Learning to Trust: Experimental Evidence of Social Learning in a Real-world Social Network of Player A's in a Trust Game

Lindsey Roots (RTSLIN001)
Supervisor: Malcolm Keswell

Half Dissertation
Master's in Applied Economics
University of Cape Town

May 23, 2013

Abstract

Using experimental data on trust from individuals within a real-world social network this dissertation considers whether perceptions of trustworthiness are malleable by the social context in which an individual operates. Contributing to the literature on social effects, I give evidence for a social learning effect in trust games, whereby individuals update their perceptions of the trustworthiness of others in accordance with the outcomes of peers in their network. The dissertation is unique in its consideration of the direct impact of social influence on the beliefs that determine economic outcomes, rather than on economic outcomes themselves. This is made possible through the use of longitudinal experimental data capturing the evolution of trust beliefs through trust games over two rounds of play.

1 Introduction

A vast literature now exists, both in sociology and in economics, dealing with social influences on individual outcomes such as drug-use, teenage pregnancy and welfare recipiency. However, there is little evidence that empirically ties these social influence outcomes back to processes of social learning in beliefs or preferences. This dissertation is a first step towards making a contribution to this research agenda. Using data from trust games prior to and after the formation of a real world social network in South Africa, I examine whether player A's in the game change their offers to player B's, by updating their priors concerning the trustworthiness of a player B in response to new information that becomes available through A's social network.

Trust games were conducted at the outset of the formation of a real world social network, within a youth leadership programme in South Africa, and were again conducted almost a year after the start of network formation. Social learning within the network is measured by considering the extent to which perceptions of trustworthiness are determined by the apparent trustworthiness of the trust game partners encountered by peers in the individual's social network.

The literature on social effects, epitomised by the work of Steven Durlauf (see, for example, Durlauf (2002)), questions whether social and economic outcomes are determined by more than just individual characteristics. The social effects literature rejects the standard approach in economics, which is to focus primarily on individual and household level characteristics, and instead considers the role of social interactions. The evidential outworking of social effects is apparent in the persistence of economic outcomes across generations, persistence of economic outcomes within geographic neighbourhoods, learning processes in the adoption of innovative technologies and so forth. The critical feature of the social effects literature, which seeks to untangle such phenomena, is the consideration of economic processes and outcomes as determined in part by social interaction and the influence of peers.

Social effects are highly relevant for the process of economic development. Group-level influences go some way to explaining the perpetuation of poverty and the apparent existence of poverty traps, inexplicable by individual characteristics alone (Bowles et al., 2011). A thorough understanding of social influences and group effects is a necessary policy consideration in many contexts. Durlauf (2002) supports this idea with an example of a scholarship programme. A scholarship programme with the objective of increasing educational achievement and considering whether to offer scholarships to individuals dispersed across schools or to concentrate scholarships within a small number of schools should consider the impact of concentration of scholarship students on non-scholarship students. Such 'social multipliers' are highly relevant for the potential overall impact of the policy.

The neighbourhood and peer effects literature has not directly considered social effects in terms of their impact on preferences or beliefs directly, only through revealed preference, for example, teenage pregnancy and drug use (Gaviria and Raphael, 2001). The use of economic experiments in this dissertation makes the observation of beliefs possible. The consideration of social effects in the context of experimental preference data allows for determination of social effects on one aspect of beliefs relevant for economic outcomes.

Another example in economics of agents operating from a set of priors is models of statistical discrimination in the labour market. In this context prior beliefs capture how someone perceives the unobservable traits of another, based on their observable characteristics, such as race. Inferences about unobservable characteristics, such as productivity, are made based on priors linked to observed characteristics. These models are used to explain discrimination in employment without relying on a taste-based discrimination argument. However, the priors from which agents operate in statistical discrimination models cannot reasonably be based on all information available in the economy and must instead be based on the very limited experiences. Therefore, prior beliefs are likely to remain relatively static, with little experimentation to determine the true relationships between the variables (Arrow, 1998). The extent to which a social learning process could cause priors to be updated is an interesting empirical question.

Though this dissertation does not consider priors of employers in job markets directly, the results speak more generally to the way in which individuals update beliefs and preferences, particularly relating to their perception of others in a context of incomplete information. A greater depth of understanding as to the malleability and evolution of preferences and priors, such as the perception of trustworthiness of others, could add greater depth and applicability to the arena of statistical discrimination models, as well as other contexts in which preferences and priors have previously often been assumed to be fixed.

Inferences about preferences and preference change based on changes in choice behaviour can be justified on the grounds of revealed preference. Observing what people do, you can infer their preferences. It is a long established result in standard consumer theory that one can infer things about utility functions (and therefore underlying preference rankings) from the observed choices that people actually make. Samuelson (1948) was the first to demonstrate that the weak axiom of revealed preferences (WARP) coupled with budget balancedness (which imply homogeneity of degree zero) generates negative definiteness and symmetry of the partial derivatives and cross-partial derivatives of given choice functions, which are exactly the same set of restrictions generated from the standard axioms on preferences directly.

2 Social Learning

Social learning, whereby individuals learn and update their preferences according to the behaviour and experiences of others, represents a specific type of social effect. Social learning is conceptually important particularly with regards to technology adoption and growth, the context in which it is usually considered (Bandiera and Rasul, 2006) (Conley and Udry, 2010). Although there is not a substantial literature considering social learning of preferences there is an evident place for such literature, particularly where policy design intends to impact economic outcomes by affecting social norms and values. This dissertation considers social learning of preferences within the realm of pro-social behaviour, trust and perceived trustworthiness.

I seek to quantify the impact of learning from others using individual level data, an objective pursued in the literature by Bandiera and Rasul (2006), Conley and Udry (2010), Foster and Rosenzweig (1995), Munshi (2004), Duflo et al. (2009) and Kremer and Miguel (2007), amongst others. The Conley and Udry (2010) approach, on which the methodology used in this dissertation is broadly based, draws on a classic approach by Coleman et al. (1957), where dissemination of a new drug amongst physicians was studied specifically in relation to the spread of the innovation through the social network. This methodology circumvents the identification challenges of measuring social effects in accordance with one solution discussed by Moffitt (2001), which requires a policy intervention that changes fundamental variables for a subset of the population, in an attempt to influence the outcomes of others in the group. The 'intervention' here, analogously to the methodology of Conley and Udry (2010), is the unexpectedly high or low actualised trust payoffs of other individuals within each individual's communication network. The unanticipated high or low outcomes of particular individuals impact only a subset of the group. Social learning is then identified by observing the impact of these outcomes more widely within the network. The model used in this dissertation captures this intervention in a social learning variable that captures each individual's exposure to surprisingly good outcomes amongst individuals in their social network.

There is limited sound empirical work within the field of social learning, due in part to a paucity of suitable data, but primarily due to complex challenges in identification (Moffitt, 2001). The following section reviews the social learning literature, and outlines some approaches taken to confronting the inherent empirical challenges.

The majority of the empirical literature on social learning centres on technology use, either adoption decisions in a new technology, in most cases, or input decisions such as acreage allocation or fertiliser quantities. Some empirical papers have attempted to draw out specific social learning effects whilst others identify a broader social effect that could encapsulate other forms of social influence, such as mimicry. Most of the empirical papers reviewed here use non-experimental data, with the exception of Kremer and Miguel (2007), who exploit an instance of the randomised allocation of a particular technology. With non-experimental data it is necessary to deal carefully with the implications of Manski's (1993) reflection problem by, for example controlling for omitted variables. The literature reviewed varies in the extent to which it succeeds in confronting these problems.

