



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Constructing Social Division to Support Cooperation:
Theory and Evidence from Nepal

James Choy

No 1011

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Constructing Social Division to Support Cooperation: Theory and Evidence from Nepal*

James P. Choy[†]

August 7, 2013

Abstract

Many societies are divided into multiple smaller groups. The defining feature of these groups is that certain kinds of interaction are more likely to take place within a group than across groups. I build a model in which group divisions are enforced through a reputational penalty for interacting with members of different groups. Agents who interact with members of different groups find that they can support lower levels of cooperation in the future. The model explains why agents may be punished by the other members of their group for interacting with members of different groups and why agents are punished for interacting with members of some groups but not others. I test the empirical implication that there should be less cooperation among members of groups that make up a larger percentage of their communities. I discuss the origin and possible future of social division.

Keywords: Cooperation, Caste, Social Institution

JEL Classification Numbers: C7, O12, O17

*I am grateful to Mark Rosenzweig, Larry Samuelson, and especially Chris Udry for their advice and support throughout this project. Treb Allen, Priyanka Anand, David Berger, Gharad Bryan, Avinash Dixit, Tim Guinnane, Melanie Morten, Sharun Mukand, Joe Vavra, and various seminar participants provided helpful comments. I acknowledge research funding from the Yale Economic Growth Center.

[†]Department of Economics, University of Warwick. E-mail: j.choy@warwick.ac.uk

1 Introduction

Many societies are divided into multiple smaller groups. These divisions are especially salient in many developing countries, where the groups have names such as castes, tribes, or clans, but developed countries are divided as well, for example by race and religion. One stylized fact about group divisions is that people are more likely to interact in certain ways with members of their own groups than with members of different groups. Interactions that take place primarily within groups include long-distance trade (Greif 1993), mutual insurance (Grimard 1997, Munshi and Rosenzweig 2009, Mazzocco and Saini 2012), and job referrals (Munshi and Rosenzweig 2006). At first glance this fact is puzzling, because the argument from the gains from trade suggests that people should seek to interact with the most diverse possible range of partners. In this paper I argue that people may have a reputational reason to avoid interacting with members of different groups, even when such interactions would be profitable in the absence of the reputation mechanism. I then test some of the empirical predictions of the model using data on interactions between caste members in rural Nepal.

In my model, agents search over the community to find partners for cooperative relationships. If an agent cheats in any relationship, then the relationship breaks up and both partners to the relationship must search for a new partner. Agents who search over a larger proportion of the community pay a lower search cost. Cooperation is maintained by the threat that any cheating agent will have to pay this search cost, and the level of cooperation that any agent can support is inversely related to the search cost that the agent is expected to incur at the end of the relationship. Thus an agent who is expected to search over a larger proportion of the community can support a lower level of cooperation. Each agent is also a member of a payoff-irrelevant group, and in equilibrium each agent searches for partners only within her own group. If an agent is observed to have interacted with a member of a different group in the past, then it is believed that the agent will continue to search for matches within that group as well as within her own group in the future. Thus, agents who are observed to have interacted with members of different groups in the past are able to support lower levels of cooperation in the future. This punishment for interacting with members of different groups is sufficient to prevent members of different groups from interacting in equilibrium, even though in the absence of reputational incentives interactions between members of different groups would be just as profitable as interactions between members of the same group. I refer to this state of affairs as group segregation. Group segregation increases the level of cooperation that each agent can support compared to the situation without segregation, and if the benefits of cooperation are sufficiently important, then group segregation is welfare improving for the community as a whole.

The idea that the division of a community into multiple payoff-irrelevant groups can help to support

cooperation and thus increase welfare first appears in Eeckhout (2006). My theoretical contribution is the reputation mechanism that punishes agents for interacting with members of different groups. This mechanism ensures that two matched agents from different groups cannot profit by jointly deviating from the equilibrium to cooperate with each other. At a technical level, the mechanism ensures that the equilibrium is renegotiation-proof. This equilibrium concept was first applied to search and matching models by Ghosh and Ray (1996).

The reputation mechanism yields novel theoretical insights. The first insight is that group members can be punished by the other members of their group for interacting with members of different groups. Many real groups inflict these kinds of punishments on their members. In the South Asian caste system, the common understanding about this system of punishments is expressed in terms of ritual purity and pollution. Caste members are prohibited from marrying, socializing with, or in some cases even touching members of certain other castes. A caste member who violates these rules is said to be polluted, and must undergo a period of ritual purification before being allowed to regain full membership in the caste community.¹ Other societies in other places have also developed similar systems.²

A second insight is that the reputation mechanism endogenously generates an asymmetry between groups. Consider two groups, group 1 and group 2. If a member of group 1 interacts with a member of group 2, then it is believed that the member of group 1 will continue to interact with members of both group 1 and group 2 in the future, and this lowers the level of cooperation that the member of group 1 can support. For this reason, members of group 1 are not willing to interact with members of group 2. However, in order for this expectation to be rational, it must be the case that members of group 2 are willing to interact with members of group 1. Since members of group 1 are not willing to interact with members of group 2, if a member of group 2 does somehow manage to interact with a member of group 1 then it will be expected that the member of group 2 will not be able to continue interacting with members of group 1 in the future. As a result, members of group 2 who have interacted with members of group 1 can support the same level of cooperation as members of group 2 who have not interacted with members of group 1. Thus, group segregation is fundamentally asymmetric, with segregation being enforced by only one out of each pair of groups. This asymmetry seems to correspond to the real asymmetry in the South Asian caste system, in which members of higher castes are deemed to be polluted for interacting with members of lower castes, but

¹See, e.g., Dumont (1970) for further discussion of these rules. Höfer (2004) discusses rules specific to the Nepalese caste system.

²Berman (2000) discusses social norms among ultra-orthodox Jews in Israel. He argues that the extremely strict interpretation of the religious law among the ultra-orthodox effectively makes it impossible for the ultra-orthodox to interact in certain ways even with other Jews without breaking the law. Weyrauch and Bell (1993) study traditional law in Gypsy communities in the United States. They report that one punishable offense in Gypsy law is “familiarity with the *gaje*”, i.e. non-Gypsies. Kraybill (2001) discusses social norms among the Amish, including the social punishments inflicted on members of the Amish community who are deemed “worldly” for interacting too closely with outsiders. Austen-Smith and Fryer (2005) argue that African-American students face ostracism by their peers if they are accused of “acting white”.

members of lower castes are not deemed to be polluted for interacting with members of higher castes.

An empirical prediction of the model is that there should be less cooperation between members of groups that make up a larger proportion of their communities. The reason is that it is easier for members of larger groups to find new matches if their current matches break up, and so the penalty for cheating in any given relationship is lower, which implies that the level of cooperation that can be supported is lower as well. I test this prediction by studying borrowing and lending between caste members in rural Nepal. My prediction is that households whose castes make up a larger percentage of the community should engage in less borrowing and lending to and from members of the same caste. I also show that my model can be distinguished empirically from a number of related models in the literature.

My theory offers a new perspective on the origin and possible future of social division. Most economic theories of social division hypothesize that social divisions are defined by fundamental differences in preferences or technologies between members of different groups, and that these differences make interactions between members of different groups unprofitable. In the political science literature, these kinds of divisions are referred to as “primordial” divisions.³ In contrast, in my theory social divisions arise in equilibrium between *ex ante* identical groups as a result of the reputational penalties for interacting with members of different groups, even when such interactions would be profitable in the absence of the reputation mechanism. Thus the divisions in my theory are “socially constructed” divisions, that do not depend on fundamental differences between the preferences or technologies of different groups. I argue that primordial theories of social division have difficulty explaining the persistence of division within single communities, and thus that most if not all observed persistent social divisions are constructed divisions of the type described in my model.

Certain kinds of socio-economic changes may lead to the demise of socially constructed division. Suppose that an improvement in the rule of law allows people to profit from some kinds of interactions even in the absence of the intertemporal incentives that characterize informal cooperative relationships. This increase may make interactions between members of different groups sufficiently profitable that people are willing to accept the punishment for such interactions, thereby destroying the segregated equilibrium. This may explain why social divisions seem to be more important in developing countries with weak formal institutions. Interestingly, the breakdown of group segregation associated with improvement in the rule of law may reduce welfare by reducing the level of cooperation in society. Thus, improvements to the rule of law may face political opposition in divided societies.

³See Chandra (2012) for discussion.

2 Theory

2.1 Setup

There is a continuum of agents with mass M . This mass is divided into G groups, with group g having mass λ_g , and with $\sum_{g=1}^G \lambda_g = M$. These groups are payoff irrelevant, but group membership is observable. The groups are ranked hierarchically, with group 1 ranked the highest and group G ranked the lowest. I defer the explanation of the significance of the group hierarchy. Time is discrete and all agents have a fixed discount factor δ . In each period the following things happen:

1. Each unmatched agent i chooses a set of groups $\mathcal{G}_i \subseteq \{1, \dots, G\}$ over which to search for new partner.

I refer to this set as the search set. Each unmatched agent pays a search cost $c(\theta_i)$, where $\theta_i = \left(\sum_{g \in \mathcal{G}_i} \lambda_g \right) / M$. I assume that $c(\theta_i) \geq 0$ and $c'(\theta_i) < 0$ for all θ_i .

2. Each agent is provisionally matched with another agent according to a uniform probability function weighted by the measure of groups over which each agent chooses to search. Specifically, suppose that there are a finite number of kinds of agents, where a kind is a set of agents from the same group who all choose the same search set. Let the proportion of kind κ within the pool of unmatched agents be α_κ , and let $\mathcal{G}(\kappa)$ be the search set of agents of kind κ . Then given an unmatched agent of kind κ , the probability that the agent is provisionally matched with an agent of kind κ' is given by:

$$\frac{\mathbb{1}_{\mathcal{G}(\kappa)}(\kappa') \frac{\alpha_{\kappa'}}{\sum_{\kappa'' \in \mathcal{G}(\kappa)} \alpha_{\kappa''}} + \mathbb{1}_{\mathcal{G}(\kappa')}(\kappa) \alpha_{\kappa'} \frac{\alpha_\kappa}{\sum_{\kappa'' \in \mathcal{G}(\kappa')} \alpha_{\kappa''}}}{1 + \sum_{\kappa'' \text{ s.t. } \kappa \subseteq \mathcal{G}(\kappa'')} \alpha_{\kappa''} \frac{\alpha_\kappa}{\sum_{\kappa''' \in \mathcal{G}(\kappa'')} \alpha_{\kappa'''}}}$$

Here $\mathbb{1}_{\mathcal{G}(\kappa)}$ is the indicator function for the set $\mathcal{G}(\kappa)$. This expression represents the following intuition. An agent can be provisionally matched either by finding a partner through her own effort, or by being found through her partner's search effort. Suppose that conditional on an agent finding her partner through her own effort, the agent has a uniform probability of provisionally matching with any agent in her search set. Then the conditional probability that an agent of kind κ finds a partner of kind κ' is given by the first term in the numerator of the expression above. Similarly, conditional on the agent being found through her partner's search effort, the probability that she is found by a partner of kind κ' is given by the second term in the numerator. When normalized by the factor in the denominator, the sum of these terms then represents the total probability that an agent of kind κ is provisionally matched with an agent of kind κ' . Conditional on an agent of kind κ being matched with an agent of kind κ' , I say that the agent of kind κ is the “finder” with probability

$$\frac{\mathbb{1}_{\mathcal{G}(\kappa)}(\kappa') \frac{\alpha_{\kappa'}}{\sum_{\kappa'' \in \mathcal{G}(\kappa)} \alpha_{\kappa''}}}{\mathbb{1}_{\mathcal{G}(\kappa)}(\kappa') \frac{\alpha_{\kappa'}}{\sum_{\kappa'' \in \mathcal{G}(\kappa)} \alpha_{\kappa''}} + \mathbb{1}_{\mathcal{G}(\kappa')}(\kappa) \alpha_{\kappa'} \frac{\alpha_{\kappa}}{\sum_{\kappa'' \in \mathcal{G}(\kappa')} \alpha_{\kappa''}}}$$

and that the agent of kind κ is the “findee” with probability

$$\frac{\mathbb{1}_{\mathcal{G}(\kappa')}(\kappa) \alpha_{\kappa'} \frac{\alpha_{\kappa}}{\sum_{\kappa'' \in \mathcal{G}(\kappa')} \alpha_{\kappa''}}}{\mathbb{1}_{\mathcal{G}(\kappa)}(\kappa') \frac{\alpha_{\kappa'}}{\sum_{\kappa'' \in \mathcal{G}(\kappa)} \alpha_{\kappa''}} + \mathbb{1}_{\mathcal{G}(\kappa')}(\kappa) \alpha_{\kappa'} \frac{\alpha_{\kappa}}{\sum_{\kappa'' \in \mathcal{G}(\kappa')} \alpha_{\kappa''}}}.$$

3. Provisionally matched agents observe all observable information about their partners, where what is observable is discussed below. Each agent may then choose to accept or reject the match. If either agent rejects the match, then both agents return to step 1, and in addition, the finder must pay a large penalty. The point of the penalty is to prevent any agent from reducing her search cost by searching over groups whose members would reject matches with that agent. Otherwise a match forms and both agents continue to step 4.
4. Every agent i receives a payoff $y b(\theta_i)$, where $b(\theta_i) \geq 0$ for all θ_i and $b'(\theta_i) > 0$.
5. All matched agents play a stage game, described below.
6. For each matched pair of agents i and j , let a_i and a_j be the actions chosen in the stage game. If $a_i \neq a_j$, then the match breaks up and both players begin the next period unmatched. Otherwise, continue to the final step.
7. All remaining matches break up exogenously with probability p . Any surviving matched pairs remain matched at the beginning of the next period.

The stage game is as follows.⁴ Both partners in the relationship simultaneously choose a stage game action $a \in [0, \infty)$. An agent’s payoff is $\Pi(a, a')$, where the agent chooses action a and her partner chooses action a' . Define $v(a) = \Pi(a, a)$ and $d(a) = \Pi(0, a)$. I make the following assumptions on Π , v , and d :

Assumption 1. 1. For all $a > 0$ and all a' , $\Pi(0, a') > \Pi(a, a')$.

2. $v(a)$ is bounded.

3. $v(0) = d(0) = 0$

4. $v(a)$ and $d(a)$ are continuous, twice differentiable, and strictly increasing in a .

5. $v'(0) = d'(0)$

⁴This stage game was first described in Ghosh and Ray (1996).

6. $v(a)$ is strictly concave in a and $d(a)$ is strictly convex in a .

