as the agent and the partner as the principal.

Define an "attendance spell" to be a sequence of periods of continuous attendance in a given market. A "joint attendance spell" for a pair of individuals is a sequence of periods of continuous attendance for both individuals.

A 8. *Each market maintains a record of an agent's prior attendance in the current attendance spell for each individual i that is costlessly observable to all who attend it.*

The record-sharing is local: each market has no record of which of the other markets an individual attended. There is no market-level record of whether an individual cheated or cooperated during this period- that information is only available to partners to an exchange.

Definition 3.1. *A (period)* ***trading strategy*** $s_i(j_i^t, h_{i,-\mathbf{i}}^t, \theta_{it})$ *is a mapping from: market attendance records* $j_i^t = \{j_{it-1}, \ldots j_{i\tau_{i0}}\}$*; bilateral trading histories for each i and* $k \neq i : h_{i,k}^t \in H^{2(t-1-\tau_{i0})}$*; location in city today* $\theta_{it} \in [0, 2\bar{c}]$*; to: the market to attend today* $j_{it} \in J$*; as principal– how much to trust each agent in the market, if matched*

$\lambda_{it} = (\lambda_{ikt})_{k\in I} \in [0,1]^N$; *as agent– to work or shirk if trusted by all potential partners* $e_{it} = (e_{ikt})_{k\in I\setminus i} \in \{W,S\}^{N-1}$.

As an equilibrium concept, we focus on subgame perfect Nash equilibria.

4 Benchmarks

Without the possibility of repeated interactions between partners or the presence of institutions supporting group formation, the unique equilibrium is that no trust is given and no trade takes place.

4.1 Peddler's equilibrium

Not surprisingly, providing individuals are extremely patient (long-lived), it is possible to sustain cooperation even with large populations and alternative trading partners. In the *peddler's equilibrium*, traders choose the most profitable venue to trade, attending the nearest market every period.

Definition 4.1. *A* ***peddler's equilibrium*** *is a subgame perfect Nash equilibrium where players employ the following strategies:*

1. *All players choose the closest market every period.*
2. *If shirked on by an agent when acting as a principal, never trust him in the future and shirk if trusted. Having cheated a principal once when acting as an agent, always shirk if trusted.*
3. *As principals: trust all agents fully except as given by 2.*
4. *As agents: players always work if trusted except as given by 2.*

Let $B_{it} \subseteq I\setminus\{i\}$ denote the set of currently living agents with whom either of these outcomes has occurred in the past. Abusing notation, we also let B_{it} denote the dimensionality of that set. In the peddler's equilibrium, the player perceives it as equally likely to be matched with all other individuals in the population upon attending any given market. Thus, a player expects to be matched as an agent with a principal from B_{it} with probability $B_{it}/(N-1)$ and expects to matched as a principal with an agent from B_{it} with probability $B_{it}/(N-1)$.

The "no-shirking" incentive constraint is

$$s - w \leq \frac{\delta^2}{1-\delta^2}\frac{1}{N-1}(R+w), \tag{2}$$

where $\frac{\delta^2}{1-\delta^2}$ weights the lost stream of profits by the probability that both agents are alive. $\frac{1}{N-1}$ is the probability that the player matches the partner again as an agent or as a principal, conditional on both partners being alive. Clearly, population size reduces the effectiveness of the sanction on cheaters. Note that modifying the scale of trade does not affect this constraint–if all trades take place at a smaller or larger scale, the key ratio $(s-w)/(R+w)$ remains unchanged.

There are two potential sources of inefficiency in any equilibrium- the costs of attending markets and the loss of social surplus when individuals are forced to engage in self-production. The peddler's equilibrium is the most economical of transportation costs, as traders simply go to the market most favoured by their opportunities.[3]

4.2 Loyalty equilibrium

A big drawback of the peddler's equilibrium is its fragility when faced with increases in the population size. If smaller groups of individuals can coordinate upon returning period after period to certain markets, they may be able to increase their frequency of interaction enough to be able to sustain trade. However, they will be overwhelmingly tempted to cheat other agents in the market and then start over elsewhere.

To see a knife-edge set of conditions where it may be possible to sustain cooperation without additional institutions, consider parameter values where the following holds: if there were no transportation costs and a single potential market, cooperation could be sustained if and only if there are less than or equal to N/J^A players the population. With these parameter values, there is a credible threat by existing markets to exclude newcomers: taking as given the movements of other players, a player moving markets would increase the market size to the critical level where she would upset cooperation and thus would lose the possibility of trading in the new market. Clearly, such an equilibrium would be fragile to any other type of uncertainty that might perturb the number of players in each market.

An alternative set of assumptions to consider would be that markets can set a capacity constraint, so that no more than a specified number of agents could attend in any period. If newly born agents arrive last, they would be prevented from remaining in a market that was already "full." Then, older agents who moved could "take up" slots in a market, and

[3]We made a couple of simplifying assumptions that ensured that all markets have at least two players present in each period. If there were more markets than traders, there would be a benefit to coordinating on a smaller number of markets, effectively shutting down some of the potential markets. Doing so does not change the probability an individual meets a particular partner, so long as there are enough markets that $R+w \geq \bar{c}/J^A$, where J^A is the number of active markets and $\bar{c}/J^A$ is the maximum distance a newly born individual might need to travel to attend an active market.

newly born agents would choose to move on to the market vacated by the older agent. In such a model, the loyalty equilibrium would not survive, because an agent could always move to a new market and start again after cheating without upsetting the balance of the number of players.

5 Communities with Institutions

In this section, we consider institutions that potentially allow cooperation for arbitrary population sizes. We begin by analyzing equilibria based on specific investments, and then we turn to seniority-based hierarchies.

5.1 Heterogeneous communities and investments in identity

One way of sustaining cooperation as populations rise is to make these communities "culturally" distinct. These distinctions allow individuals to make specific investments in a cultural identity that marks them as members of a community.

Since it is costly to invest in and relatively difficult to reverse, the development of human capital in identity-specific attributes is a natural candidate for sustaining group formation. A useful example is provided by Fryer [15], who discusses the importance of educational investments in distinguishing blacks who are seeking to leave from those who committed to the community. More generally, *appreciation capital* investments in a range of cultural dimensions- e.g. food, language, clothing and etiquette- can be viewed as strategic decisions that define a person's group membership.[4] Fryer and Levitt [16] explore the phenomenon by comparing the labour market outcomes of assuming an overtly black name. In the case of the Orma, the dimension of investment was made explicit by the competitive religiosity of the new traders, who having converted to Islam, sought to outdo one another through the overt investment in mosques.

Let the cost of the identity investment for community j be m_j.

Definition 5.1. *An* ***identity investment equilibrium*** *is a subgame perfect equilibrium with the following properties:*

1. *All players choose the active market at birth that minimises $c_{ij} + m_j$. They make the corresponding investment.*

2. *As principals: players in market j maximally trust all present who have not shirked in the past and who have made investment m_j in the current attendance spell.*

[4] The concept of appreciation capital originates with Becker and Stigler [7], though they do not consider the potentially strategic dimension of such investments.