In the early literature in this field Bikhchandani et al. (1992) present a theoretical paper demonstrating how informational cascades can explain conformity in social behaviour such as national identities, as well as the rapid spread of new behaviours. In the model it is optimal for agents to follow the behaviour of others rather than taking advantage of the full set of first hand information available to them. This work, together with Banerjee's (1992) work on herd behaviour, set the stage for the empirical expansion of the social learning field from the early 1990s. Besley and Case (1994) is the first example in the economics literature of testing for the presence of information externalities. Foster and Rosenzweig (1995), another key early work, investigates adoption of high yielding variety seeds in India. Foster and Rosenzweig's initial work is limited, like many empirical papers in this literature, by its use of network data based on group characteristics rather than network data defined in an individual specific manner.

Munshi's (2004) prominent work identifies social learning processes in rice and wheat production in India. The two crops differ in the heterogeneity of the impact of unobserved individual characteristics on output. By regressing farmers allocations of land to high-yielding variety crops on lagged own acreage and village acreage allocations, Munshi demonstrates that wheat growers respond strongly to neighbours lagged acreage decisions and yield realisations, whilst rice growers place relatively more weight on their own lagged decisions. He attributes this to the greater heterogeneity of rice productivity, given individual unobserved characteristics, relative to wheat. Growers of the less heterogeneous crop respond more strongly to their neighbours' experiences because the homogeneity of experience makes social learning from others relatively more useful. Munshi posits this as an explanation for individuals' apparent disregard for their neighbours' positive experiences of new technologies in certain contexts. In line with the challenge of Manski (1993), Munshi's (2004) results could reflect something other than social learning, for example if acreage decisions were a response to an unobserved signal then neighbours' behaviour could proxy for this signal and give a false impression of social learning. As with much of the literature. Rather than defining the actual network, Munshi uses village membership to proxy for the learning network, an issue that Moffitt (2001) suggests is a weakness in the literature.

Of the papers reviewed here, Kremer and Miguel (2007) are able to deal in the most convincing way with the econometric challenges of identifying social effects. They consider peer effects in the adoption of worm treatment in Kenya with the objective of determining whether peer effects could be large enough to justify a large temporary investment in worm treatment to move society from low adoption to high adoption equilibrium. By exploiting experimental exposure to a new technology through a randomised evaluation they are able to overcome the problem of endogeneity of adoption with the network structure, since exposure to adoption is randomly determined. In their identification strategy they test whether households with more social links to schools randomly chosen for early treatment were more likely to take deworming drugs, conditional on their total number of links to all project schools. The theoretical model accommodates peer effects through peer imitation, social learning about the benefits of new technologies and epidemiological externalities. They find that additional social links to early treatment schools reduce the probability that children take deworming drugs and increase the probability that parents do not think the drugs are effective. The perceived negative social effect is possibly due to lack of recognition that the benefits of the drugs are also experienced through externalities, which reduces the apparent benefit to the individual and might only be learnt with experience. Individuals are found to learn from both close contacts and distant contacts (Granovetter's (1973) 'weak ties'). Using non-experimental analysis Kremer and Miguel (2007) find that individuals are more likely to take the drug if they have greater social contact with others who have recently been exposed to deworming. This result is contrary to the findings of the experimental analysis and is highly suggestive of substantial omitted variable bias in the non-experimental analysis. This raises great concern about the results in the literature more generally and the ability of standard methodologies to appropriately deal with the reflection problem presented by Manski (1993).

Van den Broeck and Dercon (2011) use census data, together with survey data of the complete network of a Tanzanian village, to identify various levels of individual specific reference groups, challenging the widespread use of broad level reference groups for networks in the literature. Geographic proximity, racial group, religious group and similar identifiers are frequently used to define interaction groups, whilst the relevant interaction group may actually be much more limited. They find that kinship networks, but no other type of network, are associated with social externalities in banana output in the context of a Tanzanian village. This raises concern with the tendency in the literature to proxy for networks with shared characteristics in the absence of network data. Additionally the study uses exogenous group controls to tackle the standard omitted variable bias problem in analysing social effects.

As in Munshi (2004), the literature largely relies on the correlation of behaviour within social networks, usually by considering number of adopters in a network, to identify 'learning', which leaves other confounding social effects as a possible or partial explanation. The literature has variously tried to confront the challenge raised by Hogset and Barrett (2010) regarding the identification of distinct, rather than general, social effects.

Moser and Barrett (2003) attempt to separately identify social learning and social pressures in the adoption of a new rice growing technique in Madagascar. They achieve this by assuming initial adoption to be a concave function of cumulative village adoption, implying that the effect of time t experience of village farmers should be exceeded by the cumulative experience of farmers up to that point. Finding that this is not the case, they attribute the result to social pressures. Bandiera and Rasul (2006) attempt to separately identify imitation effects by including key individuals as regressors, but do not find evidence of imitation.

Kohler et al. (2001) attempt to distinguish between social learning and social influence empirically in the context of contraceptive use. They do this by analysing the density of the social network and how this interacts with the proportion of network partners using contraceptives. Indirect connections are assumed to increase social influence effects. If social learning dominates, the expected result is that both dense and sparse networks experience an increased probability of family planning use with an increased number of family planning users in the network. The expected result if social influence dominates is that variation in the proportion of family planning users in a sparse network should have a smaller impact on the probability of family planning use than in a dense network. They show social learning to be prevalent where there is high market activity and social influence to be predominant where there is low market activity.

In a widely cited paper Bandiera and Rasul (2006), and more recently Liverpool-Tasie and Winter-Nelson (2012) using a similar methodology, claim to identify a process of social learning specifically, rather than other network effects such as social influence or other confounding factors. Bandiera and Rasul (2006) look at adoption of a new crop, sunflowers, by farmers in Mozambique. Propensity to adopt the technology is considered as a function of the number of adopters in the individual's network. Critical to Bandiera and Rasul's (2006) social learning conclusion is the finding of an inverse U-shape in the number of adopters, such that the propensity of the individual to adopt increases in the number of adopters in the network, up to a point, after which it decreases. Bandiera and Rasul (2006) attribute this to the combination of two effects. A learning externality incentivises adoption as the number of adopters increase because additional adopters in the network increase the available learning informatio,. A strategic delay encourages a delay in adoption when the number of adopters increases because costly own experimentation can be avoided by free-riding on information from the experience of ample other adopters. With a sufficient number of adopters the strategic delay effect could outweigh the learning externality leading to an inverse U-shape of propensity to adopt in the number of adopters. As Liverpool-Tasie and Winter-Nelson (2012) also argue, adoption probabilities that peak and then fall with increasing network size are more consistent with social learning than, for example, mimicry or contextual effect. Although unobserved characteristics could drive the association between network size and probability of adoption the authors suggest this is an intrinsically linear problem such that non-linearity encountered in their work allows for the identification of endogenous network effects. Bandiera and Rasul (2006) confirm that the inverse U-shape is not spuriously generated by the composition effect of differential impact on uninformed and informed farmers. Possible alternative explanations for the inverse U-shape do exist. It could be explained by conflation of network adoption rates with farmer ability, for example, if both the most and least able farmers have the weakest incentives to grow sunflowers and ability correlates with network size. Bandiera and Rasul (2006) dispel this by demonstrating that the result is robust to dropping the most productive cashew farmers from the sample, however it is not possible to eliminate all possible alternative explanations (Liverpool-Tasie and Winter-Nelson, 2012). Liverpool-Tasie and Winter-Nelson (2012) also point out that whilst an inverse U-shape is likely to be indicative of social learning it is possible for social learning could to occur without an inverse U-shape and thus this is not a universally applicable methodology.