Part 1 of the assumption states that 0 is the strictly dominant action in the stage game, which can be interpreted as a generalized prisoner's dilemma with a continuum of actions. If both players play a then both receive a payoff $v(a)$, and I will sometimes refer to this as the value of cooperation at level a . If one player plays a and the other plays 0, then the player who plays 0 gets $d(a)$, and I will sometimes refer to this as the value of cheating at level a . Part 2 is required to rule out Ponzi schemes, in which any level of cooperation can be attained through the promise of ever higher levels of cooperation in the future. Parts 3 through 6 imply that the temptation to cheat is small for a small, and that the temptation to cheat grows large as a gets large. These assumptions ensure that the solution to each agent's maximization problem is interior.

The unique Nash equilibrium of the stage game is for both players to play $a = 0$. However, it may be possible to sustain higher levels of cooperation through intertemporal incentives, as is standard in the literature on repeated games. Thus $v(a)$ represents the value of a long-term relationship at level a . In contrast, players automatically receive the payoff $yb(\theta)$ merely from being matched, and so these parameters represent the value of a short-term relationship, which does not depend on intertemporal incentives. Many kinds of relationships have both a long-term component and a short-term component. For example, the short-term component of a trade relationship might be spot exchange of goods whose quality is observable. A long-term component of this relationship might be exchange of goods whose quality is not observable. In the long-term trade relationship there is greater opportunity to cheat by providing goods of low quality, and so intertemporal incentives are more important for maintaining the relationship. Similarly, the short-term component of a credit relationship might be a collateralized loan. The collateral provides the incentive to repay the loan, and so intertemporal incentives are not necessary. In contrast, the long-term component of a credit relationship might be an uncollateralized loan. In this case the incentive to repay the current loan is the prospect of receiving additional loans in the future. A relationship may also combine short-term and long-term components of different kinds. For example, two agents may exchange goods of observable quality, and they may also provide each other with uncollateralized trade credit. The multiplier y indexes the relative value of short-term and long-term relationships.

The search cost $c(\theta)$ for an agent to find a partner and the short-term benefit from being matched with that partner $yb(\theta)$ depend on the proportion θ that measures the size of the search set relative to the size of the population. The $c(\cdot)$ function is decreasing in θ , which expresses the idea that it is easier to find a match when searching over a larger set. The $b(\cdot)$ function is increasing in θ , which expresses the benefit of diversity. If an agent is choosing a partner from a larger and hence more diverse set, then she has more different kinds

of partners to choose from and so she is more likely to end up with a match who is well-suited to her needs.

One way to justify the forms of the $c(\cdot)$ and $b(\cdot)$ functions is the following. Each agent finds a match by assembling a set of potential matches of fixed size, and then selecting the best match from the set. In order to find this set of potential matches, each agent draws partners randomly from the entire population, paying a cost for each draw, and keeping only those potential matches who are in the agent's search set. The larger the proportion of the search set in the population, the more likely the agent is to draw a member of the search set on each try and the smaller the expected cost of finding any fixed number of potential matches. This explains the form of the $c(\cdot)$ function. After assembling a set of potential matches, each agent observes the short-term relationship value of each potential match, and chooses to form a match with the potential match that has the highest short-term relationship value. Potential matches are distributed in a space with some metric, and the larger the search set, the more likely an agent is to draw potential matches that are located far away from each other. In addition, the short-term relationship values of potential matches are correlated, and the correlation is smaller between the values of potential matches that are located far away from each other. Thus, the expected correlation between the short-term relationship values of the potential matches is smaller if the agent searches over a larger set, and so the expected value of the highest value potential match is higher. This explains the form of the $b(\cdot)$ function. In appendix A I provide a formal microfoundation for the search process summarized by the $c(\cdot)$ and $b(\cdot)$ functions. However, it is also possible to imagine search processes for which $c(\cdot)$ and $b(\cdot)$ depend not only on the fraction of agents in the search set as a proportion of the population, $(\sum_{g \in \mathcal{G}} \lambda_g) / M$, but also or instead on the absolute number of agents in the search set, $\sum_{g \in \mathcal{G}} \lambda_g$. While it does not matter much for my results if $b(\cdot)$ depends on the proportion or the absolute size of the search set, whether $c(\cdot)$ depends on the proportion or the absolute size is crucial. What $c(\cdot)$ depends on is ultimately an empirical question. In the empirical section, I show that the data support the hypothesis that, at least to a first approximation, $c(\cdot)$ depends on the proportion of agents in the search set and not on the absolute size of the search set, and I discuss the theoretical significance of this finding.

Each agent can observe her group and the group of any other agent with whom she is matched. Each agent can also observe the history of play within her current match, but she cannot observe the history of play in any match in which she does not participate. However, each agent can observe something about with whom her partner has matched in the past. Specifically, for each group g , an agent can observe whether her current partner has ever been matched with any agent in group g . Let $\mathcal{H}_i \subseteq \{1, \dots, G\}$ be the set of groups g such that agent i has been matched with a member of group g in the past. I refer to the set \mathcal{H}_i as agent i 's past match set.

Let a_i^τ be the action played in the τ th period of the current match by player i , denote agent i 's partner by $\mu(i)$, and let $a_{\mu(i)}^\tau$ be the action played by agent i 's partner $\mu(i)$. Then the history of a match that has lasted for τ periods is $h^\tau = \{(a_i^1, a_{\mu(i)}^1), \dots, (a_i^\tau, a_{\mu(i)}^\tau)\}$.

A (pure) strategy for agent i is a tuple $s_i(h^\tau, g_i, g_{\mu(i)}, \mathcal{H}_i, \mathcal{H}_{\mu(i)}) = \{\mathcal{G}_i(g_i, \mathcal{H}_i), m_i(g_i, g_{\mu(i)}, \mathcal{H}_i, \mathcal{H}_{\mu(i)}), a_i(h^{\tau-1}, g_i, g_{\mu(i)}, \mathcal{H}_i, \mathcal{H}_{\mu(i)})\}$. Here \mathcal{G}_i is agent i 's search set, $m_i \in \{A, R\}$ is agent i 's decision to accept or reject the match after observing her partner's past match set as well as her group, and a_i is the stage game action. Note that I do not allow agents to condition their strategies on the history of any match prior to the current match.

2.2 The Information Assumption

Before continuing, I briefly discuss the information assumption. Most models in the community enforcement literature follow Kandori (1992) and Okuno-Fujiwara and Postlewaite (1995) in assuming that players can observe something about what their partners have done *within* their past relationships. In contrast, I assume that players have no information about what their partners have done within their past relationships, but that players can observe something about *with whom* their partners have interacted in the past. I claim that this information assumption reflects the true village information environment better than do the assumptions of the previous literature.

Empirical evidence supports my characterization of the village information environment. Udry (1990) notes that villagers have good information about their neighbors' social and economic situations. In particular, he writes that villagers are able to report accurately about the ceremonies that their neighbors have given and attended, which supports my claim that villagers keep track of with whom their neighbors interact. In contrast, it is much harder to observe what happens within many kinds of relationships. For example, consider a credit relationship. In this kind of relationship, the stage game actions correspond to financial transactions between the relationship partners. However, various papers have shown that it is frequently very difficult for people to observe their neighbors' financial transactions. Goldstein (2000) and Anderson and Baland (2002) show that people in village environments are frequently unaware even of their own spouses' financial transactions. So my information assumption seems like a reasonable approximation to the true information environment in villages, at least for many kinds of relationships.⁵

⁵It is also helpful to discuss the difference between my information assumption and the information assumption in Pęski and Szentes (forthcoming). In Pęski and Szentes, agents have information not only about with whom their partners have interacted in the past, but also about with whom their partners' partners have interacted, with whom their partners' partners' partners have interacted, and so on to infinity. Pęski and Szentes use this information structure in a model that is otherwise related to mine to derive quite different results. In contrast, in my model agents have information about with whom their partners have interacted in the past, but that is all. The information requirements of my model are thus much less stringent and as a result possibly more plausible than the information requirements in Pęski and Szentes. The requirement of knowledge of an infinite chain of past interactions is also a feature of the caste equilibrium in Akerlof (1976).

2.3 Equilibrium Concept

An equilibrium of my model is defined by three properties. First, an equilibrium must be a steady state. In a steady state, no agent's past match set can change from one period to the next on the equilibrium path. Second, an equilibrium must satisfy an individual incentive compatibility condition analogous to subgame perfection. This is the familiar requirement that the strategy profile must be robust to the possibility of individual deviations. However, intuitively this requirement is not sufficient to prevent all plausible deviations. In an environment in which matched partners can communicate prior to choosing their actions, it is plausible to suppose that two matched agents could agree to deviate simultaneously from any given strategy profile. This intuition motivates my third requirement, that an equilibrium strategy profile must be robust to the possibility of joint deviations by any pair of matched agents. I adapt this requirement from Ghosh and Ray (1996), and following their lead, I call the requirement *bilateral rationality*. It is closely related to the various renegotiation-proofness concepts discussed in Bernheim and Ray (1989) and Farrell and Maskin (1989).

Formally, an equilibrium consists of an assignment of a past match set to each agent and a strategy profile s . Let γ_g be a vector that denotes, for each possible past match set \mathcal{H} , the proportion of agents in group g who are assigned that past match set, and let γ denote the matrix of all the γ_g 's. I assume that γ is commonly known. I say that a past match set is reachable for an agent from an equilibrium if it would be possible for that agent to have that past match set after some sequence of deviations from the equilibrium.⁶ Define the discounted expected utility for player i by $EU_i[s_i, s_{\mu(i)}, s_{-i}, h^\tau, \gamma]$, which is the expected utility that player i receives while playing strategy s_i when matched with player $\mu(i)$ playing strategy $s_{\mu(i)}$, if all other players are playing strategy profile s_{-i} , given γ , after match history h^τ . Note that, due to the steady state condition, the expected utility does not depend on the period or on any previous history from any past match. In addition, define $\frac{s'_i}{s_i}$ to be the strategy of playing s'_i for the duration of the current match and s_i thereafter. Such a collection is a bilaterally rational equilibrium if:

1. No agent's past match set changes from one period to the next on the equilibrium path.
2. $EU_i[s_i, s_{\mu(i)}, s_{-i}, h^\tau, \gamma] \geq EU_i[s'_i, s_{\mu(i)}, s_{-i}, h^\tau, \gamma]$ for all strategies s'_i and for all $h^\tau, g_i, \mathcal{H}_i, g_{\mu(i)}$, and $\mathcal{H}_{\mu(i)}$ such that \mathcal{H}_i and $\mathcal{H}_{\mu(i)}$ are reachable for agents i and $\mu(i)$.
3. There do not exist any $h^\tau, g_i, \mathcal{H}_i, g_{\mu(i)}, \mathcal{H}_{\mu(i)}, s'_i, s'_{\mu(i)}$ such that \mathcal{H}_i and $\mathcal{H}_{\mu(i)}$ are reachable for i and $\mu(i)$ and such that

⁶In particular, the null past match set is not reachable for any agent from any equilibrium that I consider.

$$\begin{aligned}
EU_i \left[\frac{s'_i}{s_i}, \frac{s'_{\mu(i)}}{s_{\mu(i)}}, s_{-i}, h^\tau, \gamma \right] &\geq EU_i \left[s''_i, \frac{s'_{\mu(i)}}{s_{\mu(i)}}, s_{-i}, h^\tau, \gamma \right] \text{ for all } s''_i \text{ and} \\
EU_{\mu(i)} \left[\frac{s'_{\mu(i)}}{s_{\mu(i)}}, \frac{s'_i}{s_i}, s_{-i}, h^\tau, \gamma \right] &\geq EU_{\mu(i)} \left[s''_{\mu(i)}, \frac{s'_i}{s_i}, s_{-i}, h^\tau, \gamma \right] \text{ for all } s''_{\mu(i)}
\end{aligned}$$

and

$$\begin{aligned}
EU_i \left[\frac{s'_i}{s_i}, \frac{s'_{\mu(i)}}{s_{\mu(i)}}, s_{-i}, h^\tau, \gamma \right] &\geq EU_i [s_i, s_{\mu(i)}, s_{-i}, h^\tau, \gamma] \text{ and} \\
EU_{\mu(i)} \left[\frac{s'_{\mu(i)}}{s_{\mu(i)}}, \frac{s'_i}{s_i}, s_{-i}, h^\tau, \gamma \right] &\geq EU_{\mu(i)} [s_{\mu(i)}, s_i, s_{-i}, h^\tau, \gamma]
\end{aligned}$$

with at least one of the last two inequalities strict.

Condition 1 is the steady state condition. Condition 2 is the individual incentive compatibility condition. Any agent i with any possible combination of group and reachable past match set must prefer to follow the strategy profile instead of playing any alternate strategy when matched with a partner with any other possible combination of group and reachable past match set, after any possible match history. This includes combinations of groups and past match sets and match histories that cannot occur on the equilibrium path. Note that I do not require the individual incentive compatibility condition to hold for any possible γ , only for the actual steady state γ . Implicitly I am requiring the strategy profile to continue to be optimal after any countable number of deviations, which can produce matches between agents with any possible combination of groups and reachable past match sets. However, I do not require the strategy profile to be optimal after a positive mass of agents have deviated, which could change γ , or for agents with past match sets that are not reachable from the equilibrium past match set assignment. Condition 3 is the bilateral rationality condition. It states that it should not be possible for any two agents to renegotiate to a new strategy profile that is individually incentive compatible for both of them and that is weakly Pareto superior. Once again, I require the bilateral rationality condition to hold for any possible combination of groups, reachable past match sets, and match histories, including combinations that only occur off the equilibrium path, but I only require the condition to hold for the actual γ .

2.4 A Benchmark Equilibrium

I will begin my analysis by discussing a benchmark strategy profile in which agents do not condition their actions on their own or their partner's group membership or past match set. If the benchmark strategy profile is part of an equilibrium, I will refer to that equilibrium as a benchmark equilibrium.

In the benchmark strategy profile, every agent chooses to search over every group each period. Thus, $\theta = 1$ for all agents in all periods. Every agent accepts every match. In each period of each match, each agent chooses action \bar{a} . If either player unilaterally deviates, then the match breaks up, so it is not necessary to specify continuation strategies after such deviations. Joint deviations are ignored.

Since groups and past match sets are strategically irrelevant under the benchmark strategy profile, it is unnecessary to specify the probabilities of matches between agents of different groups or past match sets.