3. *As <u>agents</u>: players in market j always cooperate with all individuals who have made investment m_j in the current attendance spell.*

4. *Principal and agent revert to (zero trust, shirk) in trust subgame if either individual shirks in the current joint attendance spell, or if either individual has not made investment m_j in her current attendance spell.*

5. *If an individual leaves a market, she chooses the market that minimises $c_{ij} + m_j$.*

6. *After cheating, an individual moves to a new market. After being cheated, a principal selects markets as if no cheating had occurred.*

We can consider two types of identity investment equilibria, one where individuals always return to their home market along the equilibrium path, and another where individuals may change markets in equilibrium. Here we focus on strategies where individuals return on the equilibrium path. Note that in both types of equilibria, there is no advantage to reducing the scale of trade below 1.

The crucial constraint is that of the agent shirking and then leaving. If the agent shirks, he faces a choice: endure the sanction of no trade with the current agent, who he is likely to run into in the current market; or else pay the investment cost elsewhere and start over. Note that part 7 of the definition of the equilibrium specifies that individuals must leave after cheating. If individuals return with some probability after cheating, then a principal will find it less attractive to return, undermining the sanction to the individual that cheated. Although it is possible to relax part 7 somewhat (for example, so that an agent who cheated returns with low probability, and a principal who was cheated returns with high probability), we will not pursue that generalization.

An interesting question concerns the group size that most easily sustains cooperation. Decreasing group size by increasing the number of active markets has several beneficial effects. First, it minimizes transportation costs in the first period. Second, it diminishes the attractiveness of a cheat-and-return strategy. However, it also makes a pure cheat-and-leave strategy more attractive since it reduces the extra transportation costs required in expectation after abandoning a market. The following proposition first identifies parameter values for which an identity investment equilibrium exists for $J^A = N/2$. For fixed δ, the parameter restriction can be satisfied for large N, implying that the number of agents can grow arbitrarily large without upsetting cooperation and without requiring greater identity investments, so long as the number of communities increases proportionately. The second part of this proposition shows that the sustainable group size varies with parameter values. The proofs suggest that when $\bar{c}$ is large, so that individuals are often tempted to return to the home community rather than make a specific investment of size $\bar{c}$ in a new community, then small group sizes may be required to sustain equilibrium. When $\bar{c}$ is small, so that the cheat and leave strategy is dominant, the large group sizes are better, because of the extra transportation costs imposed on agents who cheat and leave.

Proposition 5.1. *(i) Consider the class of identity investment equilibria where players always return to their original home market in equilibrium. For any $\bar{c}$, δ, N, R, w, and s satisfying our basic assumptions, if $J = N/2$,*

$$\bar{c} \leq \frac{\delta}{1-\delta^2} \frac{N}{N+4}(R+w) \tag{3}$$

and

$$s - w \leq \delta \left(1 + \frac{8}{N}\right) \bar{c}, \tag{4}$$

there exists an identity investment equilibrium with $J^A = N/2$ and $\bar{c} = m_j$ for all j.
(ii) There are parameter values for which an efficient equilibrium fails to exist with $J^A = N/2$ while it does exist for some $J^A < N/2$. If the conditions of part (i) are satisfied and (3) holds with equality, then an efficient equilibrium exists in the specified class with $J^A = N/2$, but not for $J^A < N/2$.

Proof: (i) Agents must be persuaded to return to the home market when their initial location is $\bar{c}$ away:

$$\bar{c} \leq m_j.$$

The most efficient equilibrium in this class will thus have $\bar{c} = m_j$ for all j, if the other constraints can be satisfied. We drop the j subscript since symmetric investments will be efficient.

The following condition guarantees that even when an agent arrives at the home market, she still prefers to travel rather than return to the market where she cheated someone,

$$m + \frac{2\bar{c}}{J^A} \leq \frac{\delta}{1-\delta^2} \frac{1}{N/J^A - 1}(R+w),$$

or, when $\bar{c} = m$,

$$\bar{c} \leq \frac{\delta}{1-\delta^2} \frac{J^A}{(N/J^A - 1)(J^A + 2)}(R+w). \tag{5}$$

The right-hand side is increasing in J^A, which implies that if cheat-and-leave dominates cheat-and-return for a small J^A, it also will for large J^A. In this case, an individual with home market j is deterred from cheating if

$$s - w \leq \delta \left(m + \frac{2\bar{c}}{(J^A)^2}\right) = \delta \left(1 + \frac{2}{(J^A)^2}\right) \bar{c}. \tag{6}$$

(ii) First, we show an example where an identity investment equilibrium fails to exist with $J^A = N/2$ but does exist with $J^A = N/4$. Consider the following conditions:

$$\bar{c} \leq \frac{\delta}{1-\delta^2} \frac{N}{3(N+8)} (R+w), \tag{7}$$

and

$$\delta \left(1 + \frac{8}{N^2}\right) \bar{c} < s - w \leq \delta \left(1 + \frac{16}{N^2}\right) \bar{c}. \tag{8}$$

Condition (7) guarantees that individuals always start over in a new market after cheating if $J^A = N/4$, which implies that they also do if $J^A = N/2$. But then, the cheat-and-leave constraint (6) binds first for $J^A = N/2$, since the transportation costs to start over are on average higher when there are fewer communities.

If the conditions of part (i) are satisfied and (3) holds with equality, the sanction following cheating is sufficient to induce leaving after cheating when $J^A = N/2$, but insufficient to induce agents to leave after cheating when $J^A < N/2$. $\square$

For part (i) of the proposition, note that as N grows, an equilibrium continues to exist for many parameter values so long as the number of potential markets grows too. Note also that the conditions of the proposition are sufficient but not necessary for existence of an equilibrium; larger specific investments can deter cheating and leaving more effectively, though if they get large enough, cheating and staying may become more attractive. In that case smaller group sizes help.

Part (ii) follows because when transportation costs are high and potential partners require large specific investments, it is easiest to sustain cooperation with a "marriage" or "exclusivity": you make a specific investment that commits you to a partner, and you return loyally to that partner. The small group size implies that returning to the same market after cheating is especially unattractive, but starting over in a new community is also prohibitively costly. On the other hand, when transportation costs are low and potential partners require small specific investments, individuals are not tempted to return to the home market after cheating, but a smaller number of groups increases the costs of avoiding a group where a partner has cheated in the past.

Now consider equilibria where players leave in equilibrium. Note that for this type of equilibrium to exist, we require part 6 of the definition of an identity investment equilibrium: players who leave must pay the cost of joining anew if they return. Otherwise, if a player accumulates investments in multiple markets, it will be tempting to cheat a principal in one market and go to another market where the agent has a membership. An equilibrium with equilibrium-path market switching is potentially more efficient, because it economizes on transportation costs. After the initial period, agents compare two wasteful expenditures: transportation costs, and the costs of making a new specific investment in a market. A lower specific investment leads agents to attend closer markets, reducing transportation costs expended in equilibrium. The lowest possible specific investment that still induces cooperation will be efficient.

An unattractive feature of the identity investment equilibria is that the role of identity investments is to act as a barrier to members of other markets, not to retain members of one's own. By making visible and irreversible investments in a specific identity, individuals are revealing more about which communities they are not part of, than which they are. Furthermore, the equilibrium requires that all markets impose membership barriers. Implicitly, the strategies used by members of one community are constructed in order to sustain the cohesiveness of another. If the members of even one community fail to impose barriers to entry, cooperation will fail in other communities. Another unattractive feature of the equilibrium is that it relies on exogenous specific investments that may or may not be available in a given context.