Additionally, Liverpool-Tasie and Winter-Nelson (2012), in the context of Ethiopian agriculture, contribute to the literature by considering in greater depth than elsewhere the role of network type in determining social learning. They find social learning to be more evident for more complex technologies, amongst households not in persistent poverty and within networks of intentional relationships (i.e. friendships) rather than those based on physical proximity. The finding of a weaker effect in proximal networks adds an additional argument to Moffitt's (2001) point that it is not always optimal to proxy for a network with geographic proximity. In addition Liverpool-Tasie and Winter-Nelson (2012) include a range of controls, including dummy variables for peasant association and various demographic controls, to attempt to ameliorate possible omitted variables that could confound the learning effect. Their behavioural data are proxy reported, as in Bandiera and Rasul (2006) and others. They attempt to mitigate the criticisms of Hogset and Barrett (2010) regarding the pitfalls of proxy-reported behaviour data by controlling for farmer characteristics, such as age and education, that could affect their ability to properly identify their network sizes. Despite this, results based on proxy reported behavioural data still present a measurement error concern.

Conley and Udry (2010), studying fertiliser use on pineapples in Ghana, use behavioural data that is not proxy reported, and network data at the individual level, rather than proxied by physical proximity. These are two strong characteristics of their work that are also present in my methodology. The consideration of updates in input quantities, rather than a binary variable for adoption more generally, allows for a more specific identification of the learning process. By exploiting the variation in surprising profit outcomes, conditional on growing conditions, by farmers within each network, they identify a process of social learning. In order to do this, a social learning variable is constructed capturing the relative input levels of surprisingly profitable fertiliser adopters within the individuals' networks. Using the assumed spatial correlation of growing conditions a variable is constructed that attempts to control for variation in growing conditions. Additional correlated effects are controlled for using information on various location specific variables, including soil types. Using the predicted network from logistic regressions of network link existence the authors are able to test the robustness of their social learning result to the possibility of network endogeneity.

An avenue not widely explored as yet in the literature is the role of network structure and characteristics in social processes, though Kremer and Miguel's (2007) consideration of the role of distant contacts is relevant for this agenda. Often the focus of social learning studies has been on direct interactions. However, network structure characteristics, such as size, heterogeneity and density, are also relevant (Marsden and Friedkin, 1993) (Valente, 1996). Kohler et al. (2001) find that network structure, specifically density, influences behaviour in the use of contraceptives. Though this dissertation does not directly consider network structure in the learning process, several network characteristics are controlled for, which is made possible by the precise identification of network structure in the data.

2.1 Empirical Challenges to Identifying Social Effects

One common way to test for social effects is to look for correlation between behaviour and social networks, by testing for a positive relationship between behaviour (such as technology adoption decisions) and the density of other adopters in the social network. However, Manski (1993) identifies many possible explanations for correlation of behaviour with networks. These are identified as a) correlated effects, due to similar individual characteristics within network or because individuals face the same local conditions, b) exogenous effects, where group formation is based on the similarity of individuals, for example both using a particular technology, c) endogenous effects, which affect behaviour through social interaction and influence within the group. The social effects literature is interested in the third of these. However, there are substantial challenges in distinguishing between these effects empirically. By way of example, Kohler et al.'s (2001) study on contraceptive use suggests residence in same neighbourhood (giving access to the same family planning programmes) as a possible correlated effect which could confound the network effect in the adoption of contraceptives. Women seeking to gather information from other women who are known to use birth control could be a confounding exogenous factor. It is not easy to observe whether the correlation of contraceptive use with social networks is a social effect or the results of one of these (or other) confounding factors.

Maertens (2010), in a study on the adoption of a new type of cotton in India, succinctly outlines the challenges to an empirical identification of social effects in four parts: Firstly, getting an accurate picture of a social network with a limited sample, secondly, separating social interaction effects from correlated effects, thirdly dealing with the simultaneity problem, where actions are jointly determined in equilibrium, and fourthly constructing identifying assumptions to make it possible to separate out various kinds of social interaction effects, including social learning, social pressures/influence and imitation. The second and third of these are derived from Manski's (1993) theory. The first, regarding network samples, is a product of constraints of data collection and availability in large networks. A case of this is the problem of ill-defined information networks, which often lead to reliance on potentially misleading assumptions. For example, it is often assumed that an individuals information network consists of anyone within close geographic proximity (Moffitt, 2001). The fourth issue, of differentiating between social effects, remains a persistant challenge in the literature. Various identification strategies have been employed to attempt to identify and differentiate between the various social effects, with varying degrees of success.

In broaching the issue identified here as the fourth empirical concern, Hogset and Barrett (2010) are highly critical of the generality with which economists typically approach social processes. Hogset and Barrett (2010) differentiate between social learning (seeking out information from others to construct rational beliefs) and social influence, which they define as harmonising with prevailing beliefs, for example in determining whether a population appears to find an innovation interesting in general. Arguably both effects are very different, with social influence generally based on aggregate observations, and social learning based on observations of individuals, particularly opinion leaders. It is possible for individuals to have access to population aggregates without knowing individual behaviour and vice versa. The distinction is important not only for policy but also in research methodology. These two mechanisms, characterised by potentially very different processes, are often amalgamated by the general consideration of adoption rates within a broad network.

Another, less widely applicable, challenge to empirical work, relates to the way in which network behavioural data are often collected. Hogset and Barrett (2010) criticise the widespread use of proxy-reported behavioural data in the social learning literature, and highlight Bandiera and Rasul (2006) as an example of this. The reporting by individuals on the behaviour of others in their network (for example whether another individual uses a particular technology) can lead to measurement error if adopters (non-adopters) tend to systematically over or underreport adoption rates in their network. Hogset and Barrett (2010) demonstrate empirically that measurement error resulting from proxy reporting can lead to spurious correlation of network adoption rates and adoption, leading to the false identification of social learning processes.

Though my work is not able to deal fully with these empirical concerns, the precise network definition in the data and direct observation of behavioural data (which is therefore not proxy-reported) amerliorates some of the major challenges often faced in the social effects literature.

Identifying the effect as social learning specifically (rather than other types of social effect) remains a challenge. However, the methodology of my dissertation allows for more direct linking of individual behavioural responses with specific individual experiences of peers in the network than is possible in many contexts. Additionally there is greater precision in observation of effects because the behavioural measure is a continuous variable which can move in either direction, unlike the binary response variable (such as adoption decisions) in many social effects analyses. This allows for greater confidence in the opportunity to specifically identify a social learning process, although confounds of other social learning effects remain a possibility.

3 Data Collection and Description

The data used in this dissertation come from two surveys (a baseline survey and a social networks survey) and two rounds of experimental outcomes, conducted between February and October 2012. Data collection took place amongst young people enrolled on a youth leadership programme. The study was conducted over time as the social network formed.

3.1 Baseline Survey

Participants were surveyed within 24 hours of their first arrival at the leadership programme. It is relevant that these data were collected prior to the formation of the social network as this allows for their use in controlling for the endogeneity of the network in the later analysis. Participants were surveyed on their demographic characteristics, education, parental education, employment history, health and emotional well-being.