Let V^m be the value of being an agent in a relationship that is not in the punishment phase at the beginning of a period. Let V^u be the value of being an agent who is unmatched at the beginning of a period. Finally, let V^f be the value that an agent believes that she will receive from any future match with any other agent. It is helpful to distinguish V^f from V^m since agents may be able to affect V^m through renegotiation, but they cannot affect V^f . Bilateral rationality dictates that each pair of agents chooses the level of cooperation that maximizes their joint utility. That is, V^m must satisfy:

$$V^m = \max_a (1 - \delta)[v(a) + yb(1)] + \delta[pV^u + (1 - p)V^m] \quad (1)$$

subject to the constraint

$$V^m \geq (1 - \delta)[d(a) + yb(1)] + \delta V^u \quad (2)$$

Equation (1) says that a matched agent gets $v(a) + yb(1)$ in the current period. The match then breaks up exogenously with probability p , in which case the agent gets the payoff to being unmatched V^u next period, and otherwise the match continues providing the player with payoff V^m . The constraint (2) is the individual incentive compatibility constraint. It states that the value of cooperating must be greater than the payoff that the agent receives from cheating, in which case the agent receives $d(a) + yb(1)$ this period and then becomes unmatched and receives payoff V^u with probability 1 next period. The payoff to being an unmatched agent V^u is defined by:

$$V^u = -(1 - \delta)c(1) + V^f \quad (3)$$

Equation (3) says that an unmatched agent must pay the search cost in the current period before being

matched with a new partner.

It must also be the case in equilibrium that each agent optimally chooses to search over all groups, and that it is optimal for each agent to accept every match. These conditions are trivial in the benchmark case, since in the absence of concerns about past match sets or asymmetry between the strategies of members of different groups it is always optimal to search over the entire community, as this both minimizes search costs and maximizes short-term relationship quality. All match partners are identical, so it is also optimal to accept all matches.

A benchmark equilibrium is a tuple $\{V^m, V^u, V^f, \bar{a}\}$ such that V^m , V^u , and V^f satisfy equation (1) subject to (2) and equation (3), such that \bar{a} maximizes (1) subject to (2), and such that $V^m = V^f$.

Define $\hat{a}(p)$ to be the value of a that solves

$$\max_a v(a) - (1 - \delta + \delta p)d(a).$$

The following proposition provides conditions under which a benchmark equilibrium exists, and derives the level of cooperation in a benchmark equilibrium:

Proposition 1. *A benchmark equilibrium exists if and only if c satisfies*

$$c(1) \geq \frac{1}{\delta(1-p)}[d(\hat{a}(p)) - v(\hat{a}(p))]. \quad (4)$$

If a benchmark equilibrium exists, then the equilibrium level of cooperation \bar{a} solves

$$d(\bar{a}) - v(\bar{a}) = \delta(1-p)c(1) \quad (5)$$

All proofs are in appendix B.

The interpretation of the expression for the level of cooperation in the benchmark equilibrium is straightforward. If an agent cheats in the current period, her net gain in the period is the difference between the value of cheating $d(\bar{a})$ and the value of cooperating $v(\bar{a})$. The cost of cheating is that in the next period she will have to pay the search cost to find a new partner with certainty, rather than with probability p , so the net cost, discounted for one period, is $\delta(1-p)[c(1)]$. The maximum level of cooperation that can be sustained is the level of cooperation such that the net cost of cheating is equal to the net benefit. The bilateral rationality condition ensures that all agents will renegotiate up to the highest possible level of cooperation, so only the maximum sustainable level of cooperation is consistent with equilibrium.

I briefly discuss the intuition for the fact that no bilaterally rational equilibrium exists unless $c(1)$ is sufficiently large. I consider strategy profiles in which all agents choose the same level of cooperation every

period. Since all agents accept all matches, any agent can cheat in her current relationship, break up the relationship at the end of the period, pay the search cost $c(1)$, and find a new partner in the next period. Since all agents choose the same level of cooperation, the deviating agent will be able to cooperate at the same level in her new relationship as she did in the old relationship. Thus, if $c(1)$ is low, then the penalty for cheating in any given relationship is low, and so the common sustainable level of cooperation is low. However, if all agents are cooperating at some common low level, then any two matched agents can jointly deviate to a higher level of cooperation. This higher level of cooperation does not violate the individual incentive compatibility constraint, so long as only two agents are cooperating at the high level, because the penalty for breaking up this deviant relationship is high: if either agent breaks the relationship, both agents must go back to cooperating at the low common level of cooperation. Thus the individual incentive compatibility requirement rules out all strategy profiles except those strategy profiles with a low common level of cooperation, and the bilateral rationality requirement rules out strategy profiles with a low common level of cooperation, so that there are no remaining equilibrium strategy profiles. As $c(1)$ gets larger, higher levels of cooperation become compatible with the individual incentive compatibility constraint, and for $c(1)$ sufficiently large there exist levels of cooperation that are high enough to satisfy the bilateral rationality requirement while still satisfying the individual incentive compatibility constraint.⁷

2.5 The Segregated Equilibrium

In this subsection I propose what I will call the segregated strategy profile. As before, if the segregated strategy profile is part of an equilibrium, I refer to the equilibrium as a segregated equilibrium. The key difference between the segregated equilibrium and the benchmark equilibrium is that on the segregated equilibrium path agents interact only with members of their own group, while on the benchmark equilibrium path agents interact with members of every group in the community.

The segregated strategy profile is as follows. An agent from group g with past match set \mathcal{H} chooses to search over all groups g' such that $g' \geq g$ and $g' \in \mathcal{H}$. The agent accepts all matches from members of groups g' such that $g' \leq g$ or $g' \in \mathcal{H}$. This asymmetry between higher and lower ranked groups under the segregated equilibrium strategy profile is the only point where the group hierarchy is significant. I discuss the hierarchy in more detail in section 2.7 below. Each combination of group and past match set is associated with a level of cooperation $\bar{a}_{g,\mathcal{H}}$. If an agent with group and past match set (g, \mathcal{H}) is matched with an agent with group and past match set (g', \mathcal{H}') , then in each period of the match both agents choose action $\min\{\bar{a}_{g,\mathcal{H}}, \bar{a}_{g',\mathcal{H}'}\}$. As in the benchmark equilibrium, any unilateral deviation causes the match to break up, and joint deviations are ignored.

⁷A similar issue arises in Ghosh and Ray (1996), and the proof of proposition 1 draws on ideas from the proofs in that paper.

Let $\mathcal{H}(g)$ be the past match set of group g on the equilibrium path. I look for steady states in which $\mathcal{H}(g) = \{g\}$ for all groups g . Each agent's search set is defined by $\mathcal{G}(g, \mathcal{H}) = \{g'' | g'' \geq g \text{ and } g'' \in \mathcal{H}\}$. That is, each agent searches over her own group and any lower ranking group whose members she has interacted with in the past. Then $\theta_{g, \mathcal{H}} = \sum_{g' \in \mathcal{G}(g, \mathcal{H})} \lambda_{g'} / M$ for all agents. The matching probabilities are also simple so long as there have been no more than a countable number of deviations. I define $q_{g, \mathcal{G}}^{g', \mathcal{H}'}$ to be the probability that an unmatched agent from group g who searches over a set \mathcal{G} will match with an agent from group g' with past match set \mathcal{H}' . On the equilibrium path and after any countable number of deviations, we have:

$$\begin{aligned} q_{g, \mathcal{G}}^{g, \{g\}} &= \frac{2\lambda_g}{\lambda_g + \sum_{g'' \in \mathcal{G}} \lambda_{g''}} \text{ if } g \in \mathcal{G} \\ &= \frac{\lambda_g}{\lambda_g + \sum_{g'' \in \mathcal{G}} \lambda_{g''}} \text{ if } g \notin \mathcal{G} \\ q_{g, \mathcal{G}}^{g', \{g'\}} &= \frac{\lambda_{g'}}{\lambda_g + \sum_{g'' \in \mathcal{G}} \lambda_{g''}} \text{ if } g' \neq g \text{ and } g' \in \mathcal{G} \\ &= 0 \text{ if } g' \neq g \text{ and } g' \notin \mathcal{G} \end{aligned}$$

I begin by deriving expressions for the levels of cooperation $\bar{a}_{g, \mathcal{H}}$, where $\bar{a}_{g, \mathcal{H}}$ can be thought of as the maximum symmetric cooperation level that could be sustained by an agent of group g and past match set \mathcal{H} if the agent chose her search set and match acceptance strategy as specified in the strategy profile, but if the agent expected to be matched each period with a partner who had an infinitely negative outside option. Define $V_{g, \mathcal{H}}^m$ to be the value of being an agent of group and past match set (g, \mathcal{H}) who expects to be matched each period with a partner who has an infinitely negative outside option. Define $V_{g, \mathcal{H}}^u$ to be the value of being an unmatched agent with the same expectations. We have:

$$V_{g, \mathcal{H}}^m = \max_a (1 - \delta)[v(a) + yb(\theta_{g, \mathcal{H}})] + \delta[pV_{g, \mathcal{H}}^u + (1 - p)V_{g, \mathcal{H}}^m] \quad (6)$$

subject to the constraint

$$V_{g, \mathcal{H}}^m \geq (1 - \delta)[d(a) + yb(\theta_{g, \mathcal{H}})] + \delta V_{g, \mathcal{H}}^u \quad (7)$$

where we have

$$V_{g, \mathcal{H}}^u = -(1 - \delta)c(\theta_{g, \mathcal{H}}) + V_{g, \mathcal{H}}^m \quad (8)$$

These conditions are analogous to the conditions for the benchmark equilibrium discussed earlier. The

following lemma follows from simple manipulation of (6), (7), and (8), and so I present it without proof:

Lemma 1. *Define $\bar{a}_{g,\mathcal{H}}$ to be the value of a that solves (6) subject to (7) and (8). Then $\bar{a}_{g,\mathcal{H}}$ is the value of a that solves:*

$$d(\bar{a}_{g,\mathcal{H}}) - v(\bar{a}_{g,\mathcal{H}}) = \delta(1-p)c(\theta_{g,\mathcal{H}}). \quad (9)$$

The expression for $\bar{a}_{g,\mathcal{H}}$ can be compared to the very similar expression for the equilibrium level of competition \bar{a} in the benchmark equilibrium. Notice that the maximum level of cooperation that can be sustained by an agent of group and past match set (g, \mathcal{H}) is decreasing in $\theta_{g,\mathcal{H}}$, because agents with higher $\theta_{g,\mathcal{H}}$ face a lower penalty for breaking a relationship.

I can now define a segregated equilibrium. Let $V_{g,\mathcal{H}}^{g',\mathcal{H}'}$ be the value of being an agent of group and past match set (g, \mathcal{H}) matched with an agent of group and past match set (g', \mathcal{H}') , and let $V_{g,\mathcal{H}}$ be the value of being an agent of group and past match set (g, \mathcal{H}) who is unmatched at the beginning of a period. Let $V_{g,\mathcal{H}}^{g',\mathcal{H}',f}$ be the value that an agent of group and past match set (g, \mathcal{H}) expects to receive if matched with an agent of group and past match set (g', \mathcal{H}') in the future. A segregated equilibrium is a tuple $\{V_{g,\mathcal{H}}, V_{g,\mathcal{H}}^{g',\mathcal{H}'}, V_{g,\mathcal{H}}^{g',\mathcal{H}',f}, \bar{a}_{g,\mathcal{H}}\}$ which satisfies the following conditions:

1. The levels of cooperation $\bar{a}_{g,\mathcal{H}}$ are defined by the expression in Lemma 1.
2. For all g and \mathcal{H} , the values $V_{g,\mathcal{H}}, V_{g,\mathcal{H}}^{g',\mathcal{H}'}$, and $V_{g,\mathcal{H}}^{g',\mathcal{H}',f}$ satisfy

$$V_{g,\mathcal{H}}^{g',\mathcal{H}'} = V_{g,\mathcal{H}}^{g',\mathcal{H}',f} = (1-\delta)[v(\min\{\bar{a}_{g,\mathcal{H}}, \bar{a}_{g',\mathcal{H}'}\}) + yb(\theta_{g,\mathcal{H}})] + \delta[pV_{g,\mathcal{H}} + (1-p)V_{g,\mathcal{H}}^{g',\mathcal{H}'}] \quad (10)$$

$$V_{g,\mathcal{H}}^{g',\mathcal{H}'} \geq (1-\delta)[d(\min\{\bar{a}_{g,\mathcal{H}}, \bar{a}_{g',\mathcal{H}'}\}) + yb(\theta_{g,\mathcal{H}})] + \delta V_{g,\mathcal{H}} \quad (11)$$

and

$$V_{g,\mathcal{H}} = -(1-\delta)c(\theta_{g,\mathcal{H}}) + \sum_{g',\mathcal{H}'} q_{g,\mathcal{G}(g,\mathcal{H})}^{g',\mathcal{H}'} V_{g,\mathcal{H}}^{g',\mathcal{H}'}. \quad (12)$$

3. Agents prefer to accept matches with agents from the same or higher ranking groups rather than rejecting those matches and continuing to search. That is, for all g, g', \mathcal{H} , and \mathcal{H}' such that $g' \leq g$, $g \in \mathcal{H}$, and $g, g' \in \mathcal{H}'$,

$$V_{g,\mathcal{H} \cup \{g'\}}^{g',\mathcal{H}'} \geq V_{g,\mathcal{H}} \quad (13)$$

4. Agents who have not yet interacted with members of a lower ranking group prefer to reject matches with members of that group, rather than accepting such matches and adding the group to their past match set. That is, for all $g, g', \mathcal{H}, \mathcal{H}'$ such that $g' > g, g \in \mathcal{H}$, and $g, g' \in \mathcal{H}'$,

$$V_{g, \mathcal{H}} \geq V_{g, \mathcal{H} \cup \{g'\}, \mathcal{G}(g, \mathcal{H})}^{g', \mathcal{H}'} \quad (14)$$

where $V_{g, \mathcal{H}, \mathcal{G}}^{g', \mathcal{H}'}$ is defined as the value of being an agent from group g with past match history \mathcal{H} who has searched over a set \mathcal{G} and subsequently formed a match with an agent from group g' with past match history \mathcal{H}' . That is, if $g'' \geq g$ for all $g'' \in \mathcal{G}$,

$$V_{g, \mathcal{H}, \mathcal{G}(g, \mathcal{H})}^{g', \mathcal{H}'} = (1 - \delta) \left[v(\min\{\bar{a}_{g, \mathcal{H}}, \bar{a}_{g', \mathcal{H}'}\}) + yb \left(\sum_{g \in \mathcal{G}} \lambda_g \right) \right] + \delta \left[pV_{g, \mathcal{H}} + (1 - p)V_{g, \mathcal{H}, \mathcal{G}}^{g', \mathcal{H}'} \right] \quad (15)$$