Returning to the motivating example of the Orma, becoming a Muslim trader required specific investments, but these were not high enough relative to the trading opportunities available to Muslims to prevent the demise of the Orma culture. We do not see specific investments used widely in online communities, possibly because it is too difficult to observe and implement such investments, or perhaps because there are too many potential communities and not all will make use of them.

5.1.1 Extension: High temptations and overlapping identities

Another possible class of equilibria relaxes Condition 6 of the identity investment equilibrium, and allows individuals to have overlapping identities. Suppose we consider the case with high $\bar{c}$, and allow individuals to assume more than one identity- then we can find equilibria where communities set up membership barriers to discourage defection to *nearby* markets, but allow individuals to adopt dual identities that allow them to trade in far away markets. (This requires that membership fees for other markets are all observable.) In the simplest case, an individual pays a fee and identifies with two markets, both on opposite sides of the circle. Thus he alternates between these markets, depending on his realization. The reduced frequency of interaction among group members makes cooperation harder to sustain.

5.2 Purgatorial Equilibria

We now consider the creation of community attendance records. A community attendance record that is common knowledge acts as an institutional memory of members' presence. In a purgatorial equilibrium, players are excluded from trade in their first few periods of attendance at a new market. In this way, time spent at the market without trade plays a role analagous to a specific investment.

Definition 5.2. *A (symmetric)* ***purgatorial equilibrium*** *is a subgame perfect equilibrium with the following properties:*

1. *All individuals choose the active market at birth that minimises c_{ij}. In subsequent periods, they may select new markets if it is profitable to do so.*

2. *Principals in market j maximally trust all present who have not shirked in the past and whose current attendance spell is at least t^*. They offer trust $\lambda^* \in [0, 1)$ to agents who have been present for less than t^* periods. They offer zero trust if the agent has shirked in the current joint attendance spell.*

3. *Agents work when trusted unless (i) either the principal or agent has shirked in the current joint attendance spell, or (ii) a level of trust other than λ^* is offered to an agent present less than t^* periods.*

4. *If an individual leaves a market, he chooses the active market that minimises transportation costs, and he subsequently plays as if that is the new home market.*

Unlike the identity investment equilibrium, each agent imposes an externality on all others in his market when he arrives, since they may be matched with him during purgatory. This force also affects the optimal number of players per market, since if there are only two players per market, a player who has lost his partner and received a new one sees the same payoffs as one who has completed just a year of purgatory. This player may be likely to leave himself with an adverse location shock, which imposes new purgatory costs on other agents. In contrast, when there are more players per market, each individual shares the cost of the new entrant's purgatory among all of the other senior members of the community. Thus, the presence of a group of individuals all committed to a market creates additional stability.

The following result relies on a parameter restriction in order to simplify the analysis, but it can be generalized with similar logic. For simplicity, it focuses on the case where purgatory can be only one period. Then, we establish the existence of purgatorial equilibria under the same conditions as identity investment equilibria. We then show that purgatorial equilibria are less efficient than identity investment equilibria, because withholding trust from junior community members entails costs for their partners, too. Finally, we show that efficiency increases in group size, because larger groups enable steeper incentives to seniority. For fixed turnover rates, the probability that a junior matches with a senior decreases with group size, while the probability that a senior matches with another senior increases with group size. That increases the gains to being senior.

Proposition 5.2. *Suppose that $w - R \geq \bar{c}$. (i) Whenever an identity investment equilibrium exists with specific investments $m = \bar{c}$ and where all individuals return to their initial home market in equilibrium, there exists a purgatorial equilibrium with the same market sizes and where all individuals return to their initial home market in equilibrium. All purgatorial equilibria in this class are less efficient than the corresponding equilibrium with specific investments.*
(ii) Efficiency in the class of purgatorial equilibria where all individuals return to their initial home markets increases in the community size (decreases in J^A).

Proof: When $w - R \geq \bar{c}$, purgatory of length one in each market is sufficient to induce individuals to return to their home markets for all market sizes. Let $N_m = N/J_A$ be the number of individuals per market.In each period, $(1-\delta)N_m$ individuals die, and in equilbrium all senior individuals return, so there will be $(1-\delta)N_m$ new members. Then, the value for a returning member is

$$W^R = \frac{1}{1-\delta}\left(w + R\frac{\delta N_m - 1 + \lambda^*(1-\delta)N_m}{N_m - 1} - \bar{c}/2\right).$$

The value for member beginning purgatory is

$$W^P = \lambda w + R\frac{\delta N_m + ((1-\delta)N_m - 1)\,\lambda^*}{N_m - 1} - \bar{c}/2 + \frac{\delta}{1-\delta}\left(w + R\frac{\delta N_m - 1 + \lambda^*(1-\delta)N_m}{N_m - 1} - \bar{c}/2\right).$$

Then, the cost an agent will pay in order to return is

$$W^R - W^P = (1-\lambda^*)\left(w - \frac{1}{N_m - 1}R\right),$$

which is increasing in N_m. An agent will return from the most distant market if $W^R - W^P \geq \bar{c}$, and so it is efficient (ignoring other constraints) to select λ^* such that

$$W^R - W^P = (1-\lambda^*)\left(w - \frac{1}{N_m - 1}R\right) = \bar{c}.$$

Note that when λ^* is determined this way, it is increasing in N_m. Note further that W^P is increasing in λ^*, and holding λ^* fixed, W^P is constant in N_m. W^R is increasing in N_m and in λ^*. Thus, social welfare is increasing in N_m. So it remains to check whether other constraints can be satisfied. When$W^R - W^P = \bar{c}$, the analysis of the shirking constraints is exactly identical to the proof of Proposition 5.1. (Note that agents in purgatory are trusted less than agents that have finished purgatory, yet their continuation values are the same.)

To compare efficiency with an identity investment equilibrium, note that an identity investment equilibrium requires a single wasteful investment of $\bar{c}$ for each newly born individual. The purgatorial equilibrium also requires that each newly born individual receives $\bar{c}$ less than W^R. However, W^R is lower than the value to returning in an identity investment equilibrium, where all individuals are fully trusted in each period, since W^R is reduced by the necessity of trading with individuals still in purgatory. □

We maintained the restriction $w - R \geq \bar{c}$ so that we could focus on the case where only one period of purgatory for all market sizes. Now suppose that $w < \bar{c}$, so that more than one period of purgatory is definitely required for all market sizes. Then, junior agents will not necessarily return in equilibrium, because the difference in continuation value from the first period to the second period of purgatory will be less than w. Without additional assumptions, the age distribution of each community varies over time, significantly complicating the analysis. However, the main ideas carry over. In terms of equilibrium existence, as long as purgatory is set so that $W^R - W^P = \bar{c}$, where W^R is the value from having finished purgatory and W^P is the value at the start of purgatory, and as long as agents receive zero trust while in purgatory, a purgatorial equilibrium exists under the same conditions as an identity investment equilibrium. If junior agents are not trusted, the fact that they are mobile does not affect their incentive constraints, but it does increase the equilibrium fraction of the population that is junior, which in turn reduces the aggregate trade in the economy as well as the value of being senior. The mobility saves in transportation costs, but a junior's movement has an externality: the junior agent's additional time in purgatory reduces trade for other members of the group that are matched with her.

Following the same logic as the proof of Proposition 5.2, increases in group size increase efficiency, by reducing the length of purgatory required to generate the gap $W^R - W^P = \bar{c}$.