This section gives an overview of participant characteristics in the baseline survey. Table 1 presents summary statistics of characteristics of network participants. The network is predominantly black African, though the share of coloured individuals in the Western Cape training centre is 13%. Overall there are slightly more males than females in the network, and this is most pronounced in KwaZulu Natal, where 63% of participants are male. In the Western Cape group 58.7% of participants speak isiXhosa as their home language, 10.9% English, 19.6% isiZulu and 8.7% Afrikaans. In KwaZulu Natal the home language of 33.3% of participants is isiZulu, 27.5% isiXhosa and 15.7%

Sesotho. In Gauteng 33.9% speak isiZulu as a home language, 20.3% speak Sepedi and 16.9% Sesotho. The higher proportion of coloured participants in the Western Cape group as well as the variation in language by province is reflective of the varying demographic make-up of the various regions, due to the fact that individuals were recruited to attend the programme in the region closest to their home.

Of all participants 82.7% were 'employed' at the start of the programme. This definition captures employment in its broadest sense, including part-time or casual work, volunteer work paid with a stipend, assistance in another person's business and self-employment, as well as full-time paid employment. Of those listed as 'unemployed' by this definition, many are in fact students. For 45.9% of participants their primary activity one year prior to starting on the programme was 'working for pay' This figure is slightly higher for the Western Cape centre, and lower in the other two programme centres (KwaZulu Natal and Gauteng). Students and full time scholars constituted 24.8% of participants and 12.7% unemployed and actively searching for a job. Those in volunteer positions make up 9.6%, although this is likely to be an underestimate due to the reporting of volunteers in receipt of stipends as 'working for pay'.

Table 2 presents summary statistics on some of the relevant continuous variables of participants in the network. The mean age of participants is 24.7, with a standard deviation of just 3.39, reflecting the target recruitment age for the programme, which was 20 to 30 year olds. The mean level of education of participants is 12.8 years. The mean take-home monthly income (after tax and payment of business expenditures) is R2976. This includes wages from all paid work including casual work, payment from self-employed business activities and any money received from assisting in other people's business activities. Income is not usually normally distributed. It tends to have a long right tail with a few high earners and a large bulk of individuals earning little or nothing. It is therefore necessary to transform the income variable for use in an OLS regression. Due to the challenges of log-transforming a variable that contains zeros, the transformation used here is an Inverse Sine Transformation, which can accommodate zero income values. Summary statistics on the transformed income variable are also shown in Table 4, with the label 'Income (monthly) IST'.

Maternal education is often used as an indicator of socioeconomic background. In the sample the mean years of mother's education is 10.17, with a standard deviation of 3.443. This is lower than the education levels of the participants themselves. Father's level of education was excluded from the regressions (and hence from the summary statistics) due to a lack of data, a problem likely linked to absent fathers.

The 'socioeconomic status (ladder)' variables capture an individual's subjective perception of their socioeconomic status, relative to other South Africans. A ladder of six rungs was described to participants, on which the poorest people in South Africa are on the first rung and the richest people in South Africa on the sixth. Individuals were asked both on which rung their household was when they were 15 years old and on which rung they are today. The mean measurement for self-perceived socioeconomic status aged 15 years is 2.414, with a standard deviation of 1.068. The mean for the current position is 2.904 with a standard deviation of 1.036. The standard deviations on these variables are quite small, possibly indicating that the socioeconomic backgrounds and situations of participants are on the whole fairly homogenous. There is a slight increase in mean perceived relative socioeconomic status between aged 15 and the present, demonstrating a slight improvement in perceived socioeconomic position over time.

The final variable in Table 2 is a dummy variable equal to 1 if an individual considered it either 'likely' or 'very likely' that a lost wallet containing R200 would be returned to them by a complete stranger who had found it. The variable is equal to 0 if the individual considered the wallet unlikely to be returned. This variable measures some aspect of how trustworthy an individual perceives an unknown person to be. This is relevant for consideration in this context given that the perceived 'trustworthiness' considered in this dissertation is that of an anonymous partner. In the estimation sample the mean of the wallet return variable is 0.223, indicating that 22.3% were positive about the prospect of the wallet's return. The standard deviation of this variable is 0.418.

3.2 Trust Game

Experimental games are widely used in the economics literature for eliciting various behaviours in a controlled experimental environment. The trust game is one such game. In this section I discuss the basic concept of the trust game, summarise the literature that has considered trust games in the context of social networks and discuss the trust game data.

The simple trust game, based on the seminal trust game of Berg et al. (1995), captures two mechanisms: selfless altruism and anticipation of reciprocal exchange in the future (Cox, 2004). A basic anonymous version of the trust game is used in this research. Participants are given an endowment, from which they can offer an amount of their choosing to the other player. The offer is multiplied by a predetermined factor and transferred to the other player. This player can then choose how much, if any, they wish to return.

Trust games have been conducted in real world social networks to show that the effect of directed altruism is roughly twice as strong as future interaction or reciprocity effects (Leider and Mobius, 2009, p. 1817), and in ascertaining the role of social ties and network position on insurance opportunity in a variation on the trust game (Chandrasekhar et al., 2011).

Elsewhere in the literature, social interaction within the game set-up has allowed for the consideration of social effects on trust. Lev-On et al. (2010) consider various communication media and group sizes to develop social interaction within the experimental context prior to game play. They observe that an expectation of reciprocation is better predicted by communication group size than by communication medium, with the finding that dyadic communication is most conducive to trust.

Di Cagno and Sciubba (2010) consider network effects directly by administering a game that builds social networks into lab experiments either before or after trust game play. This allows them to differentiate between the impact of social networks on trust through the mechanisms of social interaction and reputation and to compare these two mechanisms. Additionally they distinguish between absolute trust and differential trust, which shows preference to those with whom the individual shares a social link. Though experimental networks enable some observation of network effects on trust they do not capture the depth of relationship and interaction that characterises many real world networks and thus social network experiments bear only limited relevance.

As the data used in this dissertation are based on an anonymous game they do not allow for the investigation of trust in relationship specific contexts, which is the primary concern of the trust-networks literature in general. However, the focus of this dissertation is belief updating more generally, for which data from an anonymous game, which captures the amalgamated effects of general altruism and a response to trust the reciprocal behaviour of an anonymous partner, is sufficient. Diverging from the agenda of the trust game literature more generally, the primary concern of this dissertation is the impact of social effects within networks on perceptions of the trustworthiness of others, rather than on strategic game play within the network.

Where my data are unique amongst the literature is in their longitudinal nature across a real world network as the network forms. First round game play took place on the day that programme participants met, when the network amongst them was sparse and limited to a few prior relationships. The game was repeated at the end of the programme once the social network had formed. The social network was measured in an intermediate period. As discussed above, the anonymity of game partners, who participants were explicitly informed were not a part of the programme, limits observation of the development of personal trust. However, the longitudinal nature of the data does crucially allow for the study of the process of social learning.