5. Agents prefer to search only over groups of equal or lower rank that are in their past match set, rather than searching over lower ranked groups not in their past match set and possibly adding those groups to their past match set. That is, for all g, \mathcal{H} , and \mathcal{G} such that $g \in \mathcal{H}$ and $\mathcal{G} \subseteq \{g, g + 1, \dots, G\}$, we have

$$V_{g, \mathcal{H}} \geq -(1 - \delta)c \left(\sum_{g'' \in \mathcal{G}} \lambda_{g''} \right) + \sum_{g' \in \mathcal{G}} \left[q_{g, \mathcal{G}}^{g', \{g'\}} V_{g, \mathcal{H} \cup \{g'\}, \mathcal{G}}^{g', \{g'\}} \right] \quad (16)$$

The following proposition gives necessary and sufficient conditions for the existence of a segregated equilibrium:

Proposition 2. *A segregated equilibrium exists if and only if*

1.

$$c(1) \geq \frac{1}{\delta(1 - p)} [d(\hat{a}(p)) - v(\hat{a}(p))] \quad (17)$$

2. For all g and all \mathcal{H} such that $g \in \mathcal{H}$,

$$(1 - \delta)c(\theta_{g, \mathcal{H}}) \geq \frac{1 - \delta}{1 - \delta + \delta p} [v(\bar{a}_{g, \mathcal{H}}) - v(\bar{a}_{1, \{1, \dots, G\}})] \quad (18)$$

3. For all $g, g', \mathcal{H}, \mathcal{H}'$ such that $g' > g, g \in \mathcal{H}$, and $g, g' \in \mathcal{H}'$,

$$\begin{aligned} & (1 - \delta)c(\theta_{g, \mathcal{H}}) + \delta p [c(\theta_{g, \mathcal{H}}) - c(\theta_{g, \mathcal{H} \cup \{g'\}})] + \frac{\delta p}{1 - \delta + \delta p} y [b(\theta_{g, \mathcal{H} \cup \{g'\}}) - b(\theta_{g, \mathcal{H}})] \\ & \leq [v(\bar{a}_{g, \mathcal{H}}) - v(\bar{a}_{g, \mathcal{H} \cup \{g'\}})] \end{aligned} \quad (19)$$

4. For all g, \mathcal{H} and \mathcal{G} such that $g \in \mathcal{H}$ and $\mathcal{G} \subseteq \{g, \dots, G\}$,

$$\begin{aligned}
& y \left\{ \frac{1-\delta}{1-\delta+\delta p} [b(\theta_{g, \mathcal{H} \cup \mathcal{G}}) - b(\theta_{g, \mathcal{H}})] + \frac{\delta p}{1-\delta+\delta p} \sum_{g' \in \mathcal{G}/\mathcal{H}} q_{g, \mathcal{G}}^{g', \{g'\}} [b(\theta_{g, \mathcal{H} \cup \{g'\}}) - b(\theta_{g, \mathcal{H}})] \right\} \\
& + (1-\delta) [c(\theta_{g, \mathcal{H}}) - c(\theta_{g, \mathcal{H} \cup \mathcal{G}})] + \delta p \sum_{g' \in \mathcal{G}/\mathcal{H}} q_{g, \mathcal{G}}^{g', \{g'\}} [c(\theta_{g, \mathcal{H}}) - c(\theta_{g, \mathcal{H} \cup \{g'\}})] \\
& \leq \sum_{g' \in \mathcal{G}/\mathcal{H}} q_{g, \mathcal{G}}^{g', \{g'\}} [v(\bar{a}_{g, \mathcal{H}}) - v(\bar{a}_{g, \mathcal{H} \cup \{g'\}})]
\end{aligned} \tag{20}$$

If a segregated equilibrium exists, then the level of cooperation $\bar{a}_{g, \{g\}}$ achieved by agents from group g on the equilibrium path is defined by

$$d(\bar{a}_{g, \{g\}}) - v(\bar{a}_{g, \{g\}}) = \delta(1-p)c(\lambda_g). \tag{21}$$

The first necessary condition for the existence of segregated equilibrium is the same as the condition for the existence of a benchmark equilibrium, and it is necessary for the same reason. Under the segregated equilibrium the individual incentive compatibility constraint and the bilateral rationality constraint must be satisfied for each combination of group and past match history, and the constraints cannot be jointly satisfied unless the search cost is sufficiently large. It is most difficult to satisfy both constraints for agents from group 1 whose past match history includes all groups. These agents search over the entire community (i.e. they choose $\theta = 1$) and thus have the lowest search costs. If both constraints can be satisfied for these agents, then they can be satisfied for agents with all other combinations of group and past match history.

The second necessary condition is the condition that agents must prefer to accept matches with agents from the same or higher ranking groups rather than rejecting those matches and continuing to search. The condition is most difficult to satisfy when the agent is matched with a partner from group 1 whose past match history includes all groups, since these partners can sustain the lowest levels of cooperation. The gain from rejecting such a partner is that with probability 1 the agent will then be matched with a partner from her own group whose past match history includes only that group, and will be able to achieve a higher level of cooperation for a period of time that depends on the expected length of the relationship, parameterized by p . The loss is that she will have to pay the search cost in the current period. The condition states that the loss from rejecting the agent with the larger past match history is greater than the gain.

The third necessary condition is the condition that an agent must prefer to reject matches with agents from lower ranking groups with whom she has not previously interacted and to continue to search rather than accepting those matches and increasing the size of her past match set. If an agent accepts a match with a lower ranking agent, she benefits by not having to pay the search cost again, and by searching over a larger

group in the future she will also pay lower search costs and gain higher quality short-term relationships. However, she will be able to support lower levels of cooperation in the future. The condition states that the gain from rejecting a partner from a lower ranking group is greater than the loss. Notice that conditions 2 and 3 can only be satisfied simultaneously if p is sufficiently large. This reflects the fact that in order for conditions 2 and 3 to be satisfied simultaneously, it must be the case that agents are willing to interact with members of higher ranking groups but not with members of lower ranking groups. In both cases, the agent sacrifices the ability to cooperate at a high level for the duration of the current relationship in exchange for the benefit of not having to pay the search cost again. However, by interacting with a member of a lower ranking group an agent also sacrifices the ability to cooperate at a high level in all future relationships. Thus agents are only willing to behave differently towards members of higher ranking and lower ranking groups if the benefit of cooperation in future relationships is sufficiently important, which is true only if the probability p of moving on to a new relationship is sufficiently large.⁸

The final necessary condition is the condition that agents must prefer to search over their own group rather than searching over their own and any set of lower ranking groups and incurring the possibility of expanding their past match sets. The condition states that the multiplier y must be sufficiently small, that is, the value of short-term relationships must not be too large in relation to the value of long-term relationships. Otherwise, agents would be willing to forgo the benefits of higher levels of cooperation with smaller past match sets in order to achieve higher quality short-term relationships by searching over a larger set of potential partners. I defer further discussion of this condition to the final section of the paper.

While the notation involved in defining the segregated equilibrium is somewhat complex, the idea is simple. When an agent is in a relationship, she would like to commit to interacting only with members of her own group in the future, since by doing so she can lower the outside option to the relationship and therefore support a higher level of cooperation within the relationship. However, when an agent is out of a relationship, it is optimal for her to search over the entire community to find a new partner, since by doing so she both minimizes her search cost and maximizes the short-term value of her next relationship. The segregated equilibrium generates a reputation effect that allows agents to solve this commitment problem. If an agent is observed to have interacted only with members of her own group in the past, then it is believed

⁸One way to relax this condition is to allow agents to have more than one relationship at any given time. In this case, by interacting with an agent from a lower ranking group in one relationship, an agent immediately loses the ability to cooperate at a high level with all of her other relationship partners, while by interacting with a member of a higher ranking group the agent only loses the ability to cooperate at a high level in one relationship. Thus the difference between the payoffs for interacting with members of lower ranking groups and with members of higher ranking groups remains large even for arbitrarily small p .

While analytical convenience is the only reason to restrict agents to having only one relationship partner at a time in my model, it is important that agents be restricted to some fixed, finite number of relationships at any given time. Here my model is complementary to the model of Genicot and Ray (2003), which gives a reason why agents may not be able to cooperate with an unlimited number of partners simultaneously. The complementary purpose of my model is to show why each agent's set of potential future relationship partners might also be restricted relative to the universe of all other agents.

that she will continue to interact only with members of her own group in the future, and so she can support a high level of cooperation. However, if an agent is observed to have interacted with a member of a lower ranking group in the past, it is believed that she will continue to interact with members of that group in the future. Thus it is believed that the agent has a higher outside option to any relationship, and so she can only support a lower level of cooperation within any relationship. This punishment is sufficient to deter agents from interacting with members of lower ranking groups if they have not done so in the past, while agents who have interacted with members of a lower ranking group in the past cannot be subject to any further punishment and so cannot be deterred from interacting with members of that group in the future. Thus, all expectations are fulfilled in equilibrium.

2.6 Welfare, Optimal Group Size, and Group Formation

An immediate corollary of proposition 2 is:

Corollary 1. *For any group g with $\lambda_g < 1$,*

1. *The segregated equilibrium level of cooperation $\bar{a}_{g,\{g\}}$ is higher than the benchmark equilibrium level of cooperation \bar{a} .*
2. *The segregated equilibrium search cost $c(\lambda_g)$ is higher than the benchmark equilibrium search cost $c(1)$ and the segregated equilibrium short-term relationship quality $b(\lambda_g)$ is lower than the benchmark equilibrium short-term relationship quality $b(1)$.*

Corollary 1 demonstrates both the advantage and the disadvantages of the group segregation norm. On the one hand, group segregation allows group members to support higher levels of cooperation. On the other hand, group segregation both increases the search costs paid by each agent and reduces the quality of each agent's short-term relationships. Mathematically, define the average welfare under segregation W_g^S in group g by taking the mean of per-period utility in a given period across all members of group g . Similarly, let W^B be the average welfare in the benchmark equilibrium. Then we have:

$$W_g^S = v(\bar{a}_{g,\{g\}}) + yb(\lambda_g) - pc(\lambda_g) \quad (22)$$

and

$$W^B = v(\bar{a}) + yb(1) - pc(1) \quad (23)$$

The net gain from segregation is:

$$W_g^S - W^B = [v(\bar{a}_{g,\{g\}}) - v(\bar{a})] - y[b(1) - b(\lambda_g)] - p[c(\lambda_g) - c(1)] \quad (24)$$

Welfare is higher under group segregation if p and y are sufficiently small.

Notice that the disutility from high search costs under segregation decreases as the exogenous breakup probability p falls, while the disutility from low quality short-term relationships does not depend on p . Thus in situations where relationships are relatively stable, it seems likely that the disadvantage of segregation from low quality short-term relationships is more important than the disadvantage from high search costs.

It is also possible to derive a result on the optimal group size under segregation. In general, W_g^S is not concave in λ_g , and so there is not necessarily a unique maximizer of $W_g^S(\lambda_g)$ on $[0, 1]$. However, W_g^S is supermodular in (λ_g, y) . Therefore, by Topkis' theorem (Topkis 1998, theorem 2.8.3), we have

Proposition 3. *The set $\arg \max_{\lambda_g \in [0, 1]} W_g^S(\lambda_g)$ has a greatest element and a least element, and both the greatest and the least element are increasing in y .*

Proposition 3 gives a more precise statement of the intuition that the optimal group size is increasing as the relative importance of short-term relationships increases.

Proposition 3 need not imply a positive theory explaining observed group sizes. Group membership must be common knowledge in the community, and so agents cannot unilaterally change groups even if doing so would increase welfare. For this reason, shocks that change group sizes or that make salient new group divisions may have persistent consequences. Understanding the determinants of observed group sizes would be an interesting project for future research.

2.7 Equilibrium Selection and the Hierarchy

I have shown that there exists an equilibrium in which agents interact only with members of their own groups, and that this equilibrium allows agents to support higher levels of cooperation than the benchmark equilibrium. Without the bilateral rationality condition, this result would be trivial. Consider the strategy profile in which agents search only over their own groups and accept matches only with members of their own groups, regardless of past match history, and then choose the highest sustainable level of cooperation within each match. This strategy profile satisfies the individual incentive compatibility requirement, because no agent has a profitable individual deviation. However, the strategy profile does not satisfy the bilateral rationality requirement. Given that all other agents are following the strategy profile, any two provisionally matched agents from different groups can profitably renegotiate to accept their match instead of rejecting

it.⁹ Thus, all of the theoretical interest of the model comes from the need to satisfy the bilateral rationality requirement.¹⁰

However, even though the bilateral rationality requirement rules out many possible alternative strategy profiles, bilateral rationality does not uniquely select one strategy profile as the equilibrium. Thus it is necessary to explain why the segregated equilibrium would be selected. In particular, the hierarchical structure of the segregated equilibrium may seem to be an arbitrary addition to the other assumptions of the model. However, in this subsection I will show that given plausible additional assumptions, any equilibrium in which all agents interact only with members of their own groups on the equilibrium path must have the hierarchical structure of the segregated equilibrium. That is, I will show that segregation necessarily implies hierarchy.

My first additional assumption is:

No irrelevant updating (NIU): *An equilibrium satisfies NIU if, under the equilibrium strategy profile, whenever an agent's past match set does not contain group g , the agent's search set does not contain group g . Formally, if $g \notin \mathcal{H}$, then for all g' , $g \notin \mathcal{G}(g', \mathcal{H})$.*

The point of this principle is to rule out strategy profiles such as the strategy profile under which if a member of group 1 is observed to have interacted with a member of group 2, then it is believed that the member of group 1 will search for matches from group 3 in the future. The idea is that observing a change in how the member of group 1 behaves towards members of group 2 should not affect beliefs about whether the member of group 1 will interact with members of group 3. In other words, the principle states that each agents' willingness to interact with members of any group is in some sense independent of her willingness to interact with members of any other group.

NIU implies that the hierarchical structure of the segregated equilibrium must apply between any two groups in any equilibrium, given that agents interact only with members of their own group on the equilibrium path. Roughly speaking, the intuition for this result is that an agent from group g only rejects matches with members of some group $g' \neq g$ if matching with a member of group g' would lower the level of cooperation that the member of group g could expect to achieve in the future. The level of cooperation that the agent can support in the future is inversely related to the expected size of the agent's search set. By NIU, the fact that the agent has interacted with a member of group g' in the past cannot change expectations about

⁹The bilateral rationality condition also rules out "starting small" strategies such as the strategy profile in Ghosh and Ray (1996), in which agents cooperate at a lower level at the beginning of their relationship and then build up to a higher level of cooperation. With symmetric information, the low initial level of cooperation is Pareto sub-optimal and hence not renegotiation proof.