5.3 Hierarchical Equilibria

Like the identity investment equilibrium, in the purgatorial equilibrium each market's enforcement of purgatory requires agents to refrain from trade purely for the purpose of supporting social conventions that help other markets sustain cooperation. We now explore the extent to which this feature can be relaxed.

We consider a "hierarchical structure" for each market to be a number of levels, L, and a vector of numbers, $(\gamma_1, \gamma_2, ...\gamma_L)$, where γ_l is the fraction of agents in market j at level l. We consider only cases where $\gamma_l N/J^A$ is an integer for each l, where J^A is the number of active markets. Empty places in the hierarchy are populated as follows. If any agents at level L have died, then agents from level $L-1$ are promoted, where if there are more agents at level $L-1$ than empty places, agents are selected randomly from level $L-1$. If there are still empty places at level L, the procedure is repeated at successively lower levels until the bottom level is reached. Next, if there are empty places in level $L-1$, an analogous procedure is used to fill them. Finally, new agents are born, and they fill in the remaining places in the hierarchy. To keep things simple, suppose that there are capacity constraints in each market, so that new agents will select another market if more than N/J^A agents are in attendance already. Also for simplicity, assume that $J = J^A$.

A 9. *Capacity constraints can be exogenously imposed on markets.*

In equilibrium, members of each level of the hierarchy are treated symmetrically, but different levels give and receive different levels of trust.

We will let $\lambda(\tau_i, \tau_k)$ be the equilibrium scale of trade between a principal i and an agent k, where τ_i is the level of the principal and τ_k is the level of the agent, when the demographic structure of that market is γ_{jt}. We let $q^*(\tau_i)$ be the probability that an agent with of level τ_i leaves market j after period t. $c^*(\tau_i)$ represents the maximum travel cost an agent is willing to incur to return to the home market.

Definition 5.3. *A **hierarchical equilibrium** is a subgame perfect Nash equilibrium with the following properties:*

1. **IC**$_L$: *Players attend the closest market j^* at birth- henceforth referred to as the home market. In each subsequent period t, a player with level τ_i continues to attend market j^* as long as the travel is less than $c^*(\tau_i)$. Otherwise, the player goes to the closest market and treats it as the new home, starting at the lowest level.*

2. **IC**$_P$: *as principals, in period t in market j players with seniority τ_i trade with a trading partner with seniority τ_k at the scale $\lambda^*(\tau_i, \tau_k) \in [0, 1]$, unless cheated by an agent, in which case they trade at scale $\lambda = 0$ in subsequent trade with that agent in the same joint attendance spell. When a new spell begins, they behave as if the cheating never took place.*

3. **IC**$_A$: *as agents, players do not cheat. If they cheat a player once, they cheat throughout the remainder of the spell. When a new spell begins, they behave as if the cheating never took place.*

4. **IC**$_{AL1}$: *If a player of level τ_i cheats another, in the subsequent period the player attends the closest market j' such that $j' \neq j^*$ and thenceforth plays as if j' was the player's home market.*

5. $\mathbf{IC}_{AL2}$: *If a player of level τ_i is cheated by another, the player continues to select markets as specified in* $\mathbf{IC}_L$.

6. *Off-equilibrium strategies:*

   (a) *If at the end of period t a player with level τ_i has cheated one or more players that are present in his home market at period t, in the subsequent period $t+1$ the player attends the closest market j' such that $j' \neq j^*$ and thenceforth plays as if j' was the player's home market and as if he had never cheated.*

   (b) *If a player cheats one or more others and returns to his home market anyway, and sees that there are no players there he has ever cheated, he then cooperates with all present and subsequently behaves as if he had never cheated.*

   (c) *If a player cheats one or more others and returns to his home market anyway, and sees that there are players present that he has cheated in the past, then he cheats with anyone who trusts him and in the subsequent period $t+1$ the player attends the closest market j' such that $j' \neq j^*$ and thenceforth plays as if j' was the player's home market and as if he had never cheated.*

7. *The scales of trade $\lambda^*(\tau_i, \tau_k)$ have the following properties:*

   (a) $\tau_i > \tau_i' \Rightarrow \lambda^*(\tau_i, \tau_k) \geq \lambda^*(\tau_i', \tau_k)$

   (b) $\tau_k > \tau_k' \Rightarrow \lambda^*(\tau_i, \tau_k) \geq \lambda^*(\tau_i, \tau_k')$

The strategies specify a limited form of retaliation against shirking–agents only retaliate in the same attendance spell for the pair, otherwise the record of cheating is "erased." This allows us to avoid complex calculations of the probability that an individual leaves a market after cheating, but then by chance encounters the same agent again in a different market after the individual has left, or by chance returns to the same market after leaving another. The latter is not an issue when individuals stay in the same market forever in equilibrium, but it arises when individuals leave in equilibrium.

As in the identity investment and purgatorial equilibria, another restriction incorporated in these strategies is that individuals leave the home market after cheating, rather than returning. If we allowed individuals to return, then the individal that was cheated would find it optimal to leave more often in equilibrium, undermining the sanction that follows cheating. This might be particularly problematic when a senior agent cheats on a more junior agent, since the senior agent would rarely or never leave, while the junior agent has little to gain from staying. The restriction we impose is stronger than necessary, but it simplifies the analysis.

Before providing a formal analysis of the incentive constraints, we first provide some interpretation of how the equilibrium works. The increasing scales of trade with the agent's level implies that an agent's value for being in a market is increasing in his seniority. This provides an incentive to return. Taking as given this structure, agents lower in the hierarchy are more tempted to leave, and thus trusting them fully might lead to a cheat-and-leave defection. Thus, it is possible that in an individual pairing, the agent is trusted as far as possible.

Now consider a more formal analysis. To analyze this class of equilibria, we introduce notation for the value functions of agents $W(B_{it}, \tau_i)$, which depend on an agent's history of cheating behavior B_{it} and his level τ_i in the current home market. Our assumptions ensure that demographic structures are always the same across markets when players follow the specified strategies. We let $\beta_{l,l'}$ be the probability that an agent at level l transitions to level l' in a given period, conditional on surviving. Note that $\beta_{L,L} = 1$.

Now consider an analysis of the hierarchical equilibrium with a given hierarchy. We begin by noting that (incorporating IC_L) $c^*(\tau_i)$ is determined by

$$c^*(\tau_i) = \max(0, \min(\bar{c}, \textstyle\sum_{l=\tau_i}^{L} \beta_{\tau_i,l} W(\emptyset, l) - W(\emptyset, 1))).$$

This is the largest distance an agent would be willing to travel to stay in the same group, having never cheated. Of course, the value functions and transition probabilities themselves depend on c^* (through q^*). It then follows that the probability of returning is given by

$$q^*(\tau_i) = \frac{1}{J} \min\left(1, \max\left\{\alpha : \alpha \in \{1, 3, .., (J+1)\},\ \frac{(\alpha - 1)\bar{c}}{J} < c^*(\tau_i)\right\}\right).$$

We let the period payoffs from a given level of seniority be written

$$\Pi(\tau_i) = \frac{1}{N-1} \left(\begin{array}{c} N \sum_{l=1}^{L} \gamma_l [\lambda^*(\tau_i, l)R + \lambda(l, \tau_i)w] \\ -\lambda^*(\tau_i, \tau_i)(R + w) \end{array} \right).$$