The set-up of the specific trust game employed in this research is as follows. Participants are told that they have been given R50 and are asked to decide how much they will give to player B. Though physical cash is not given to the participants at this stage they are made aware that the money exchanged in the task is real and that they will be paid the resultant sum in cash. The identity of player B is not known to player A, however A is told that player B is a real person who is not part of the programme. Participants are told that any amount of their R50 that they choose to send to player B will be doubled before it is given to player B and added to the R50 that player B is also given in the game. Player B then has the opportunity to return any amount they wish to the participant. In practice player B decides in advance how much they will return for each potential offer that the participant could make. Having played the trust game participants were paid in cash later that day, receiving an amount in accordance with their offer and player B's amount returned. It is pareto efficient for the participant to send the entire amount of their R50 and for player B to return at least R50 to the participant. However, it is not rational (from the perspective of maximising monetary outcomes) for player B to return anything to the participant, since there is no repeated game play in this context. Trust is expected to increase the offer made over that which would be made in context of maximising monetary payoff where there was no expected reciprocity. Therefore the behaviour of the player As, who are the subject of the study, can thus be considered as incorporating the extent to which they perceive their anonymous partner as trustworthy. When playing in round 3 participants were not paired with the same player B that they had played with in round 1 and players were made aware of this. Therefore, reciprocity based on repeated play with the same opponent should not impact decisions in the game.

Table 3 reports summary statistics on trust game offers. The statistics in Table 3 and Table 4 are calculated only over the sample of individuals for whom both round 1 and round 3 trust game data are available. This is the estimation sample used in the subsequent analysis.

The mean offer in round 1 of the trust game was R19.22, which decreased to R18.23 in third round play, although the difference is not statistically significant. The amount that individuals expected to get in return for their offer also decreased from R29.31 to R27.28. This is probably in part due to the lower mean offer made, for which participants are likely to expect a lower average amount returned. Table 4 shows the proportion of individuals who increased, decreased or did not change their offer between the two rounds of play. Between the two rounds of play 26.1% did not change their offer, 32.5% increased their offer and 41.4% decreased their offer. The objective of this dissertation is to indentify whether these changes, including their magnitude and direction, or the absence of change in offer, can be explained by a process of social learning.

The expected own profit of individuals (calculated here as the difference between the amount sent to player B and the amount the individual nominated as expecting to be returned by player B) also decreased from R10.18 to R8.98. Of all participants 27.4% expected to receive zero profit on their offer to player B, 54.1% expected to make a positive profit and 17.8% expected to make a negative profit. That some people play offers for which their subjective expectations of profits are negative would seem to suggest that there is some degree of altruism that goes beyond an attempt purely to maximise profits in the game. Outcomes in this regard were very similar in round 3 play, with 53.5% expecting a positive profit, 26.8% a zero profit and 19.7% a negative profit. The influence of pure altruism rather than pure profit-seeking behaviour could dilute the observed social learning effect, if the trustworthiness of the partner (the factor determining profit potential), is not the primary concern of player A when making an offer decision.

3.3 Attrition

There is a relatively high level of attrition in the sample between rounds 1 and 3. Table 5 reports the differences between the two groups along various characteristics, together with their standard deviations. Of participants who played the trust game in the first round 28.4% were not present for play in round 3. If this attrition is non-random there is a possibility that results will be biased. The group who left the sample are not significantly different from the sample group in terms of age, years of education, racial composition (though the change in the proportions of some racial groups is nearly significant) or trust game variables from round 1 play. There is a significant difference between the two groups in terms of income, the mean of which is R2322 higher for those who left the sample, relative to the mean income of R2976 in the sample. The difference in proportion of the groups who are employed (in the broad sense of the term) is almost significant, with an 8.8 percentage point higher share of employed individuals amongst those who left the sample. There is also a significantly higher proportion of home language English speakers amongst those who left the sample, with a mean of 14.8% relative to 5.8%. Those who left the sample also have significantly higher levels of maternal education, a mean of 7.0 years relative to 5.1 years of education. Some of these differences are expected, in the sense that they relate to an individual's opportunity cost of remaining in the programme. This is particularly obvious in the case of income and employment status, although less clear for home language English or maternal education, though these may also proxy for opportunity cost of programme attendance.

Results observed in the sample remaining in the programme may not be reflective of those of the full population, since those who left the programme are inherently different. There may therefore be some bias in the results, though the direction of this potential bias is not clear since there is not clear theoretical justification for individuals with a higher income, for example, being more or less inclined to engage in a process of social learning with regards to trustworthiness.

3.4 Social Networks Survey

Social network data were collected in a survey conducted at the start of the second gathering of programme participants (that is, the second of three 10-day residential sessions). Each participant in attendance at the second gathering was shown photos, each labelled with a name, of all participants who were part of the programme in their region including those from the two other programme groups in the region. Respondents were asked of each person whether they had ever communicated with that person. The use of photos overcame potential difficulties in data collection that may otherwise have arisen due to the fact that some participants go by names other than the name they gave on their application form for the programme. 'Communication' was defined to respondents as any form of communication, including via phone or internet, for example, but only including direct communication and not recognition of a face or mutual presence at an event.

Having identified the set of individuals with whom the participant had ever communicated they were asked further questions about their relationship with each individual they had identified as a communication link. These questions included whether they had known the person prior to the programme, their frequency of communication outside of programme sessions, conversation topics, approaching to work together on a project or for assistance with a project, strength of relationship and language of communication.

Participants were then asked to nominate anyone else who they knew who was taking part in the programme whom they had not already been asked about, for example participants who were part of the programme in another province. Only 15.3% of participants nominated anyone outside of their region, and of these the mean number of additionally nominated contacts was 1.9. The maximum number of contacts nominated outside of the individual's programme region was 4.

Social network data are often sampled, such that the full network is not observed. This is usually necessitated by large networks, networks of high average degree or networks without clear boundaries (Wasserman and Faust, 1994, p. 34). Networks can be randomly sampled or alternatively sampled in a rolling manner, where a random sample of the population is surveyed on their links and thereafter their nominated contacts are contacted and also surveyed on their links. This sampling method has been discussed widely in the literature, for example by Frank (2005), and tends to be used when the research question considers network centrality (Fafchamps and Gubert, 2007, p. 331). Additionally research in the small worlds and 'strength of weak ties' literature (Granovetter, 1973) is aided by ego-centric sampling of this type. However, multi-stage sampling presents inference problems because of non-random selection of individuals after the first stage.

Additionally network surveys often limit the number of contacts that respondents can nominate. For example, in the widely used AddHealth data) individuals were asked to nominate their five closest male friends and five closest female friends (Alexander et al., 2001). This process also restricts the collection of a complete network and can lead to concerns regarding the stability of some centrality measures, amongst other issues (Costenbader and Valente, 2003).

The network sample used here surveys the full relevant network and does not restrict the number of nominations individuals can make. Given the relatively small size of the network (each programme region had less than a hundred participants) it was possible to ask participants about everyone they may potentially know within their training region and to allow them to nominate as many additional participants as they knew in other regions. This mitigates most of the problems that arise from sampling procedures and negates considerations of dealing with the biases that arise from sampled network data. Inference from the network data is therefore more straightforward than it might otherwise be.

The sampling boundaries of the network data gathered for this dissertation could potentially have been extended to include friendships outside of the programme. However, in addition to the practical considerations of limiting data collection to programme participants, it makes sense to restrict the sample in this context to those who would have been exposed to trust game play. Social learning relating to trust game outcomes is likely to be well-captured solely within the programme participant network, since it is unlikely that individuals would know others outside of the programme who have taken part in the trust game. This is less clear cut for the evolution of pro-social behaviour more broadly, which may be influenced by an individual's wider social network. However, that is beyond the scope of this study to consider.

4 Network Characteristics

This section seeks to paint a descriptive picture of the social network, as well as to present definitions and summary statistics for the relevant network centrality measures that feature as control variables in subsequent analysis. The network considered here relates in all cases to responses to the question 'Have you ever communicated with...?'. Though other questions were asked, characterising a variety of networks, for example the network of requested assistance on community projects, the network description here is limited to the communication network, since that is the network considered in the social learning analysis.