¹⁰One other strategy profile that has been mentioned in the literature as a way to support high levels of cooperation in random matching environments is the "contagious" strategy profile proposed by Kandori (1992) and Ellison (1994). This strategy profile can only be an equilibrium in an environment with a finite number of agents, and so it is ruled out in my model in which the number of agents is infinite.

whether the agent will search over any third group $g'' \neq g'$ in the future. Thus, if interacting with members of group g' lowers the level of cooperation that the agent expects to achieve in the future, it must be because after interacting with a member of group g' , the agent is expected to search over group g' in the future. In order for this expectation to be rational, members of group g' must accept matches with members of group g on the equilibrium path. Thus, on the equilibrium path, if members of group g reject matches with members of group g' , then members of group g' must accept matches with members of group g .

NIU does not rule out the possibility of cycles, in which on the equilibrium path members of group 1 reject matches with members of group 2, members of group 2 reject matches with members of group 3, and members of group 3 reject matches with members of group 1. In order to rule out this possibility I introduce a second assumption:

Transitivity: *For all g, g', g'' , if $m(g, \{g\}, g', \{g'\}) = R$ and $m(g', \{g'\}, g'', \{g''\}) = R$, then $m(g, \{g\}, g'', \{g''\}) = R$.*

Transitivity implies that there is a well defined ordering of groups in which one group is prior to another group in the order if members of the first group reject matches with members of the second group on the equilibrium path. Given the conditions on model parameters from proposition 2, the segregated equilibrium is the unique equilibrium that satisfies NIU and transitivity, and in which all agents interact only with members of their own groups on the equilibrium path. Proposition 4 states this result formally:

Proposition 4. *Consider an equilibrium, and let $\mathcal{H}(g)$ be the past match set of group g on the equilibrium path. Suppose that:*

1. *For all groups g , $\mathcal{G}(g, \mathcal{H}(g)) = \{g\}$.*
2. *The equilibrium satisfies NIU and transitivity.*
3. *The parameters of the model satisfy the conditions in proposition 2.*

Then the equilibrium is the segregated equilibrium.

Note that if $G = 2$, then NIU and transitivity are automatically satisfied in all equilibria, and so the second condition of proposition 4 is redundant for $G = 2$.

The result that segregation implies hierarchy is unexpected. It is not baked into the motivation of the paper. Nevertheless, the result does seem to be relevant to understanding real world segregation. Casual empiricism suggests that segregation is frequently associated with hierarchy and inequality, although interestingly the group that is at the top of the ritual purity hierarchy need not be the group at the top of the socio-economic scale. For example, while in the South Asian caste system the Brahmins are traditionally at

the top of both the ritual purity hierarchy and the socio-economic scale, groups such as the Gypsies and the ultra-orthodox Jews are at the top of the ritual purity hierarchy even though they are near the bottom of the socio-economic scale in their respective societies.

At a more abstract level, the key takeaway from proposition 4 is that segregation is fundamentally an asymmetric relationship, enforced by only one group out of each pair of groups. This suggests, for example, that in order to implement policies to combat segregation between two groups it is first necessary to identify which group is enforcing segregation and which group is merely accepting it.

2.8 From Theory to Empirics

The following is another immediate corollary to proposition 2:

Corollary 2. *The segregated equilibrium level of cooperation $\bar{a}_{g,\{g\}}$ for members of group g is decreasing in λ_g , the percentage of members of group g in the community.*

I test this proposition by studying informal credit relationships between villagers in rural Nepal. Informal credit relationships benefit households by allowing them to borrow in response to unforeseen shocks or opportunities. However, because formal contract enforcement is unavailable in rural Nepal, households have the option of refusing to repay their loans. Households are willing to repay their loans because by doing so they increase the likelihood that they will be able to receive additional loans in the future.¹¹ The informal credit relationship therefore meets my definition of a cooperative relationship. My measure of the level of cooperation engaged in by each household is the logarithm of the sum of the outstanding amount owed by the household to friends and family members and the outstanding amount owed by friends and family members to the household. My definition of a community is an administrative unit called a ward. My main empirical hypothesis is thus the following:

Empirical Hypothesis: *The amount of informal credit that a household engages in is decreasing in the percentage of the household's caste in the ward.*

The next section is devoted to testing this hypothesis.

¹¹See Ligon, Thomas, and Worrall 2002 for a much more in-depth discussion of the theory of informal insurance under limited commitment.

3 Empirics

3.1 Data

My dataset is the first round of the Nepal Living Standards Survey (NLSS), collected in a joint effort by the Nepalese government and the World Bank between 1995 and 1996. The survey consisted of a nationally representative sample of administrative units called wards, stratified by the major regions of the country. Within each sampled ward, 12 households were chosen to be interviewed. A rural ward typically contains between 50 and 500 households, and consists of a single village in those parts of the country in which houses are clustered into villages. In some parts of the country houses are distributed more or less evenly across the countryside rather than being clustered into villages, and in those parts of the country the ward is a somewhat arbitrary administrative division. Table 1 presents summary statistics. Monetary amounts are in Nepalese rupees; there were approximately 50 rupees to the dollar in 1996. Two additional rounds of data were also collected, from 2003-2004 and from 2010-2011. These rounds were affected by the fighting and aftermath of the Nepalese civil war, which began with a few minor incidents in 1996, escalated in 2001, and ended in 2006. The civil war appears to have had a dramatic effect on the structure of the Nepalese caste system, and the results from the second and third rounds of the survey are quite different from the results from the first round. I discuss the civil war and its effects in more detail in Appendix C.

One advantage of the Nepalese context for my study is the Nepalese government has a fairly short, standardized list of castes that it uses for all surveys. This contrasts markedly with the situation in India, where there are tens of thousands of castes, and where no official data on caste membership has been collected since 1931. The NLSS 1 includes 14 named caste categories, which are the 14 largest castes from the 1991 census making up about 80% of the Nepalese population, and a 15th category, “other”. I check that my regressions are robust to dropping the “other” category.

The NLSS 1 did not include a census of the surveyed wards. An official in each ward was asked to provide estimates of the total number of households in the ward and the number of households in each caste in the ward. I use these data to construct the percentage of each caste in each ward. However, I am only able to match the caste of the household to a caste named by the ward official in 67% of cases. As a robustness check, I also estimate the percentage of each caste in each ward by the percentage of each caste within the set of sampled households in each ward. This specification allows me to generate a caste percentage measure for all households. I check that my main result is robust to using this alternative measure of caste percentage.

3.2 Methodology and Results

My basic identification strategy will be to use a difference-in-difference regression across castes and wards to examine the effects of the percentage of a household's caste in the ward on different outcomes for the household. Thus my identification assumption is that there are no unobserved caste-ward specific factors that affect informal borrowing and lending and that are correlated with the percentage of the caste within the ward. I partially test this assumption by including various covariates in some of my regressions.

My main empirical hypothesis is that the amount of borrowing and lending between caste members is decreasing in the percentage of that caste in a ward. To test this hypothesis, I estimate the following regression:

$$\begin{aligned} C_{icw}^* &= \exp(\beta_1 P_{cw} + \beta_2 X_{icw} + \nu_w + \nu_c + \eta_{icw}) \\ C_{icw} &= \ln C_{icw}^* \text{ if } \ln C_{icw}^* > \xi, 0 \text{ otherwise} \end{aligned} \tag{25}$$

Here C_{icw} is the logarithm of the sum of all outstanding loan balances owed by the household to friends and family and outstanding balances owed by friends and family to the household, where i indexes the household, c indexes the caste, and w indexes the ward. I refer to this quantity as $\ln(\text{CasteCredit})$. If a household has an outstanding loan that has been partially repaid, I calculate the outstanding loan balance as the amount of the principal minus the amount of repayment. Unfortunately it is not possible to calculate the amount of interest, if any, still owed on partially repaid loans from my data. I use the logarithm because CasteCredit appears to be approximately log-normally distributed conditional on $\text{CasteCredit} > 0$. Figure 1 shows the distribution of $\ln(\text{CasteCredit})$ conditional on $\text{CasteCredit} > 0$.

On the right hand side of the regression, P_{cw} is the percentage of caste c in ward w , X_{icw} are covariates, ν_w and ν_c are ward and caste fixed effects, and η_{icw} is an error term. The censoring point ξ is unknown, as it depends on factors such as the unobserved fixed cost of borrowing. Following Carson and Sun (2007), I estimate ξ by $\hat{\xi} = \min_{C_{icw} > 0} C_{icw} - 10^{-6}$.

I estimate (25) using the Honoré panel tobit estimator (Honoré 1992), which unlike the regular tobit estimator can consistently estimate censored regression models with fixed effects. As a robustness check, I also estimate a linear probability model, taking my dependent variable to be 1 if a household owes or is owed anything to or by friends or family:

$$Pr(C_{icw} > 0) = \beta_1 P_{cw} + \beta_2 X_{icw} + \nu_w + \nu_c + \eta_{icw} \tag{26}$$

For this regression, I cluster the standard errors by ward in order to account for possible correlation in

the error structure.

Table 2 presents results on the effect of caste percentage on $\ln(CasteCredit)$. Column 1 shows the regression without covariates, using the Honoré estimator. The effect is negative and strongly significant, as predicted by my model.

A potential concern with my identification strategy is that caste percentage and informal credit are both correlated with some third factor that varies by ward-caste. For example, maybe it is the case that large castes are able to exert disproportionate political influence within the ward, so that they are able to acquire more wealth through political connections, and because they have more wealth they need less insurance and so they engage in less informal credit. In this case it is wealth that has the true effect on informal credit, and the correlation between caste percentage and informal credit is spurious. In order to control partially for such omitted third factors, I include a number of additional covariates in my regression in column 2. The value of each household’s land holdings proxies for the household’s wealth, and the variable “migrant” takes the value 1 if the head of the household was born in a district other than the district he or she is currently living in. Occupation fixed effects include fixed effects for whether the household head’s primary activity is farming, agricultural wage labor, any non-agricultural activity, or not working. The coefficient on caste percentage is unchanged, which indicates that there are no observable third factors that are biasing my estimate.

Columns 3 and 4 of table 2 repeat the previous exercise using the linear probability model instead of the Honoré estimator. The estimate of the effect of caste percentage on caste credit is negative in all specifications.

Table 3 presents some robustness checks. Approximately 80% of the population of Nepal is Hindu and most of the remainder are Buddhist. While the Buddhist part of the population is also divided into groups that are similar in some ways to the Hindu castes, it is possible that the Buddhist groups are not sufficiently similar to the Hindu castes to make the comparison meaningful. Thus in column 1 of Table 3 I restrict my sample to the Hindu population. In column 2 I drop the “Other” castes and restrict my sample only to the 14 named castes that also make up about 80% of the population. In column 3 I drop the top 5% of households by *CasteCredit*, to check that outliers are not driving my results. Some of the ward population data appear to be affected by missing digits and other data entry errors. For example, in one ward the sum of the number of households in each caste is 84, but the total number of households in the ward is 4. In the main dataset, I do my best to correct these errors by hand. I correct the population data for 6 wards. In column 4 check that these corrections do not unduly affect my results by dropping the corrected wards. In column 5 I use the alternative measure of caste percentage estimated as the percentage of households of a given caste among the set of sampled households in each ward. In all of these regressions I use the Honoré

estimator. In each case the coefficient on caste percentage remains negative and strongly significant.

As a placebo regression, table 4 presents the same regressions as in table 2, but instead of using $\ln(CasteCredit)$, I use the log of lending and borrowing to and from all other sources, namely banks, NGOs, shopkeepers, employers/employees, and landlords/tenants, as my dependent variable. I refer to this variable as $\ln(OtherCredit)$. If there were some bias that was affecting my regressions in table 2, this bias would likely affect the regressions in table 4 as well, leading to a negative coefficient on caste percentage in these regressions. In fact, if anything caste percentage has a positive effect on credit to and from other sources, although this effect is not statistically significant. This suggests that members of castes that make up a larger percentage of the community substitute away from caste credit and towards other sources of credit, which is consistent with my theory.

In Table 5, I run regressions to test whether the true effect on cooperation is from caste percentage, as predicted by my model, or from absolute caste size. Column 1 repeats my basic regression estimating the effect of caste percentage on $\ln(CasteCredit)$, again using the Honoré estimator. Column 2 runs the same regression using the absolute number of households in the caste-ward instead of the caste percentage as the independent variable. Absolute caste size also has a negative effect on caste credit, which is not surprising since absolute caste size and caste percentage are correlated. Column 3 includes both absolute caste size and caste percentage in a single regression. The standard errors increase, which again is not surprising due to the correlation between caste size and caste percentage. It is more interesting to look at the change in the size of the estimated coefficients. The coefficient on caste percentage falls by a little under 12%, while the coefficient on absolute caste size falls by a factor of 5. This indicates that it is caste percentage that has the true effect on caste credit, and the positive effect of caste size in column 2 is spurious and due to the correlation between caste size and caste percentage.

4 Alternative Theories

My results can distinguish my model from a number of other models. My first result is that the level of cooperation is decreasing in caste percentage. This result distinguishes my model from models with perfect contracting, either with perfect contracting between any two members of the community or with perfect contracting only between members of the same caste. If there were perfect contracting between any two members of the community, then we would expect that the level of cooperation would not be affected by caste percentage. If there were perfect contracting only between caste members, then we would expect that the level of cooperation would be increasing in caste percentage, as it would be easier for members of larger castes to find relationship partners and the level of cooperation within any given relationship would

be unaffected by caste size. So my results reject models with perfect contracting.

A different class of models that predict that the level of cooperation should be increasing in caste percentage are models in which the social network functions as collateral, as in Greif (1993). In these models, the penalty for cheating in the current relationship is that the cheater will lose her ability to cooperate not just with her current partner but also with all of her other social connections. Thus an agent's social connections act as collateral to guarantee her honesty in her current relationship. In these models agents with large social networks have more to lose from cheating and so they can support higher levels of cooperation. If we assume that each household's social network is identical to its caste, then these models predict again that the level of cooperation should be increasing in the size of the caste. Even if the household's social network is not identical to its caste, we might expect that the size of a household's social network would be increasing in the size of its caste, as in Currarini, Jackson, and Pin (2009). In this case too, the theory of social networks as collateral predicts that the level of cooperation should be increasing in the size of the caste. My results are thus evidence against the social networks as collateral theory.