Let the period payoffs from shirking today, having cheated an agent of (today's) level l' yesterday be

$$\Pi^s(\tau_i, l') = \frac{1}{N-1} \left(\begin{array}{c} N \sum_{l=1}^{L} \gamma_l \left(\lambda^*(\tau_i, l)R + \lambda^*(l, \tau_i)s\ \right) - \lambda^*(\tau_i, \tau_i)(R + s) \\ -[\lambda^*(l', \tau_i)R + \lambda^*(\tau_i, l')s] \end{array} \right)$$

Now consider the constraints that must be satisfied for the specified strategies to be best responses. Begin with the constraint that says individuals prefer cooperation to cheating and leaving ($\mathbf{IC}_A$). The average cost of returning to the home market, conditional on returning from the closest $J \cdot q^*(\tau_i)$ markets, can be calculated as

$$\hat{c}(\tau_i) = 1\{q^*(\tau_i) < 1\} \cdot \frac{(q^*(\tau_i))^2 J^2 - 1}{2J^2 q^*(\tau_i)} \bar{c} + 1\{q^*(\tau_i) = 1\} \cdot \frac{\bar{c}}{2},$$

while if a player leaves he will attend the closest market in equilibrium. The expected cost of travel to attend the closest market while avoiding one particular market (after cheating) is $\frac{1}{J} \cdot \frac{2\bar{c}}{J}$. Then, the constraint is

$$\lambda^*(\tau_k, \tau_i)(s-w) \leq \delta \left(\begin{array}{c} q^*(\tau_i)\left(\sum_{l=\tau_i}^{L} \beta_{\tau_i,l} W(\emptyset, l) - \hat{c}(\tau_i)\right) + (1-q^*(\tau_i))W(\emptyset,1) \\ -\left(W(\emptyset,1) - \frac{2\bar{c}}{J^2}\right) \end{array} \right). \tag{IC$_A$}$$

Finally, we need to consider the constraint (IC$_{AL}$) that deters agents from returning to the home market after cheating. This constraint is most binding when the principal cheated was junior, since a junior player is most likely to leave and also had the least potential trade to withhold as a sanction, and when the agent who cheated is senior, since that agent has the most to gain from returning. This constraint specifies that even when transportation costs to the new market are high ($2\bar{c}/J$) (that is, when the agent lands in the original home market), it is still better to start over rather than return to the home market. The payoffs in the home market depend on whether the principal that was cheated in the past has left. If the principal has left, the individual continues as if he had never cheated. If the principal is present, then the strategies specify leaving tomorrow. The transportation cost in that case is incurred with probability $1/J$, which is the probability that the agent lands in the original home market. Let $\xi_{1,L}$ be the probability that a player of level 1 survives and returns, conditional on one other player of level L surviving and being initially placed at the home market (recall that survival and locations are not independent across players). Let $\tilde{\beta}_{1,l}$ be the probability that a player of level 1 transitions to level l, conditional on surviving and on one other player of level L surviving and being placed in the home market (which implies that one position in level L in the home market is definitely full).

$$W(\emptyset,1) - 2\bar{c}/J \geq (1-\xi_{1,L})W(\emptyset,L) + \xi_{1,L}\left(\sum_{l=1}^{L} \psi_{1,l}\Pi^s(\tau_i,l) + \delta \cdot \left(W(\emptyset,1) - 2\bar{c}/J^2\right)\right). \tag{IC$_{AL}$}$$

Note that off of the equilibrium path, strategies specify that an individual leaves after returning and finding the individual she cheated is still there. If (IC$_{AL}$) holds, then she will also be willing to leave in that contingency.

Clearly, this model is quite complex, and we have not succeeded in finding closed form solutions for the value functions in the case where players leave in equilibrium. When players do not leave in equilibrium, it is possible to simplify the above expressions and write the value functions and incentive constraints explicitly in terms of primitives, but the expressions are cumbersome.

In either case, it is possible to numerically calculate whether or not a hierarchical equilibrium exists for a given hierarchical structure and set of parameter values. We have done this for two- and three-level hierarchies. We find that typically there are many equilibria with different amounts of leaving in equilibrium. The lower the level of trust at each level of the hierarchy, the more individuals wish to leave; but the more they wish to leave, the less trustworthy they are. Thus, there are low-trust, high leaving equilibria and high-trust, low-leaving equilibria co-existing for the same parameter values.

Example Parameter values: δ =0.8, $J = J^A = 130$, $N = 1300$, $w = 10$, $s = 10.3$, $R = 50$, $r = 1$, $\bar{c} = 1.44$. Note that $s - w > \frac{\delta^2}{1-\delta^2}\frac{1}{N-1}(R+w)$, so that the peddler's equilibrium does not exist. There exists a hierarchical equilibrium with a two-level hierarchy, with four places in the lower level and six in the upper level. Upper level agents are fully trusted, while lower level agents are trusted at scale $\lambda = .77$. The no-shirking constraint (IC_A) binds for lower level agents. The probability that an individual of level one who returns to their home market transitions to level two is .76. Low level agents return to their home market with probability .64, while high level agents return with probability .84. In this equilibrium, individuals strictly prefer to leave after cheating–the individual markets are small enough, and individuals return frequently enough, to provide a strong sanction for cheaters if they return.

In the example, requiring the incentive constraints for low level agents bind puts a non-trivial tax on efficiency: more efficient equilibria exist with higher levels of trust for the lower level of the hierarchy, together with more frequent market-changes and the associated lower transportation costs.

Summarizing, the hierarchical equilibrium is attractive because the social conventions are less artificial, and it is possible to construct equilibria where each pair trusts as much as possible, taken as given mobility. The social hierarchy induces higher mobility for young agents, which renders them less trustworthy.

6 Choice-Based Hierarchical Equilibria

Now consider modifying the model so that, once players have arrived in a market, they can freely choose among partners. This modeling choice captures a different type of interaction than with the random matching model, and may be more appropriate for some kinds of trading environments than others.

Assume that each individual can act as agent to many others, but as principal to only one other player. Assume that matches and trades are private information, so that one individual cannot discern the absence or presence of a match in other agents. However, there exists a public randomization device that can be used to coordinate principals that are selecting among the same set of agents, so that the trades can be spread out among them rather than "accidentally" concentrated.

We begin by considering how the equilibria we have studied already would change when individuals can select among partners within a community. Suppose that when indifferent among potential partners, individuals randomize. For the case of the peddler's equilibrium, where individuals always attend the closest active market, then $J^A = N/2$ implies that there is no effect of the ability to select partners, since each market only has two participants. With $J^A < N/2$, the peddler's equilibrium would break down for parameter values where it could be sustained under random matching. To see why, note that if individual i cheats on individual j, if individuals i and j happen to meet again in the future, individual i can avoid selecting individual j. There is still some sanction, however, since j will not select i to be an agent.

Now consider modifying the purgatorial and hierarchical equilibria, where the ability to choose partners produces more substantive changes. Principals will strictly prefer to select senior individuals as agents, rather than junior ones, since junior agents can not be trusted as much and therefore generate less surplus. Thus, the central source of inefficiency in the purgatorial and hierarchical equilibria disappears. However, as we proceed to show, a number of incentive constraints are modified.