There is a large literature, spanning economics, sociology and physics, that considers the theory of networks. Wasserman and Faust (1994), a prominent work in this field, brings together the key areas of social network analysis, and is the primary work on which the descriptive network section of this dissertation is based. Various statistics, including network density, geodesic distance, centrality and so forth, characterise features of networks and can be used in explaining various outcomes that may result from network behaviour. O'Malley and Marsden (2008) offer a general synthesis of this field.

4.1 Graphical representations of the Network

The network here is presented in the form of sociograms, graphical representations of the network where each individual is represented by a point (or node) on the graph and identified by a unique number. Each communication relationship is represented by a line (in the case of a graph of reciprocal relationships) or an arrow in the case of directed relationships.

Figure 1 shows each nomination of a communication relationship with another individual, giving the full network of all individuals across all three programme regions. The three clear clusters within this network reflect the three programme regions, with the smaller clusters within these reflecting the three programme groups within each region. The cluster to the top left of the diagram is the KwaZulu Natal group, the cluster to the right is the Western Cape group and the cluster to the bottom left, the Gauteng group. This graph reflects all nominations of communication relationships, even if that link was not reciprocated. The direction of the nomination is reflected by the arrows in the graph. Beyond very small networks it becomes impossible to represent social distance by line length in a graph in two dimensions; therefore the distance between nodes is determined by the practicality of layout. Distance between two individuals on the graph is therefore not reflective of a socially distant relationship, in terms of geodesic distance (the number of links traversed to reach each other by the shortest route).

Figure 2 depicts the same individuals as Figure 1 but shows only reciprocal communication links. A link is only shown to exist where both individuals claim to have communicated with each other. For some purposes reciprocal links might more realistically represent relationships, as a non-reciprocal 'friendship' is unlikely to be meaningful. However, for the purposes of low-value transactions, such as basic information dissemination and social learning, a one-way communication could still be relevant. Information transmission could plausibly take place along the unidirectional link.

When links are limited to reciprocal relationships, as in Figure 2, the network is sparser. There are very few reciprocal links between provinces and none at all between the Gauteng region and the other two regions. Additionally, the three groups within each region become more well-defined when only reciprocal links are included, as there are fewer reciprocal links between groups. The contrast between Figures 1 and 2 highlights the extent to which the directed network consists of one-way relationships. Given that this could impact results, if learning in fact only takes place in stronger reciprocal relationships, the results presented in the dissertation were replicated using the reciprocal network definition and are available on request.

As there are relatively few links between the programme regions hereafter the network is described at the regional level (i.e. the three provinces in which the programme groups were based) for greater clarity of network depiction. Figures 3 to 5 depict the regional networks, showing the directed (non-reciprocal) relationships only. Nodes are coloured in accordance with the individual's group in the first round. In all regions there is clear clustering evident by group, though there are also a reasonable number of inter-group links. To distinguish between groups the second digit of participant unique number represents the group that the participant was a part of in round 1 (the first digit indicates the region in which they attended the programme).

4.2 Demographic Characteristics in the Network

4.2.1 Race and Gender

Figures 8 to 10 show graphical representations of the regional networks indicating the race and gender of each individual in the network. Figures 6 and 7 indicate the key used in colour-coding the nodes for race and gender respectively.

In the Western Cape network, presented in Figure 8, the key nodes connecting the three groups are all African, and both males and females are represented among these. Within the groups there are very well connected individuals of all race and genders including notably several African women and coloured women.

The Kwazulu Natal network, shown in Figure 9, is predominantly African, however the small minority of coloured and Indian individuals in the network are extremely well connected, in particular in terms of their indegree, the number of individuals who report to have communicated with them. Contrary to the Western Cape network, the key individuals connecting the three groups within the network are female.

In the Gauteng network in Figure 10 there is only one individual who is not African and several for whom race data were not obtainable. The one Indian male in the network is very well connected, though many other group 3 members are also, including a core group predominantly made up of African females.

4.2.2 Home Language

This section presents the regional network diagrams according to home language. Figure 11 gives the key to the colour-coding used in the home language network diagrams.

In the Western Cape network in Figure 12 the key individuals connecting the three groups are Xhosa speakers. Within group 3 the key nodes are Xhosa and Zulu speakers. This is likely due to the higher proportion of Zulu speakers in this group than in other groups, due to a transfer of participants from the Kwazulu Natal base, which may have allowed the Zulu language speakers to characterise the group. In group 2 there are a number of well-connected Xhosa speakers but also three Afrikaans speaking individuals who play a key role in the network. Group 1 is dominated by Xhosa speakers. Afrikaans and English speaking individuals appear more peripheral in the network, with the notable exception of one English speaking individual who is well connected both to group 1 and group 2 participants.

The Kwazulu Natal network in Figure 13 has a greater diversity of languages, dispersed unevenly across the three groups. In group 1, where there are predominantly Xhosa and Zulu speakers, the Zulu speakers are the key nodes within the network. The one English speaking individual is well connected within the group as well as between groups. In group 3, where there are a number of Sotho and Afrikaans speaking individuals, the distinction of key nodes by language is less apparent. One Afrikaans speaker is very well connected, particularly in terms of being nominated by others. Around half of the Zulu speakers in group 3 are central whilst the other half are somewhat more peripheral. In group 2 there are well connected individuals who speak Sotho, Engish, Xhosa or Zulu, but no clear pattern. The nodes connecting the network more widely are predominantly Xhosa speaking.

The Gauteng network in Figure 14 is more linguistically diverse than the other two regional networks. The key nodes connecting the network as a whole are predominantly Zulu or English speaking. Within group 3 there appears to be some clustering of individuals according to language, with a cluster of Sotho speakers and some clustering of Zulu speakers evident. This does not appear to be the case in groups 1 and 2, where languages are well dispersed among the network. The key nodes in groups 1 and 2 are mainly Zulu, though there are also some very prominent Sepedi and Xhosa speakers.

4.3 Density

The density of a network reflects its completeness in terms the proportion of potential links in the network that do in fact exist.

$$\Delta = \frac{m}{n(n-1)} \tag{1}$$

Equation 1 gives the formula for the density of a directed graph. Density is a fraction that ranges from 0, where there are no links, to 1, a complete network where each individual is connected to each other individual in both directions. It would be unrealistic to expect a complete network across the combined three groups in each region or in the full network across regions, since participants across the full network were not required to attend programme events together. However, the programme could hope to achieve a complete network within the subgraph for each group, especially since the relationship under consideration is ‘ever communicated’, reflecting a very shallow level of relationship.

Table 6 shows the densities for the three regional networks, both for the regional network as a whole, and for the individual groups within the regional. The separate programme groups are small and participants would have met each other through programme activities, so at this level there is potential for densities to be near to 1. The overall densities of the three regional networks are similar, with approximately 30% of all potential ties that could exist in the network actually existing. The highest density is observed in the Kwazulu Natal network, where 33.6% of potential ties exist. The densities for the group networks range between 0.584 (implying that only 58% of the potential links actually exist) for Gauteng group 3, to 0.977 (implying that 98% of potential links exist) for Western Cape group 3. This variation is in large part due to the differing group sizes represented, which complicates direct comparison.

Network densities can reasonably be compared across these groups because each regional group is of a similar size, although smaller groups are likely to have a higher density due to the smaller pool of potential links. The size of a network affects the realistic density that could be achieved. For example, in a group of ten people it is reasonable to expect a complete communication network, but in a group of ten thousand a complete communication network would be virtually impossible. Comparison to values from the wider literature presents a greater challenge, as networks reported in the literature vary greatly in size and nature.