My second result is that the true effect on the level of cooperation is from caste percentage and not caste size. As discussed in the theory section and in Appendix A, it is possible to construct models in which the cost of searching over a group and hence the level of cooperation within the group depends on the absolute size of the group in addition to or instead of the percentage of the group in the community. The regressions in table 5 suggest that it is in fact the percentage of the group in the community and not the absolute size of the group that matters. This fact is important for the interpretation of the model. When the search cost depends on the percentage of the group in the community, it is advantageous for each group to share a single community with other groups, even though the groups do not interact. On the other hand, when the search cost depends only on the absolute size of the group, there is no advantage to any group to sharing the community to other groups, and so we might expect that groups would split off to form separate, homogenous communities. Thus, only when the search cost depends on the percentage of the group in the community does my theory explain the existence of social division within a single community. I expand on this point in section 5 below.

The result that caste percentage and not caste size affects cooperation also distinguishes my model from the model of Dixit (2003). In Dixit's model, agents can sometimes observe whether their partners have cheated in past relationships, but the flow of information is imperfect and is noisier in larger groups. Thus in large groups an agent's next partner is less likely to find out if she has cheated in her current relationship, and so agents in large groups are more likely to cheat. In this model, the level of cooperation is decreasing in caste size, not caste percentage, since the noisiness of information flow within one group is not affected by the size of any other group. Thus my finding that the true relationship is between caste percentage and

cooperation supports my model against Dixit’s model.

5 Persistence and Change

In this paper I have developed a model of social divisions in which a reputation effect prevents interactions between members of different groups in equilibrium, even though those interactions would be profitable in the absence of the reputation effect. A competing theory is there are fundamental differences between the preferences of members of different groups or the technologies available to members of different groups, and that these differences make interactions between members of different groups unprofitable. In the political science literature, divisions of this type are referred to as “primordial” divisions. The most obvious reason why interactions between members of different groups might be unprofitable is that the members of different groups may speak different languages. Even between members of groups that speak the same language, there might be varying preferences or technologies that make interactions unprofitable. Such differences may include taste-based discrimination (Becker 1957), preferences for expressing one’s identity by engaging in group specific behaviors (Akerlof and Kranton 2000), preferences for passing on group specific traits to children (Bisin and Verdier 2000), differing preferences across groups over public goods (Alesina, Baqir and Easterly 1999, Alesina and La Ferrara 2000), difficulty in judging the abilities of members of different groups (Cornell and Welch 1996), or information frictions between groups that make it difficult to write and enforce contracts (Fearon and Laitin 1996, Miguel and Gugerty 2005, Habyarimana et. al. 2007). All of these barriers to interaction are primordial in that they are generated by differences between groups in preferences or in the technologies for communication and information diffusion. In contrast, divisions that are not primordial are said to be “socially constructed”, and the kind of division described in the model in this paper is an example of a constructed division.

Lazear (1999) develops a simple model to study the persistence of primordial divisions over time. In his model, agents search randomly over the community to find trade partners. Upon finding a trade partner, an agent receives a fixed benefit analogous to the $b(\cdot)$ function in my model, and there is no additional benefit to cooperation. However, a natural barrier such as a language barrier prevents agents from trading with members of different groups. Each agent can pay some cost to overcome the barrier, for example by learning the language of the other group. Agents from minority groups have more of an incentive to learn the majority language, because agents from minority groups are more likely to encounter potential trade partners from the majority group than vice-versa. Alternatively, agents from the minority group can separate themselves completely to form a new community in which they are guaranteed to encounter only other members of their own group. In the model, social division is unstable for most parameter values: either members of

the minority group will assimilate to the majority culture, or members of different groups will separate completely to form distinct communities. Thus while primordial barriers may explain the persistence of differences between different communities, primordial barriers do not easily explain the persistence of group divisions within a single community.

In contrast, in my model it is optimal for members of different groups to share the same community, since the division of the community is what allows each group member to support a high level of cooperation. Paradoxically, agents want to share their communities with members of other groups precisely so that they can not interact with the members of those groups. Rabbi Joel Teitelbaum, one of the founders of the ultra-orthodox Jewish community in the United States, had a saying that perfectly illustrates this dynamic. He said, “Separatism [is] of such importance that even if a city had no wicked Jews, it would be worthwhile to pay some wicked Jews to come and live there so that the good Jews would have something to separate themselves from.”¹² Because it is optimal for members of different groups to share the same community when divisions are socially constructed, socially constructed divisions can persist over time.

Some groups may be separated from the surrounding population by both primordial and constructed barriers, but I argue that in these cases it is the constructed barriers that allow the division to persist. For example, consider the German-speaking immigrants who arrived in Pennsylvania and surrounding states in the 17th and 18th centuries. All of these immigrants faced a primordial language barrier to interacting with the English-speaking majority. One subset of these immigrants, the Amish, also developed a very strong version of the group segregation norm described in this paper.¹³ Non-Amish German speakers assimilated with the majority over time, and now Pennsylvania German has nearly vanished outside of the Amish population. The Amish, however, have maintained their distinctive language and culture despite their minority status. Thus it seems that the group segregation norm and not the language barrier was the key to maintaining the distinctiveness of Amish culture. I conjecture that most if not all persistent primordial divisions are also supported by socially constructed barriers to interaction between groups in this way, and that without the constructed barriers, the primordial divisions would soon collapse.

Despite the persistence of socially constructed divisions, my model also suggests socio-economic changes that could lead to the demise of these divisions. In particular, consider the multiplier y . One factor that may increase y is improvement in formal contract enforcement and the rule of law. As the rule of law becomes more entrenched, it becomes possible to support more and deeper relationships without relying on intertemporal incentives, and so the value of short-term relationships increases relative to the value of long-term relationships. For y sufficiently large, condition 4 for the existence of a segregated equilibrium in

¹²Quoted in Wallace-Wells (2013).

¹³Kraybill (2001) discusses the history of the Amish church.

proposition 2 fails to hold. Thus, as the role of formal contract enforcement expands, the group segregation norm may break down. Moreover, my model suggests that the breakdown of segregation may happen suddenly, as y crosses the critical threshold, even in societies that have been segregated for thousands of years in the past.

There is some empirical evidence that this kind of change does in fact reduce the salience of group membership. Munshi and Rosenzweig (2006) study the effects of increasing economic integration with the outside world on castes in Mumbai. Over the period they study, increasing trade opportunities increased the relative value of formal sector employment as compared to informal sector employment, which in the context of my model could be thought of as an increase in y . Munshi and Rosenzweig show that the percentage of people marrying outside of their castes increased dramatically as formal sector employment opportunities improved. In the context of my model, this could be interpreted as a breakdown of the segregated equilibrium.

The breakdown of socially constructed divisions allows people to take advantage of the full range of possible relationships available to them, and it may also help to ameliorate other problems caused by division such as political conflict and violence. At the same time, the breakdown of these divisions is likely to lead to the loss of traditional community values and the high levels of cooperation that they entail. Finding a balance between these conflicting sets of values is likely to be one of the key issues for many countries as they move forward in the process of development.

A A Microfounded Search Process

In this section I provide a possible microfoundation for the reduced form of the search process described in section 2. I make two simplifications. First, I describe a search process for which $c(\cdot)$ and $b(\cdot)$ are the *expected* cost of search and the *expected* short-term relationship value in each period. Thus the reduced form described in the text abstracts from the uncertainty due to the randomness of the search process. Second, I assume that the characteristics of each agent's match are determined only by that agent's own choices, and not by the choices of the other agents in the population. That is, I assume that the search process is one-sided, as opposed to the two-sided search process described in the text. Avoiding either of these simplifications would substantially complicate the exposition without adding any additional insight.

Suppose that the groups are located along a line of length 1, with the points occupied by the members of each group forming a connected set. Suppose also that each agent's search set must itself be a connected set. In each period, each unmatched agent chooses a search set \mathcal{G} subject to this restriction, as in the procedure described in the text. The agent can then evaluate exactly two potential matches from within her search set. The agent finds potential matches by drawing at random from the set of all unmatched agents. At each draw the probability of drawing a potential match who is in the search set is $\theta = \left(\sum_{g \in \mathcal{G}} \lambda_g\right) / M$. The agent continues drawing until she has found two potential matches who are members of her search set, discarding any potential matches who are not in her search set. In order to make each draw the agent must pay a cost c . The number of draws that the agent must make has a negative binomial distribution, and the expected cost to the agent is $c(\theta) = c \frac{2(1-\theta)}{\theta}$. As in the text, we have $c(\theta) \geq 0$ and $c'(\theta) < 0$.

Having found two potential matches, each agent then observes the short-term relationship value of each potential match. Short-term relationship values are distributed $N(0, \sigma^2)$, and short-term relationship values are correlated across potential matches. The coefficient of correlation ρ between the short-term relationship values of two potential matches depends on the distance w between those matches, with $\rho = 1 - w$. Assuming that potential matches are uniformly distributed over the line, as is the case under all the strategy profiles that I consider, the pdf of w is

$$f(w) = \frac{2(\theta - w)}{\theta^2}, \quad 0 \leq w \leq \theta. \quad (27)$$

The agent chooses to form a match with the potential match with the highest short-term relationship value. Thus, conditional on the distance w between her two potential matches, the expected short-term relationship value of her match is the expected value of the first-order statistic of a sample of two normally distributed random variables with coefficient of correlation $1 - w$, where I denote the realized value of the first-order statistic by b . The expected value of b given w is (see Afonja 1972 for a derivation)

$$E[b|w] = \sigma \sqrt{\frac{w}{\pi}}. \quad (28)$$

By the law of iterated expectations, we have

$$E[b] = E[E[b|w]] \quad (29)$$

$$= \int_0^\theta \sigma \sqrt{\frac{w}{\pi}} \frac{2(\theta - w)}{\theta^2} dw \quad (30)$$

Define $b(\theta) = E[b]$. The equation above implies that $b(\theta) \geq 0$. Taking the derivative with respect to θ yields

$$b'(\theta) = \int_0^\theta \sigma \sqrt{\frac{w}{\pi}} \left(\frac{4w}{\theta^3} - \frac{2}{\theta^2} \right) dw \quad (31)$$

$$= \frac{4}{15} \sigma \sqrt{\frac{1}{\pi \theta}} \quad (32)$$

which implies that $b'(\theta) > 0$. Thus the microfoundation above yields the reduced forms described in the text.

As discussed in the text, it is also possible to find a microfoundation for a search process in which the $c(\cdot)$ and $b(\cdot)$ functions depends on the absolute size of the search set, $\sum_{g \in \mathcal{G}} \lambda_g$, as well as or instead of on the proportion of agents in the search set, $(\sum_{g \in \mathcal{G}} \lambda_g) / M$. One way to do this is to suppose that the cost c of drawing each potential match is decreasing in the absolute number of agents in the search set, and to suppose that the agents are distributed over a line with length M instead of length 1. As discussed in the text, whether the $c(\cdot)$ and $b(\cdot)$ functions depend on the absolute size of the search set or the proportion of agents in the search set is ultimately an empirical question, and I explore what the data have to say about this question in the empirical section.

B Proofs

B.1 Proof of Proposition 1

Plugging equation (3) into the constraint (2) and equation (1) and rearranging yields

$$V^m = \max_a \frac{1-\delta}{1-\delta+\delta p} [v(a) + yb(1)] - \frac{\delta p(1-\delta)}{1-\delta+\delta p} c(1) + \frac{\delta p}{1-\delta+\delta p} V^f \quad (33)$$

subject to

$$V^f \leq \frac{1}{\delta(1-p)} [v(a) - (1-\delta+\delta p)d(a)] + yb(1) + (1-\delta)c(1) \quad (34)$$

Since v is strictly concave and d is strictly convex, there exists a finite value of a that maximizes $v(a) - (1-\delta+\delta p)d(a)$. Recall that $\hat{a}(p)$ was defined as the value of a that solves

$$\max_a v(a) - (1-\delta+\delta p)d(a).$$

Since $v(a) - (1-\delta+\delta p)d(a)$ has a maximum value, there exists \hat{V}^f such that the constraint (34) can be satisfied for $a \geq 0$ if and only if $V^f \leq \hat{V}^f$, with \hat{V}^f defined by

$$\hat{V}^f = \frac{1}{\delta(1-p)} [v(\hat{a}(p)) - (1-\delta+\delta p)d(\hat{a}(p))] + yb(1) + (1-\delta)c(1). \quad (35)$$

Now, define a function $\phi(x)$ by

$$\phi(x) = \max_a \frac{1-\delta}{1-\delta+\delta p} [v(a) + yb(1)] - \frac{\delta p(1-\delta)}{1-\delta+\delta p} c(1) + \frac{\delta p}{1-\delta+\delta p} x \quad (36)$$

subject to

$$x \leq \frac{1}{\delta(1-p)} [v(a) - (1-\delta+\delta p)d(a)] + yb(1) + (1-\delta)c(1) \quad (37)$$

Any fixed point of ϕ is a benchmark equilibrium. However, notice that ϕ is not well-defined for all x , since for $x > \hat{V}^f$ there is no $a \geq 0$ that satisfies (37). I will show that ϕ has a fixed point if and only if ϕ is defined on a sufficiently large domain. I prove the following lemma:

Lemma 2. 1. ϕ is continuous and differentiable.

2. Let $x = -(1-\delta)c(1)$. Then $\phi(x)$ is well defined and $\phi(x) > x$.

3. $\frac{\partial \phi}{\partial x} < 1$ for all x such that $\phi(x)$ is well-defined.

Proof. Part 1: This follows from the fact that $v(a)$ and $d(a)$ are continuous and differentiable.

Part 2: Let $x = -(1-\delta)c(1)$. Assumption 1 ensures that the constraint (37) can be satisfied with $a > 0$. Examination of (36) shows that $\phi(x) > 0$, so $\phi(x) > x$.

Part 3: By the envelope theorem, we have

$$\frac{\partial \phi}{\partial x} = \frac{\delta p}{1 - \delta + \delta p} - \psi < 1 \quad (38)$$

where $\psi > 0$ is the Lagrange multiplier on the constraint (37). \square

Lemma 2 implies that $\phi(V^f)$ has exactly one fixed point with $V^f \geq -(1-\delta)c(1)$ if and only if $\phi(\hat{V}^f) \leq \hat{V}^f$. Since $V^f = -(1-\delta)c(1)$ is the minimum individually rational value to being in any relationship, this implies that there exists a unique benchmark equilibrium if and only if $\phi(\hat{V}^f) \leq \hat{V}^f$. Plugging in the expression for \hat{V}^f from (35) into this condition and rearranging yields the condition that a benchmark equilibrium exists if and only if

$$c \geq \frac{1}{\delta(1-p)}[d(\hat{a}(p) - v(\hat{a}(p)))]. \quad (39)$$

This completes the proof.