Let us focus on the hierarchical equilibrium, since it can be analyzed most easily when individuals leave in equilibrium. We define a **choice-based hierarchical equilibrium** to be the same as a hierarchical equilibrium, with the following modifications. Principals select agents as follows: in period t in market j players with seniority τ_i always select an agent that she has not cheated on in the current spell, and among those agents, they select randomly among the individuals with the highest level of seniority. We assume that the randomization is coordinated, so that trades are spread evenly (up to integer constraints) among the seniors. Note that our modifications imply that if an individual i cheats a principal j, i is still willing to select j to be an agent in the future, as if no cheating had taken place. In contrast to the baseline model, in the choice-based model this does not reduce the sanction to individual i from cheating as long as there are three or more individuals at the highest level, since individual i would have been able to avoid matching with individual j at zero cost anyway. With this modification, individual j has no disincentive to return to the market, even if individual i returns.

Now consider an analysis of the choice-based hierarchical equilibrium with a given hierarchy. The definitions of $c^*(\tau_i)$ and $q^*(\tau_i)$ are unchanged. The period payoffs from a given level of seniority become

$$\Pi(\tau_i) = \lambda^*(\tau_i, L)R + 1\{\tau_i = L\} \sum_{l=1}^{L} \frac{\gamma_l}{\gamma_L} \lambda^*(l, L)w.$$

The period payoffs from shirking today, having shirked on k agents of (today's) levels $(l_1, .., l_k)$ in the same spell, when l^* is the highest level of a potential agent that has not been cheated, are

$$\Pi^s((l_1,..,l_k),\tau_i) = \lambda^*(\tau_i,l^*)R + 1\{\tau_i = L\}\left(\sum_{l=1}^{L}\frac{\gamma_l}{\gamma_L}\lambda^*(l,L) - \frac{1}{\gamma_L N}\sum_{s=1}^{k}\lambda^*(l_k,L)\right)s.$$

Now consider the constraints that must be satisfied for the specified strategies to be optimal. As in the baseline model, $\mathbf{IC}_L$ defines $c^*(\tau_i)$. The constraint that says individuals prefer cooperation to cheating and leaving ($\mathbf{IC}_A$) is unchanged from the baseline model, as is the constraint that deters agents from returning to the home market after cheating (IC_{AL}).

If an equilibrium exists where $\lambda^*(\tau_i, L) = 1$ for all τ_i, and if $\gamma_L N > 3$, the payoff expressions simplify to

$$\begin{aligned}\Pi(\tau_i) &= R + 1\{\tau_i = L\}\sum_{l=1}^{L}\frac{\gamma_l}{\gamma_L}w,\\ \Pi^s((l_1,..,l_k),\tau_i) &= R + 1\{\tau_i = L\}\left(\frac{1}{\gamma_L} - \frac{k}{\gamma_L N}\right)s.\end{aligned}$$

Now consider the efficiency of this equilibrium, if it exists. First, every principal can find an agent who is trustworthy, and so trade is fully efficient. Second, a player leaving no longer has any externality, because there will always be senior members of the community who can be fully trusted. However, the travel induced by the seniority structure is socially costly. Thus, the most efficient equilibrium will minimize travel, subject to being able to maintain cooperation.

Let us compare the constraints for this potential equilibrium to those for the (non-choice-based) hierarchical equilibrium. An important difference is that senior agents are now trusted by $1/\gamma_L$ players (for simplicity suppose this is an integer). This increases the incentive to shirk, since the one-period gain is larger relative to the future cost of starting over. It also increases the incentive to leave after cheating, since there are more individuals in the community that have been cheated.

Thus, IC_A for the seniors becomes

$$(1/\gamma_L)\lambda(\tau_L,\tau_L)(s-w) \leq \delta\left(\begin{array}{c} q^*(\tau_L)\left(W(\emptyset,L) - \hat{c}(\tau_i)\right) + (1-q^*(\tau_L))W(\emptyset,1) \\ -\left(W(\emptyset,1) - \frac{2\bar{c}}{J^2}\right)\end{array}\right). \qquad (\text{IC}^c_A(L))$$

Note that a junior, if trusted at level λ' by a single principal, would have the following constraint to ensure cooperation:

$$\lambda'(s-w) \leq \delta\left(\begin{array}{c} q^*(\tau_L)\left(\sum_{l=\tau_i}^{L}\beta_{\tau_i,l}W(\emptyset,l) - \hat{c}(\tau_i)\right) + (1-q^*(\tau_L))W(\emptyset,1) \\ -\left(W(\emptyset,1) - \frac{2\bar{c}}{J^2}\right)\end{array}\right). \quad (\text{IC}^c_A(1))$$

Of course, the strategies we have specified entail no trust for the juniors, and we can always specify strategies whereby a junior shirks if trusted. Thus, our interest in this constraint is that of the robustness concerns raised above: we have argued that it is of interest to construct equilibria where trade is not withheld "artificially." Note that the right-hand side of ($\text{IC}^c_A(1)$) is smaller than that of ($\text{IC}^c_A(L)$), but at $\lambda' = \lambda(\tau_L, \tau_L)$, the left-hand side is smaller. Thus, it is not necessarily true that juniors are less trustworthy in equilibrium. If ($\text{IC}^c_A(L)$) binds with $\lambda(\tau_L, \tau_L) < 1$ but ($\text{IC}^c_A(1)$) is satisfied with $\lambda' > \lambda(\tau_L, \tau_L)$, we can modify the definition of the equilibrium so that junior agents receive some trades, for example by creating multiple "low" levels where some of these are designated to trade with other low level individuals. In equilibrium, these individuals would cheat if they received extra trades. For simplicity, we focus our exposition on the case where ($\text{IC}^c_A(L)$) is satisfied at $\lambda(\tau_L, \tau_L) = 1$ while $\text{IC}^c_A(1)$ fails at $\lambda' = 1$.

Let K be a set of individuals of cardinality $1/\gamma_L$ with the lowest possible aggregate seniority. IC_{AL} becomes

$$W(\emptyset, 1) - 2\bar{c}/J \geq \Pr(\text{none of } K \text{ return})\, W(\emptyset, L) \qquad (\text{IC}^c_{AL})$$
$$+ \sum_{K' \subseteq K, K' \neq \{\}} \Pr \begin{pmatrix} K' \text{ return and transition} \\ \text{to levels } (l_1, .., l_{k'}) \end{pmatrix} \begin{pmatrix} \Pi^s((l_1, .., l_{k'}), \tau_i) \\ +\delta \cdot (W(\emptyset, 1) - 2\bar{c}/J^2) \end{pmatrix}.$$

So far, we have described forces that tighten IC_A and slacken IC_{AL}. We next describe forces that work in the opposite direction. Because all trade is concentrated on seniors, relative to the non-choice-based hierarcy, the choice-based hierarchy will have greater benefit to returning to the home market in equilibrium, implying that $c^*(\tau_i)$ is higher. This will relax the cheat-and-leave constraint (IC_A): the individual gives up more by starting over. On the other hand, the constraint that requires an individual to leave after cheating (IC_{AL}) is tightened, for two reasons. First, there is more to lose by starting over. Second, the sanction within a market is lower, since the individual can always find an agent to trade with after cheating (although it may be a more junior agent). Indeed, if both the cheater and the "cheatees" remain junior after the cheating happens, there is no sanction at all in the home market.