For interest, the densities of the reciprocal network are presented in Table 7. As expected, the densities of the reciprocal communication network are lower than those of the directed communication network.

4.4 Network size and closeness

Granovetter's (1973) theory of the strength of weak ties suggests that it is often connections with acquaintances rather than close relationships that are most relevant for economic outcomes. For example, whilst an individual's close friends may be largely similar to them (and exposed to similar information and resources) an acquaintance may know of a suitable job opportunity that the individual would otherwise not have encountered. Networks benefitting most from the strength of weak ties are those that are less closed. Karlan et al. (2009, p. 1330), however, show that whilst weak ties may be most relevant for the transmission of low-value assets, such as information about job openings, for high-value assets, such as those where trust is required, strong ties are the relevant consideration. Therefore, high-value asset transactions may be better supported by closed networks. Low-cost information as well as highly invested relational transactions could be relevant for perceptions of trustworthiness, thus both high-value and low-value asset transactions are relevant considerations here. However, it is not clear which is more relevant for social learning of preferences and pro-social behaviour. Whilst low-value asset transactions allow for widespread dissemination of information on trust game experiences and trustworthiness of partners, a sense of trust and pro-social behaviour in general may be fostered better in a context suited to high-asset transactions.

Karlan et al. (2009) connect network size to how closed a network is, concluding that small networks tend to be more closed. The link between small networks and high-value transactions has been used to explain the negative correlation observed between community size and prosocial behaviour including volunteering, working on public projects and helping friends (Putnam, 2001).

4.5 Centrality

This section reports summary statistics on various centrality measures in the network. These are important as they are used in later analysis to control for individuals' network connectedness.

There are many, and divergent, measures of centrality. Measures of local centrality consider whether a point is connected with many point's in its immediate environment. Measures of global centrality consider a points position of strategic significance in the overall structure of the network. I focus here on three measures of centrality, one local (degree) and two global (betweenness and eigenvector centrality). Importantly, centrality is measured at the individual level, not the network level, which lends these measures to controlling for individual network position in the later analysis. Table 8 presents summary statistics on these centrality measures for individuals in the network sample.

4.5.1 Degree

Degree is a basic measure of local centrality indicating how many individuals each person is connected to (in this case, has ever communicated with). As the data here are directional, both indegree and outdegree are reported and indicate the number of individuals who nominated that person and the number of individuals whom the person nominated, respectively.

As shown in Table 8 the mean indegree is 19.3 and the mean outdegree 19.5. The standard deviation of outdegree (6.10) is greater than that of indegree (5.32) implying greater variation in the number of individuals nominated than number of individuals a person is nominated by. Similarly the range of indegree values is smaller than of outdegree. Outdegree may in part reflect someone's enthusiasm for the task of nominating communication links and also the rigidity with which the concept of communication was applied. This could perhaps be the cause of greater variation in outdegree than indegree.

The degree distribution for both indegree and outdegree is presented in the kernel density diagram in Figure 18. The degree distributions of social networks are typically characterised by power law distributions, which represent scale-free networks. The distribution here is slightly right-skewed, though fairly close to symmetrical, and with thin tails.

Normalised degree circumvents the problem that comparison of raw degree score across networks of different sizes can be misleading. The normalised measures are standardised irrespective of network size. Normalised degree is calculated as degree multiplied by $\frac{100}{n-1}$, representing a percentage of the maximally possible centrality, and can be used to compare across networks of different sizes. Summary statistics on normalised degree, which is used as a control variable in later analysis, are also reported in Table 8.

4.5.2 Betweenness

This section considers the centrality (or importance) of players in the networks, according to their betweenness centrality. Betweenness is a measure of point centrality that considers the extent to which a particular point lies 'between' the various other points on a graph. A point with low degree can still have a high betweenness centrality if it plays an intermediary role in the network. Betweenness proportion is measured as the proportion of geodesics (shortest paths between two points) that pass through the point. The pair dependency is then defined as the sum of the betweenness proportions of Y for all pairs that involve X. The overall betweenness of a point is calculated as half the sum of the values in the columns of the matrix showing the dependence of each row element on each column element.

Network nodes with a high betweenness centrality are those who, if removed from the network, would cause the network to fragment into more components. Though these individuals are not necessarily 'popular', their absence has the potential to sever potential chains of communication along which information could be conveyed.

As shown in Table 8, mean betweenness in the sample is 35.20, with a standard deviation of 61.13. Figures 15 to 17 show the overall communication network and regional communication networks with the betweenness centrality of individuals reflected in the size of the nodes. Those who are most central in terms of betweenness have the largest nodes.

4.5.3 Eigenvector Centrality

Eigenvector centrality is a measure of how connected an individual is to other well-connected individuals and thus is often used as a measure of power or popularity (see, for example, Mihaly (2007)). The mean eigenvector centrality in the sample is 0.095, as shown in Table 8, and the standard deviation is 0.088. Figures 19 to 21 show the regional networks with the size of each node determined by its eigenvector centrality.

4.6 Comparison with social networks in the literature

By way of comparison across networks in the literature Newman (2003, p.182) references networks of various types, detailing descriptive statistics on these networks. The social networks section of Newman's table is replicated in Table 9, with the addition of three rows at the bottom of the table featuring the three regional networks from this research.

Table 9 presents the following statistics on the networks: total number of vertices n (the number of people in the network); total number of edges m, the number of relationship links; mean degree z, the average number of connections an individual has; mean vertex-vertex distance l, the average distance between individuals in the network; exponent α of the degree distribution if the distribution follows a power law (or "-" if not. In/ out-degree exponents are given for directed graphs); clustering coefficient $C^{(1)}$, from equation 18; clustering coefficient $C^{(2)}$ from equation 19, which both reflect the extent to which the network bunches into smaller groups; and degree correlation coefficient r.

$$C^{(1)} = \frac{3 \times number\ of\ triangles\ in\ the\ network}{number\ of\ connected\ triples\ of\ vertices} \tag{2}$$

$$C^{(2)} = \frac{1}{n} \sum_i C_i \tag{3}$$

The huge diversity in size and nature of the networks listed means that no clear meaning can be given to direct comparison on any of these statistics in attempting to determine the nature of the networks. Therefore the table is included for general interest only.

5 Analytical Framework

The analytical framework is based on a simple learning model derived by Conley and Udry (2010). Since my context is somewhat different, I present my adaptation of their model here.

In period 1 each participant, i, makes an offer to a 'player B' (outside of the network of participants) with whom he is paired in the trust game. In the second period i receives a payout from player B in response to the offer made to player B in the first period. Similarly payouts from game play in round 3 are received in round 4. The trust game was played in rounds 1 and 3 only [1].