B.2 Proof of Proposition 2

The necessity of the first condition is proved by following exactly the same steps as in the proof of proposition 1, using the individual incentive compatibility and bilateral rationality constraints for agents from group 1 with past match set $\{1, \dots, G\}$. These agents face the lowest search costs of any agent with any combination of group and past match history, so if the search cost is sufficiently high that both the individual incentive compatibility and bilateral rationality constraints can be satisfied for these agents, then both constraints can be satisfied for all agents. The remaining necessary conditions for equilibrium are derived through algebraic manipulation of inequalities (13), (14), and (16).

It is necessary to justify the fact that the segregated equilibrium level of cooperation between a pair agents with group and past match set (g, \mathcal{H}) , (g', \mathcal{H}') is $\min\{\bar{a}_{g, \mathcal{H}}, \bar{a}_{g', \mathcal{H}'}\}$. It is clear that this is a lower bound for the level of cooperation that any two agents could support. The level of cooperation that an agent can support in a relationship is inversely related to the value that the agent expects to receive from future relationships, and the construction of $\bar{a}_{g, \mathcal{H}}$ assumes that agents expect to have future partners who can support maximal levels of cooperation and hence who will yield maximum future relationship value. However, if I were to allow the possibility of a positive measure of deviations from the equilibrium, it would be possible to construct situations in which matched pairs of agents would have low expectations for future relationships and hence could support higher levels of cooperation within certain current relationships. For example, consider first the segregated equilibrium without deviation and consider two matched agents from

group 1. Each agent expects to be able to cooperate at level $\bar{a}_{1,\{1\}}$ in each future relationship, and this expectation allows the pair to cooperate at level $\bar{a}_{1,\{1\}}$ in the present relationship as well. Now suppose that a small measure ϵ of group 1 agents deviate and begin to search over the entire community, and consider again two matched non-deviant members of group 1. In expectation, these agents will be able to cooperate at a level below $\bar{a}_{1,\{1\}}$ in their future relationships, because they will sometimes encounter deviant members of group 1 who can only support a low level of cooperation. This implies that the matched non-deviant agents will be able to cooperate at a level above $\bar{a}_{1,\{1\}}$ in the present relationship. More generally, a problem arises if any agent expects to receive a lower value from her future relationships than from her present relationship. This problem does not affect the segregated equilibrium given that I only allow for at most a countable number of deviations from equilibrium. After any countable number of deviations from equilibrium, all agents including deviant agents will be able to cooperate at the same level or higher in all of their future relationships as in their present relationship with probability 1. Thus $\min\{\bar{a}_{g,\mathcal{H}}, \bar{a}_{g',\mathcal{H}'}\}$ is the highest level of cooperation that two matched agents with group and past match set (g, \mathcal{H}) , (g', \mathcal{H}') can support on the equilibrium path or after any countable number of deviations, and so it is the equilibrium level of cooperation.

B.3 Proof of Proposition 4

Consider an equilibrium, and let $\mathcal{H}(g)$ be the past match set of group g on the equilibrium path. Given the assumptions stated in the proposition, I prove that the equilibrium must be the segregated equilibrium. The proof proceeds in three steps. First I show that it is not possible for the members of any two groups both to reject matches with the other group on the equilibrium path. Second I show that it is not possible for the members of any two groups both to accept matches with the other group on the equilibrium path. This implies that on the equilibrium path, of any two groups, the members of one group must accept matches with members of the other group while members of the second group must reject matches with members of the first group. Finally, I show that if the members of one group reject matches with members of another group on the equilibrium path, then off the equilibrium path a member of the first group searches over the second group and accepts matches with members of the second group if and only if the member of the first group has interacted with a member of the second group in the past. Together these three steps imply that the hierarchical structure of the segregated equilibrium applies between any two groups. Transitivity then implies that there are no cycles and so the equilibrium is the segregated equilibrium. The steps are:

Step 1: For any two groups g and g' , it is not possible that $m(g, \mathcal{H}(g), g', \mathcal{H}'(g')) = R$ and $m(g', \mathcal{H}'(g'), g, \mathcal{H}(g)) = R$.

Proof. Suppose not. By NIU, if a member of group g accepts a match with a member of group g' , no group

other than group g' can thereby be added to the member of group g 's search set. If members of group g reject matches with members of group g' and vice versa on the equilibrium path, then members of group g prefer not to search over group g' on the equilibrium path since the match will be rejected with probability 1. Thus if a member of group g accepts a match with a member of group g' , the member of group g 's search set cannot thereby increase in size. Since the level of cooperation that any agent can expect to achieve in the future can decrease only if she adds new groups to her search set, the level of cooperation that an agent from group g can expect to achieve in the future cannot decrease as a result of accepting a match with a member of group g' . But, by the second necessary condition of proposition 2, each agent prefers to accept matches with all other agents so long as doing so does not reduce the levels of cooperation that the agent can attain in the future. Thus agents from group g prefer to accept matches with members of group g' and vice versa, contradicting the assumption. \square

Step 2: For any two groups g and g' , it is not possible that $m(g, \mathcal{H}(g), g', \mathcal{H}'(g')) = A$ and $m(g', \mathcal{H}'(g'), g, \mathcal{H}(g)) = A$.

Proof. Suppose not. Without loss of generality, suppose that $\bar{a}_{g, \mathcal{H}(g) \cup \{g'\}} \geq \bar{a}_{g', \mathcal{H}'(g') \cup \{g\}}$. From the proof of proposition 2, on the equilibrium path and after any countable number of deviations, the level of cooperation that can be supported in any match is the minimum of the maximum levels of cooperation that can be supported by two agents in the match. Since $\bar{a}_{g, \mathcal{H}}$ is the maximum level of cooperation that an agent from group g with past match set \mathcal{H} can support, and since $\bar{a}_{g', \mathcal{H}'(g') \cup \{g\}} \leq \bar{a}_{g', \mathcal{H}'}$, agents from group g with past match set $\mathcal{H} \cup \{g'\}$ can support the same level of cooperation when matched with either members of group g with past match set $\mathcal{H}(g)$ or members of group g' with past match set $\mathcal{H}'(g')$. Since matching with either a member of group g or group g' cannot affect the level of cooperation that an agent from group g with past match set $\mathcal{H}(g) \cup \{g'\}$ expects to receive in the future, and since searching over a larger search set yields a match of greater present value, agents from group g with past match set $\mathcal{H}(g) \cup \{g'\}$ strictly prefer searching over both group g and g' to searching over only one of these groups. Thus an agent from group g with past match set $\mathcal{H}(g) \cup \{g'\}$ will search over both groups g and g' . By assumption, on the equilibrium path agents do not search over any group but their own. But this means that if an agent from group g with past match set $\mathcal{H}(g)$ accepts a match with a member of group g' , then it will be believed that the agent's search set will change to include group g' in the future. The third necessary condition of proposition 2 states that each agent rejects matches with members of any other group if it is believed that the agent's search set would change to include the other group in the future after accepting the match. Thus members of group g with past match set \mathcal{H} must reject matches with members of group g' , contrary to the assumption. \square

Step 3: Suppose that $m(g, \mathcal{H}(g), \mathcal{H}(g), g', \mathcal{H}(g')) = R$ and $m(g', \mathcal{H}(g'), g, \mathcal{H}(g)) = A$. Then $m(g', \mathcal{H}', g, \mathcal{H}) = A$ for all reachable \mathcal{H} and \mathcal{H}' , while $m(g, \mathcal{H}, g', \mathcal{H}') = A$ if and only if $g' \in \mathcal{H}$. Moreover, $g \notin \mathcal{G}(g', \mathcal{H}')$ for all reachable \mathcal{H}' , while $g' \in \mathcal{G}(g, \mathcal{H})$ if and only if $g' \in \mathcal{H}$.

Proof. By NIU, if an agent from group g' accepts a match with a member of group g , then no group other than group g can thereby be added to the agent's search set. But agents from group g' do not search over group g on the equilibrium path or after any countable number of deviations because if they do so they will be rejected with probability 1. Thus accepting a match with a member of group g cannot affect the level of cooperation that an agent from group g' expects to achieve in future matches. The second necessary condition of proposition 2 states that agents accept matches with all other agents if doing so does not affect the level of cooperation that they expect to achieve in future matches. Thus, agents from group g' always accept matches with members of group g and never search over group g , regardless of past match sets.

By the second necessary condition of proposition 2, each accepts matches with members of any other group if doing so would not lower the level of cooperation that the agent expects to achieve in the future. In order for members of group g to reject matches with members of group g' on the equilibrium path, it must be the case that if accepting a match with a member of group g' would lower the level of cooperation that a member of group g could achieve in the future by increasing the size of the set that the member of group g 's search set. By NIU, after accepting a match with a member of group g' , a member of group g 's search set can only change by adding g' as a new element. Thus, it must be the case that members of group g with past match set \mathcal{H} search over group g' and accept matches with members of group g' if $g' \in \mathcal{H}$. By the third necessary condition of proposition 2 each agent rejects matches with members of any other group if doing so would lower the level of cooperation that the agent could achieve in the future, so members of group g with past match set \mathcal{H} reject matches with members of group g' if $g' \notin \mathcal{H}$. \square

Steps 1, 2, and 3 together imply that the hierarchical structure of the segregated equilibrium applies between any two groups. Transitivity then implies that there is a well-defined ordering of the groups, which implies that the equilibrium is the segregated equilibrium.

C The Nepalese Civil War

In this section I discuss the Nepalese civil war and its impact on my results. I draw my summary of the history of the war from Do and Iyer (2010). The Nepalese civil war was a conflict between Maoist insurgents based in rural areas and the government, headed by the Nepalese monarchy, based in the capital city of Kathmandu. There is no clear consensus on the causes of the conflict, and in particular there is disagreement about

whether caste grievances were a major factor leading to the conflict. The conflict began in February 1996 with an attack on a police post in the remote district of Rolpa, but the conflict remained at a fairly low level until 2001, when most of the Nepalese royal family were murdered under mysterious circumstances and a new king, Gyanendra, took the throne. King Gyanendra mobilized the Royal Nepal Army to combat the insurgents in November 2001, greatly increasing the intensity of the conflict. The conflict continued until a cease-fire in 2005, at which point the insurgents were in control of most rural areas of the country. In April 2006, Gyanendra was forced to give up power, and Nepal officially became a republic, with the Maoists winning the largest number of seats in elections for the legislature in 2008. Clearly, there was significant institutional change during this period.

I have data from three rounds of the Nepal Living Standards Survey, the first which was conducted from 1995-96 before the civil war began and which is the dataset used in the body of the paper, the second which was conducted from 2003-04, and the third which was conducted from 2010-11. The second round of the survey was affected by the civil war, and it was not possible to survey some planned sample areas. I show the results of my main regression using each of the three rounds of the survey. Column 1 of appendix table 1 repeats the results from column 2 of table 2, while columns 3 and 5 repeat the same regression using the second and third rounds of the survey. While caste percentage has a strongly negative effect on caste borrowing and lending in the first round of the survey, as predicted by my model, there are no significant effects in the second or third round of the survey. In columns 2, 4, and 6, I include as a regressor the interaction between the percentage of a household's caste in the ward and the number of battle deaths per 1000 population in the household's district during the civil war, where the number of battle deaths proxies for the intensity of violence in the district. The district is a larger administrative division than the ward; Nepal is divided into 75 districts. Data on the number of battle deaths in each district are drawn from Do and Iyer (2010), who collect the data from the Informal Sector Service Centre, a Nepalese non-governmental organization. In the first round of the survey, before the conflict began, the interaction term is insignificant, but in the second round of the survey the inclusion of the interaction term shows that caste percentage continued to have a negative effect on caste borrowing and lending, but only in those districts that were less affected by the civil war. Finally, in the third round of the survey caste percentage does not appear to have an effect on caste borrowing and lending in any part of the country. These results seem to show that the civil war somehow changed the relationship between caste and household borrowing and lending, starting in the districts most affected by the war in 2003-2004 and spreading to the entire country by the time of the post-war survey in 2010-2011.

I am not sure what is the best interpretation of these results. The civil war created substantial refugee flows, and one possibility is just that this migration scrambled traditional communities, causing me to

mismeasure the sizes of the relevant communities. Another more interesting possibility is that the civil war somehow broke down the traditional barriers between castes, perhaps by increasing the value of cross-caste interaction within communities for mutual defense, as in the mechanism discussed in section 5. My data do not allow me to distinguish between these possibilities. In any case, the data from the first round of the survey are likely the most reliable for my purposes as they were not affected by the civil war, and the results from these data strongly support my theory.

References

- [1] Afonja, Biyi (1972), "The Moments of the Maximum of Correlated Normal and t-Variates," *Journal of the Royal Statistical Society* 34:2, 251-262.
- [2] Akerlof, George (1976), "The Economics of Caste and of the Rat Race and Other Woeful Tales," *Quarterly Journal of Economics* 90:4, 599-617.
- [3] Akerlof, George and Rachel E. Kranton (2000), "Economics and Identity," *Quarterly Journal of Economics* 115:3, 715-753.
- [4] Alesina, Alberto, Reza Baqir, and William Easterly (1999), "Public Goods and Ethnic Divisions," *Quarterly Journal of Economics* 114:4, 1243-1284.
- [5] Alesina, Alberto and Eliana La Ferrara (2000), "Participation in Heterogeneous Communities," *Quarterly Journal of Economics* 115:3, 847-904.
- [6] Anderson, Siwan and Jean-Marie Baland (2002), "The Economics of ROSCAs and Intrahousehold Resource Allocation," *Quarterly Journal of Economics* 117:3, 963-995.
- [7] Austen-Smith, David and Roland G. Fryer, Jr. (2005), "An Economic Analysis of 'Acting White,'" *Quarterly Journal of Economics* 120:2, 551-583.
- [8] Becker, Gary S. (1957), *The Economics of Discrimination*, Chicago: University of Chicago Press.
- [9] Berman, Eli (2000), "Sect, Subsidy, and Sacrifice: An Economist's View of Ultra-Orthodox Jews," *Quarterly Journal of Economics* 115:3, 905-953.
- [10] Bernheim, B. Douglas, and Debraj Ray (1989), "Collective dynamic consistency in repeated games," *Games and Economic Behavior* 1:4, 295-326.
- [11] Bisin, Alberto and Thierry Verdier (2000), "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits," *Quarterly Journal of Economics* 115:3, 955-988.
- [12] Central Bureau of Statistics, Nepal (1996), *Nepal Living Standards Survey 1995-1996*, provided by the National Data Archive, www.cbs.gov.np.
- [13] Central Bureau of Statistics, Nepal (2004), *Nepal Living Standards Survey 2003-2004*, provided by the National Data Archive, www.cbs.gov.np.
- [14] Central Bureau of Statistics, Nepal (2011), *Nepal Living Standards Survey 2010-2011*, provided by the National Data Archive, www.cbs.gov.np.
- [15] Chandra, Kanchan (ed.) (2012), *Constructivist Theories of Ethnic Politics*, Oxford: Oxford University Press.
- [16] Cornell, Brad and Ivo Welch (1996), "Culture, Information and Screening Discrimination," *Journal of Political Economy* 104:3, 542-571.
- [17] Currarini, Sergio, Matthew O. Jackson, and Paolo Pin (2009), "An Economic Model of Friendship: Homophily, Minorities, and Segregation," *Econometrica* 77:4, 1003-1045.
- [18] Dixit, Avinash K. (2003), "Trade Expansion and Contract Enforcement," *Journal of Political Economy* 111:6, 1293-1317.
- [19] Do, Quy-Toan and Lakshmi Iyer (2010), "Geography, Poverty and Conflict in Nepal," *Journal of Peace Research* 47:6, 735-748.
- [20] Dumont, Louis (1970), *Homo Hierarchicus: The Caste System and Its Implications*, Chicago: University of Chicago Press.