If the leave-after-cheating constraint is not satisfied, we can consider modifying the strategies so that individuals do not always leave after cheating. We did not allow this in the baseline model because it would potentially induce an individual who has been cheated upon to leave, destroying the sanction to the cheater. As discussed above, however, this effect disappears in the choice-based model. Then, we could replace ($\mathbf{IC}_A$) and ($\mathbf{IC}_{AL}$) with a single, new constraint that specifies that cooperating is better than optimizing between leaving and returning. For simplicity, we will not formally analyze that case.

Then, the problem of finding the most efficient equilibrium reduces to the problem of finding a hierarchical structure (market size and proportion at the top level) that minimizes transportation costs subject to ($\mathbf{IC}_A$) and ($\mathbf{IC}_{AL}$). Of course, the more individuals move markets in equilibrium, the lower will be the sanction from cheating, since both the individual who was cheated as well as the individual who was not will be more likely to move on.

We analyze this model numerically, illustrating parameter values where it exists, as well as comparative statics on how the shape of the hierarchy affects efficiency and existence of equilibrium. The results of this analysis will be included in future versions of the paper.

We conclude by describing the desirable robustness properties of choice-based hierarchical equilibria. First, it is possible for some parameter values to construct equilibria whereby no principal withholds trade "artificially." Junior agents may not be fully trustworthy, because they have less to lose from a cheat-and-leave strategy. If so, principals have a strict preference to trade with senior agents, as desired.

As discussed above, finding equilibria with this feature is in some ways easier, and in some ways more difficult than for the non-choice-based hierarchy. Since junior agents do not serve as agents in equilibrium, the level of trust that could be given to junior agents does not directly affect value functions or incentive constraints. So long as they are not fully trustworthy for a single trade, others will not wish to trade with them. On the other hand, as discussed above, since principals receive more than one trade it is possible that their incentive constraints bind first.

A second robustness property concerns whether cooperation in a given market will be upset by the emergence of another market that does not have the same institutions. The choice-based hiearchy has the attractive feature that being a senior in a given market is more profitable than being in a community without a seniority structure, but where all members are fully trusted. Thus, seniors have something to lose from a cheat-and-leave strategy, even if another community trusts them immediately. In addition, if all communities have a seniority structure and IC_A is binding, then cooperation will not be sustained in a new institution-free community, because the benefit to returning to the new community is lower than the benefit to returning to hierarchical communities. This feature of the choice-based hierarchy contrasts with the identity investment, purgatorial, or hierarchical equilibria, in which the value to the most senior members of a market is just equal to that from a community that fully trusts its members immediately.

7 Existing Theoretical Perspectives

Given its importance, it is not surprising that research on engendering trust has a long and venerable tradition. In this section we relate our results to the existing literature.

Beyond the classic folk theorems, theorists have traditionally focused on two types of mechanisms that overcome the trust problem: those that signal reputations and those that require third party enforcement. A number of recent studies have focused on the importance of sacrifice and specific investments. Our analysis differs from the prior literature in its focus on group formation, group size, hierarchical structures, and the problems associated with increasing population size.

Reputation models provide a useful theory of how intrinsically honest or dishonest individuals can use costly signals to establish a reputation for honesty. Sobel [37] and, more recently, Watson [40] consider reputation models where the scale of trade increases over time. In a model that focuses more directly on outside options, Ghosh and Ray [18] examine a case where there are two types of individual - myopic opportunists and the more patient.[5] Individuals can continue to interact with the same partner should they so choose. In order to sustain cooperation, partners spend a period "trying each other out," using this as a screen for myopic opportunism. They find that societies with very few myopic members are in fact sometimes less likely to cooperate, as there the barrier to establishing new relationships is reduced.

In all of the reputation-based models, players learn about the intrinsic type of their partners as the relationship progresses. In our study, players are intrinsically symmetric, but the extent to which a player is trustworthy changes over time and is determined in equilibrium.

In other studies, third parties or groups assume the role of contract enforcement. Milgrom, North and Weingast [34] provide two examples underlying the revival of trade in late medieval Europe- the Law Merchant system of private adjudication and annual trade fairs, where justice was administered by the Counts of Champagne. Dixit [10] analyses a recent example, based on Gambetta's [17] study of Mafia practice.

Kandori [25] and Greif [20],[21],[22] analyze enforcement within groups that transmit information to each other about defectors. These defectors can then be punished by others in the group- whether through ostracism or persistent betrayal.

Dixit's [11] study of trade expansion and enforcement is similar in motivation to ours. He uses a circular world as a geographical analogy for the costliness of information flows across distances, and examines how much cooperation can be sustained as the circle grows in size. He finds that small and large worlds can sustain greater trust than their

[5] Also see Ghosh and Ray [19] for an application to micro-credit.

intermediate counterparts. In small worlds, partners to a transaction are likely to know third parties in common, and thus are able to share information about defectors. In large worlds, developing a legal system becomes economical. In contrast, our model focuses on the case where information sharing about behavior within bilateral relationships is not possible.

Klein and Leffler [26]'s seminal study analyzed the role of specific investments in a model of repeated interaction. Iannaccone [23] applies the notion of specific investments to cults that provide club goods. There is an incentive to free-ride upon others' zealousness. In order to limit participation to the truly committed, religious practices, such as stigma and self-sacrifice develop to act as screening devices. Similarly, Fryer [15] examines the issue of cultural identity and particularly of blacks " acting white."

In a series of papers, Akerlof [2], Akerlof and Kranton [3], [4] introduce social custom and group "identity" into individual's utility functions. These identities- and the concomitant "prescriptions" for behaviour that they imply- result in individual and group sanctions for violators of the appropriate group "code of conduct." Their analyses take both the set of identities and prescriptions to be exogenous.

Other studies have focused on individual partnerships rather than groups, with barriers to entry being raised to switching partners. For example, Kranton [28] attempts to find the conditions under which cooperation can be sustained in bilateral partnerships even if an alternative partner is present. She suggests the use of *bonds*- initial losses of disutility- can act to discourage partner-hopping by making it costly to form new relationships. A similar point appears in Carmichael and MacLeod [8].

Friedman and Resnick [14] analyze cooperation in online communities, and look for a force to deter the relatively costless assumption of new identities by individuals on the Internet. They argue that if the incumbent members of a chat room treat new members poorly for a requisite period of time, this will discourage fraudulent use of new identities. This idea is closely related to the purgatorial equilibrium that we analyze. Our analysis considers alternatives to purgatorial equilibrium that may be more robust, and in addition we analyze the role of group size and structure.

Lindsey, Polak and Zeckhauser [31] incorporate the notion of itinerant temptations into their study of long-term bilateral, exclusive relationships between individuals where there is "free love"- no barriers to a new start. Instead, existing relationships gain value the longer they exist. They contrast two "conventions"- a rule where bilateral relationships break up whenever either party cheats on the other and one where they only break up when both do so simultaneously. An interesting feature of their model is that cheating occurs along the equilibrium path. In a similar manner to the hierarchical equilibrium we will discuss, relationships are robust due to their own increasing value over time. However, this occurs for different reasons in our model- in the social hierarchies we construct, senior agents engender more trust.

In a related paper, Sobel [38] analyzes a model where agents in a large population form bilateral relationships. He contrasts relational contracts with formal contracting as mechanisms for sustaining trust. As in our model and that of Lindsey, Polak and Zeckhauser [31], it can be inefficient for partnerships to be exclusive in every period. Sobel's [38] model has relationships that permanently "grow stale." In his model, relationships based on relational contracts may last inefficiently long, because the institutions that support cooperation must entail costs of starting new relationships. This is similar to our result that there is inefficient travel in equilibria that require loyalty.