- [21] Eeckhout, Jan (2006), "Minorities and Endogenous Segregation," *Review of Economic Studies* 73:1, 31-53.
- [22] Ellison, Glenn (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies* 61:3, 567-588.
- [23] Farrell, Joseph and Eric Maskin (1989), "Renegotiation in repeated games," *Games and Economic Behavior* 1:4, 327-360.
- [24] Fearon, James D. and David D. Laitin (1996), "Explaining Interethnic Cooperation," *American Political Science Review* 90:4, 715-735.
- [25] Genicot, Garance and Debraj Ray (2003), "Endogenous Group Formation in Risk-Sharing Arrangements," *Review of Economic Studies* 70:1, 87-113.
- [26] Ghosh, Parikshit and Debraj Ray (1996), "Cooperation in Community Interaction Without Information Flows," *Review of Economic Studies* 63:3, 491-519.
- [27] Goldstein, Markus Paul (2000), "Intrahousehold allocation and farming in southern Ghana," Unpublished Ph.D. Dissertation, University of California, Berkeley.
- [28] Greif, Avner (1993), "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," *American Economic Review* 83:3, 525-548.
- [29] Grimard, Franque (1997), "Household consumption smoothing through ethnic ties: evidence from Cote d'Ivoire," *Journal of Development Economics* 53:2, 391-422
- [30] Habyarimana, James, Macartan Humphreys, Dan Posner, and Jeremy Weinstein, "Why Does Ethnic Diversity Undermine Public Goods Provision? An Experimental Approach," *American Political Science Review* 101:4, 709-725.
- [31] Höfer, András (2004), *The caste hierarchy and the state in Nepal: a study of the Muluki Ain of 1854*, Lalitpur: Himal Books.
- [32] Honoré, Bo E. (1992), "Trimmed LAD and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects," *Econometrica* 60:3, 533-565.
- [33] Kandori, Michihiro (1992), "Social Norms and Community Enforcement," *Review of Economic Studies* 59:1, 63-80.
- [34] Kraybill, Donald B. (2001), *The Riddle of Amish Culture*, Baltimore: The Johns Hopkins University Press.
- [35] Lazear, Edward P. (1999), "Culture and Language," *Journal of Political Economy* 107:S6, S95-S126.
- [36] Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall (2002), "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies* 69:1, 209-244.
- [37] Mazzocco, Maurizio and Shiv Saini (2012), "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review* 102:1, 428-468.
- [38] Miguel, Edward and Mary Kay Gugerty (2005), "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya," *Journal of Public Economics* 89:11-12, 2325-2368.
- [39] Munshi, Kaivan and Mark Rosenzweig (2006), "Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy," *American Economic Review* 96:4, 1225-1252.
- [40] Munshi, Kaivan and Mark Rosenzweig (2009), "Why is Mobility in India so Low? Social Insurance, Inequality, and Growth," *mimeo*, Brown University.

- [41] Okuno-Fujiwara, Masahiro and Andrew Postlewaite (1995), "Social Norms and Random Matching Games," *Games and Economic Behavior* 9:1, 79-109.
- [42] Pęski, Marcin and Balazs Szentes (forthcoming), "Spontaneous Discrimination," *American Economic Review*.
- [43] Topkis, Donald M. (1998), *Supermodularity and Complementarity*, Princeton: Princeton University Press.
- [44] Udry, Christopher (1990), "Credit Markets in Northern Nigeria: Credit as Insurance in a Rural Economy," *World Bank Economic Review* 4:3, 251-269.
- [45] Wallace-Wells, Benjamin (2013), "Them and Them," *New York Magazine*, April 29.
- [46] Weyrauch, Walter Otto and Maureen Anne Bell (1993), "Autonomous Lawmaking: The Case of the 'Gypsies,'" *Yale Law Journal* 103:2, 323-399.

Figure 1: Distribution of $\ln(\text{CasteCredit})$ conditional on $\text{CasteCredit} > 0$




Table 1: Summary Statistics
NLSS 1995-1996 round

Caste credit	3,880 (11,270)
Other credit	4,985 (22,785)
Caste percentage	0.511 (0.315)
Land value	185,995 (638,633)
HH head age	44.75 (14.46)
HH head education	2.054 (3.669)
Female HH head	0.134 (0.341)
HH size	6.114 (2.987)
Migrant HH head	0.142 (0.349)
Number of observations with caste credit > 0	804
Caste credit conditional on caste credit > 0	11,038
Number of observations with other credit > 0	897
Other credit conditional on other credit > 0	12,694

Monetary amounts are in Nepalese rupees. Standard errors in parentheses.

Table 2: Effect of Caste percentage on credit to and from friends and family
NLSS 1995-1996 round

VARIABLES	(1) Honoré panel tobit ln(CasteCredit)	(2) Honoré panel tobit ln(CasteCredit)	(3) OLS =1 if CasteCredit > 0	(4) OLS =1 if CasteCredit > 0
Caste percentage	-1.953*** (0.618)	-1.949*** (0.586)	-0.144** (0.0643)	-0.143** (0.0648)
Land value		6.13e-08 (2.63e-07)		-7.89e-09 (1.58e-08)
HH head age		-0.00997 (0.00787)		-0.00169** (0.000786)
HH head education		0.0444 (0.0339)		0.000836 (0.00300)
Female HH head		-0.856*** (0.312)		-0.0753*** (0.0288)
HH size		0.0972** (0.0396)		0.00303 (0.00352)
Migrant HH head		0.762* (0.402)		0.0762** (0.0363)
Occupation FE?	NO	YES	NO	YES
Caste and Ward FE?	YES	YES	YES	YES
Number of Wards	205	205	205	205
Observations	2279	2239	2279	2239

*** p<0.01, ** p<0.05, * p<0.1. Dependent variable in columns 1 and 2 is $\ln(CasteCredit)$, where $CasteCredit$ is the sum of borrowing and lending to and from friends and family. Dependent variable in columns 3 and 4 is equal to 1 if $CasteCredit > 0$, and 0 otherwise. Columns 1 and 2 estimated using the Honoré panel tobit estimator. Columns 3 and 4 estimated using OLS. Standard errors in parentheses in columns 1 and 2. Robust standard errors clustered by ward in parentheses in columns 3 and 4.

Table 3: Robustness checks
NLSS 1995-1996 round

VARIABLES	(1) Drop non-hindus	(2) Drop other castes	(3) Drop top 5 percent	(4) Drop corrected wards	(5) Alternative caste percentage measure
Caste percentage	-2.243*** (0.588)	-2.048*** (0.651)	-1.656*** (0.628)	-2.147*** (0.571)	-1.543*** (0.556)
Land value	1.38e-07 (2.03e-07)	7.56e-09 (3.19e-07)	-3.23e-07* (1.90e-07)	6.50e-08 (2.70e-07)	-2.39e-08 (2.24e-07)
HH head age	-0.0120 (0.00862)	-0.0111 (0.00790)	-0.0187** (0.00825)	-0.00938 (0.00795)	-0.00567 (0.00796)
HH head education	0.0254 (0.0356)	0.0426 (0.0373)	0.00700 (0.0351)	0.0454 (0.0344)	0.0225 (0.0331)
Female HH head	-0.884*** (0.340)	-0.819** (0.332)	-0.944*** (0.330)	-0.854*** (0.313)	-0.408 (0.320)
HH size	0.0964** (0.0412)	0.0548 (0.0446)	0.0690* (0.0384)	0.0993** (0.0400)	0.113*** (0.0394)
Migrant HH head	0.685 (0.459)	0.482 (0.432)	0.771* (0.403)	0.847** (0.415)	0.111 (0.397)
Occupation FE?	YES	YES	YES	YES	YES
Caste and Ward FE?	YES	YES	YES	YES	YES
Number of Wards	200	195	205	199	274
Observations	1977	1852	2143	2178	3152

*** p<0.01, ** p<0.05, * p<0.1. All regressions using Honoré panel tobit estimator. Standard errors in parentheses.

Table 4: Effect of caste percentage on credit to and from all other sources
NLSS 1995-1996 round

VARIABLES	(1) Honoré panel tobit ln(OtherCredit)	(2) Honoré panel tobit ln(OtherCredit)	(3) OLS =1 if OtherCredit > 0	(4) OLS =1 if OtherCredit > 0
Caste percentage	0.822 (0.597)	0.767 (0.544)	0.0750 (0.0548)	0.0658 (0.0546)
Land value		-2.41e-07 (2.91e-07)		-1.87e-08* (1.04e-08)
HH head age		-0.00115 (0.00844)		-0.000790 (0.000731)
HH head education		0.0700* (0.0373)		0.000860 (0.00345)
Female HH head		-0.919*** (0.345)		-0.0569* (0.0331)
HH size		0.220*** (0.0345)		0.0140*** (0.00377)
Migrant HH head		0.143 (0.406)		0.00694 (0.0438)
Occupation FE?	NO	YES	NO	YES
Caste and Ward FE?	YES	YES	YES	YES
Number of Wards	205	205	205	205
Observations	2279	2239	2279	2239

*** p<0.01, ** p<0.05, * p<0.1. Dependent variable in columns 1 and 2 is ln(*OtherCredit*), where *OtherCredit* is the sum of borrowing and lending to and from all sources other than friends and family. Dependent variable in columns 3 and 4 is equal to 1 if *OtherCredit* > 0, and 0 otherwise. Columns 1 and 2 estimated using the Honoré panel tobit estimator. Columns 3 and 4 estimated using OLS. Standard errors in parentheses in columns 1 and 2. Robust standard errors clustered by ward in parentheses in columns 3 and 4.

Table 5: Does caste percentage or caste size matter for caste credit?
NLSS 1995-1996 round

VARIABLES	(1) logcastecredit	(2) logcastecredit	(3) logcastecredit
Caste percentage	-1.949*** (0.586)		-1.718* (1.008)
Caste size		-0.0106* (0.00572)	-0.00192 (0.00689)
Land value	6.13e-08 (2.63e-07)	6.61e-08 (3.19e-07)	7.15e-08 (3.17e-07)
HH head age	-0.00997 (0.00787)	-0.0101 (0.00795)	-0.0100 (0.00789)
HH head education	0.0444 (0.0339)	0.0456 (0.0343)	0.0445 (0.0340)
Female HH head	-0.856*** (0.312)	-0.854*** (0.315)	-0.859*** (0.312)
HH size	0.0972** (0.0396)	0.0951** (0.0398)	0.0965** (0.0395)
Migrant HH head	0.762* (0.402)	0.769* (0.409)	0.759* (0.399)
Occupation FE?	YES	YES	YES
Caste and Ward FE?	YES	YES	YES
Number of Wards	205	205	205
Observations	2239	2239	2239

*** p<0.01, ** p<0.05, * p<0.1. All estimations using the Honoré panel tobit estimator. Dependent variable is $\ln CasteCredit$. Standard errors in parentheses.

Appendix Table 1: Did the civil war change the caste system?

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	1995-1996 NLSS	1995-1996 NLSS	2003-2004 NLSS	2003-2004 NLSS	2010-2011 NLSS	2010-2011 NLSS
Caste percentage	-1.949*** (0.586)	-1.534* (0.876)	-0.536 (0.528)	-1.293** (0.634)	0.405 (0.510)	0.663 (0.695)
Caste percentage \times combat deaths				0.807** (0.406)		-0.247 (0.327)
Land value	6.13e-08 (2.63e-07)	5.71e-08 (2.58e-07)	2.32e-07 (1.65e-07)	2.35e-07 (1.65e-07)	2.73e-08*** (3.72e-09)	2.73e-08*** (3.71e-09)
HH head age	-0.00997 (0.00787)	-0.0100 (0.00793)	-0.0302*** (0.00687)	-0.0296*** (0.00686)	-0.0216*** (0.00729)	-0.0218*** (0.00730)
HH head education	0.0444 (0.0339)	0.0435 (0.0335)	0.0255 (0.0273)	0.0245 (0.0273)	0.0101 (0.0239)	0.0104 (0.0239)
Female HH head	-0.856*** (0.312)	-0.860*** (0.314)	-0.855*** (0.305)	-0.845*** (0.305)	-0.0412 (0.208)	-0.0422 (0.209)
HH size	0.0972** (0.0396)	0.0978** (0.0398)	0.133*** (0.0389)	0.133*** (0.0389)	0.154*** (0.0357)	0.153*** (0.0356)
Migrant HH head	0.762* (0.402)	0.767* (0.401)	-0.375 (0.335)	-0.384 (0.335)	0.173 (0.280)	0.179 (0.278)
Occupation FE?	YES	YES	YES	YES	YES	YES
Caste and Ward FE?	YES	YES	YES	YES	YES	YES
Number of Wards	205	205	229	229	325	325
Observations	2239	2239	2526	2526	3682	3682

*** p<0.01, ** p<0.05, * p<0.1. All estimations using the Honoré panel tobit estimator. Dependent variable is $\ln CasteCredit$. Combat deaths are the number of combat deaths per 1000 population in the district. Standard errors in parentheses.