In summary, our models differs from the existing literature in several respects. First and most importantly, we focus on endogenous group formation, and we analyze alternative institutions that might arise in a group setting. It is costly for players to maintain loyalty to a group. Individuals make inferences about the trustworthiness of potential matches by the loyalty the individual has shown to the group in the past, which increases over time despite the intrinsic symmetry of individuals. An individual can maintain his seniority status independent of whether a particular partnership breaks down through cheating or death of the partner. We do not rely on group sanctions to deter cheating, since individual cooperation is observable only to the parties directly involved, and our model seeks to minimize the need for individuals to withhold trade from trustworthy partners simply for the purpose of upholding social conventions. Instead, social conventions about trust are related to the equilibrium trustworthiness of individuals. Our model provides a starting point for analyzing further questions about social hierarchies within groups.

8 Conclusions and Applications

This paper has analyzed the endogenous formation of groups, when outside options are strong and legal enforcement is not available. We have constructed a number of alternative equilibria, and we have discussed their strengths and weaknesses as well as robustness properties. A common theme is that institutions support cooperation by creating a disincentive to change markets. Seniority structures provide this disincentive without necessarily requiring artificial barriers to entry or withholding of trade. Seniority structures may also be robust to the entry of institution-free markets.

The insights from our analysis can be used to help understand how cooperation is sustained in a number of real-world settings. For example, consider online communities, such as open-source software. These communities often have formal or informal seniority structures that provide incentives for loyalty. Our model suggests that those with robust seniority structures may maintain the loyalty of senior members even in the presence of newly formed online communities with weak institutions.

References

[1] Jaleel Ahmad. Factor market dualism, small scale industry and labor absorption. *Journal of economic development*, 25(1), 2000.

[2] George A. Akerlof. A theory of social custom, of which unemployment may be a consequence. *Quarterly Journal of Economics*, 94:749–95, June 1980.

[3] George A. Akerlof and Rachel E. Kranton. Economics and identity. *Quarterly Journal of Economics*, 115(3):715–53, 2000.

[4] George A. Akerlof and Rachel E. Kranton. Identity and schooling: some lessons for the economics of education. *Journal of Economic Literature*, 40:1167–1201, December 2002.

[5] Kenneth J. Arrow. *The limits of organization*. Norton, New York, NY, 1st edition, 1974.

[6] Abhijit V. Banerjee and Andrew F. Newman. Information, the dual economy and development. *Review of Economic Studies*, 65(225):631–653, 1998.

[7] Gary S. Becker and George J. Stigler. De gustibus non est disputandum. *American Economic Review*, 67(2):76–90, March 1977.

[8] H. L. Carmichael and W.B. Macleod. Gift giving and the evolution of cooperation. *International Economic Review*, 38:485–508, 1997.

[9] M. Corominas-Bosch. *On two-sided network markets*. PhD thesis, Universitat Pompeu Fabra, 1999.

[10] Avinash K. Dixit. On modes of economic governance. *Econometrica*, 71(2):449–481, March 2003.

[11] Avinash K. Dixit. Trade expansion and contract enforcement. *Journal of Political Economy*, 111(6):1293–1317, December 2003.

[12] Jean Ensminger. *Making a market*. Cambridge University Press, Cambridge, UK, 1992.

[13] Hadi S. Esfahani and Djavad Salehi-Isfahani. Effort observability and worker productivity: Towards an explanation of economic dualism. *Economic Journal*, 99(387):818–36, 1989.

[14] Eric Friedman and Paul Resnick. The social cost of cheap pseudonyms. *Journal of economics and management strategy*, 10(2):173–199, 2001.

[15] Roland G. Fryer Jr. An economic approach to cultural capital. mimeo, University of Chicago, October 2002.

[16] Roland G. Fryer Jr. and Steven D. Levitt. What's in a name. 2003.

[17] Diego Gambetta. *The Sicilian Mafia: the business of protection*. Harvard University Press, Cambridge, MA, 1993.

[18] Parikshit Ghosh and Debraj Ray. Cooperation in community interaction without information flows. *Review of Economic Studies*, 63(3):491–519, July 1996.

[19] Parikshit Ghosh and Debraj Ray. Information and enforcement in informal credit markets,. mimeo, November 2001.

[20] Avner Greif. Reputation and coalitions in medieval trade: evidence from the Maghribi traders. *Journal of Economic History*, 49:857–82, December 1989.

[21] Avner Greif. Contract enforceability and economic institutions in early trade: the Maghribi traders' coalition. *American Economic Review*, 83(3):525–547, June 1993.

[22] Avner Greif. Cultural beliefs and the organisation of society: a historical and theoretical reflection on collectivist and individualist societies. *Journal of Political Economy*, 102(5):912–50, October 1994.

[23] Laurence R. Iannaccone. Sacrifice and stigma: reducing free-riding in cults, communes and other collectives. *Journal of Political Economy*, 100(2):271–291, April 1992.

[24] Matthew O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74, 1996.

[25] Michihiro Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59:63–80, 1992.

[26] Benjamin Klein and Kenneth B. Leffler. The role for market forces in assuring contractual performance. *Journal of Political Economy*, 89(4):615–41, 1981.

[27] Rachel Kranton and Deborah Minehart. A theory of buyer-seller networks. *American Economic Review*, 3:485–508, June 2001.

[28] Rachel E. Kranton. The formation of cooperative relationships. *Journal of Economics, Law and Organisation*, 12(1):214–33, April 1996.

[29] Murray Last. Some economic aspects of conversion in hausaland (nigeria). chapter 13, pages p.236–247. Holmes and Meier Publishers, Inc., New York, NY, 1979.

[30] Joshua Lerner and Jean Tirole. The economics of technology sharing: open source and beyond. *Journal of Economic Perspectives*, 19(2):99–120, Spring 2005.

[31] John Lindsey, Benjamin Polak, and Richard Zeckhauser. Free love, fragile fidelity and forgiveness: rival social conventions under hidden information. mimeo, October 2003.

[32] John McMillan and Christopher Woodruff. Dispute prevention without courts in Vietnam. *Journal of Law, Economics and Organization*, 15:637–58, October 1999.

[33] John McMillan and Christopher Woodruff. Interfirm relationships and informal credit in Vietnam. *Quarterly Journal of Economics*, 114:1285–1320, November 1999.

[34] Paul R. Milgrom, Douglass C. North, and Barry R. Weingast. The role of institutions in the revival of trade: the Law Merchant, private judges, and the Champagne fairs. *Economics and Politics*, 2(1):1–23, March 1990.

[35] Eric S. Raymond. *The Cathedral and Bazaar: musings on linux and open source by an accidental revolutionary*. O' Reilly, Cambridge, MA, 2001.

[36] Sonali Shah. Motivation, governance and the viability of hybrid forms in open source software development. *Management Science*, 52(7):1000–1014, 2006.

[37] Joel Sobel. A theory of credibility. *Review of Economic Studies*, 52(4):557–73, October 1985.

[38] Joel Sobel. For better or forever: formal versus informal enforcement. *Journal of Labor Economics*, 2006.

[39] Hernando De Soto. *The Other Path: the invisible revolution in the Third World*. Harper Collins, New York, 1989.

[40] Joel Watson. Starting small and renegotiation. *Journal of Economic Theory*, 85:52–90, 1999.