The outcome i receives from his partner, player B, in period 2 or 4 is:

$$y_{i,t+1} = f(x_{i,t}) + \epsilon_{i,t+1}, t = 1, 3 \tag{4}$$

where $f(x_{i,t})$ depends only on the offer made in the previous round, and represents 'pure trustworthiness', the object of learning. Epsilon encapsulates 'trembles' in player B behaviour, which incorporates any uncertainty about player B's context that idiosyncratically affects his response, in line with the game theory literature on trembles (Rajan, 1998).

i's total profit from his offer is equal to the response he receives from player B minus the amount, x, that he paid to player B initially. It is realised in periods 2 and 4 is equal to:

$$\pi_{i,t+1} = f(x_{i,t}) + \epsilon_{i,t+1} - x_i \tag{5}$$

Observing profits across his network, i takes the expectation of profits at offer level $x_{i,t}$, such that:

$$E_{i,t+1}(\pi_{i,t+1}(x_{i,t})) = E[f(x_{i,t}) + \epsilon_{i,t+1} - x_{i,t}] \tag{6}$$

Therefore, since $\epsilon_{i,t+1}$ has a mean of zero:

$$E_{i,t+1}(\pi_{i,t+1}(x_{i,t})) = f_{i,t}(x_{i,t}) - x_{i,t} \tag{7}$$

Similarly, for any other $\tilde{x}$:

$$E_{i,t+1}(\pi_{i,t+1}(\tilde{x}_t)) = f_{i,t}(\tilde{x}_t) - \tilde{x}_{i,t} \tag{8}$$

In time t, i chooses his offer to player B, $x_{i,t}$, to maximise his subjective expected profits in the subsequent period, such that:

$$E_{i,t}(\pi_{i,t+1}(x_{i,t})) = f_{i,t}(x_{i,t}) - x_{i,t} \geq f_{i,t}(\tilde{x}) - \tilde{x} \tag{9}$$

[1] In practice payment was received in round 1 but simultaneously and after all participants had played. Thus the subscript 2 for payouts is to incorporate the fact that information on payouts was only available after all participants had played and received payment, such that learning could not take place before that stage.

where $f_{i,t}(x)$ is i's subjective expectation of the return on x at time t and $E_{i,t}(\pi_{i,t+1})$ is i's subjective expectation of profits in the subsequent period.

After profits are realised by all players in period 2, there is a process of social learning. Each player i observes both the offer and profit level of each person j who is in his information network. Thus, for j we have (analogous to equation (5)):

$$\pi_{j,t+1} = f(x_{j,t}) + \epsilon_{j,t+1} - x_j \tag{10}$$

Observing j's outcome of $\pi_{j,2}$ from an offer of amount $x_{j,1}$, i updates his subjective expectation of the trustworthiness function from $f_{i,1}$ to $f_{i,3}(x_{j,1})$, where $f_{i,3}(x_{j,1})$ is equal to the profit that j experiences in period 2.

In summary, the update in i's expectation of player B's response function is:

$$\Delta f_{i,3}(x_{j,1}) \equiv f_{i,3}(x_{j,1}) - f_{i,1}(x_{j,1}) \tag{11}$$

in response to observing event $(\pi_{j,2}, x_{j,1})$.

5.1 Assumptions

There are various models of learning in the literature (see Chamley (2003) for a discussion of these). However, I do not restrict learning to a particular model but rather make assumptions about the learning process, which bear certain implications for the empirical model. This is in keeping with the model of Conley and Udry (2010).

Assumption 1 $\Delta f_{i,t}(x_{j,t-1})$ *has the same sign as* $\pi_{j,t}(x_{j,t-1}) - E_{i,t-1}(\pi_{j,t}(x_{j,t-1}))$ *and increases without bound as* $\pi j, t(x_{j,t-1})$ *exceeds* $E_{i,t-1}(\pi_{j,t}(x_{j,t-1}))$.

This assumption states that if j's profits are higher than i's expectation of j's profits in the first period then the learning effect of this is positive on i's change in offer between the two periods. Similarly if j's profits are lower than i's expectation of j's profits then the learning effect of this on i's change in offer between the two periods will be negative. The extent to which profits exceed expectations is reflected in the magnitude of the change in offer in the third period.

Assumption 2 $\Delta f_{i,3}(x) = 0$ *for all* x *other than* $x_{j,1}$.

An observation by i of j's profit level at offer level $x_{j,1}$ only affects i's beliefs about the profitability of that offer level.

5.2 Implications

Implication 1 *Observation of profit above expectation leads to adjustment in offer towards the surprisingly successful offer.*

Profit higher than expected implies: $\Delta f_{i,t}(x_{j,t-1}) > 0$

Higher than expected profits at a particular input level make it less likely that the offer will change from that level in the next round.

Implication 2 *Observation of profit below expectation leads to adjustment in offer away from the surprisingly unsuccessful offer.*

Profit lower than expected implies: $\Delta f_{i,t}(x_{j,t-1}) < 0$

This implication is not indicative of the theoretical direction of magnitude of change in response to the surprisingly unsuccessful offer.

Implication 3 *If the profit is higher than what was expected the offer will adjust to the surprisingly successful level. Lower than expected profits at a particular input level make it more likely that the offer will change from that level in the next round.*

$$\Delta(x_{i,t}) = 1\{\pi_{j,t}(x_{j,t-1}) - E_{i,t-1}(\pi_{j,t}(x_{j,t-1})) > 0)\}(x_{j,t-1} - x_{i,t-1}) \tag{12}$$

Equation 12 gives the expression for the change in offer implied by the theoretical model. The logical condition holds when j's profit from the previous round of trust game play, minus i's subjective expectation of j's profit at j's offer level in the first round, is greater than zero. This is a 'good news' occurrence, where profits are higher than expected. When the logical condition holds the change in offer, the left hand side of Equation 12, is equal to the difference between j's offer in the first round and i's offer in the first round. That is, i, when observing j's surprisingly successful outcome, adjusts his offer in the subsequent round to the offer that j made in the first round.

6 Empirical Models

Two basic empirical models are considered, followed by an extension to the second model that deals with the problem of network endogeneity.

The first model is a logit model that considers predictors of whether there is a change in offer between periods 1 and 3, using as predictors the shares of good and bad news at offer levels the same as and different from i's period 1 offer. Social learning is determined by considering the signs on the variables for share of good and bad news at the same and alternative offer levels within the individual's network.

The second model considers the precise outcomes of individuals in the network. The surprisingly good outcomes experienced by peers within the individual's network are incorporated into a social learning variable, M, which has a sign and magnitude in accordance with the average of successful offers made amongst peers, relative to i's offer. The update in individual i's behaviour between rounds is considered as a function of the social learning variable and a range of control variables. The model allows for the identification of learning effects by exploiting the variation in exposure to surprisingly successful offers within the network and considering the impact of this, not only on the individual whose offer was successful but also in the network more widely.

6.1 Basic model 1: Logistic model for whether i's offer changed between periods

This model tests implications 1 and 2 of the theoretical model (that offers move away from surprisingly unsuccessful offers and towards surprisingly successful offers).

This is a logit model of the probability that i's offer to player B changes between periods 1 and 3, that is that $\Delta x_i \neq 0$. The dependent variable here is Δx_i, where Δx_i is defined as:

$$\Delta x_i = x_{i,3} - x_{i,1} \tag{13}$$

$$\begin{aligned} Pr\{\Delta x_{i,t} \neq 0\} = \Lambda[&\alpha_1 S(good, x = x_{i,t-1}) + \alpha_2 S(good, x \neq x_{i,t-1}) \\ &+\alpha_3 S(bad, x = x_{i,t-1}) + \alpha_4 S(bad, x \neq x_{i,t-1}) + (controls)'\alpha_6 \end{aligned} \tag{14}$$

$S(good, x = x_{i,t-1})$ is the share of observed events within i's information network in the previous periods that were good news events at input level $x_{i,t-1}$. Good news is defined as profits above expectations, where expectation is defined as the median profits received by participants within the same group in the programme whose offer was within R5 above or below the offer. According to implications 1 and 2 of the theoretical model we expect $\alpha_1 < 0$, $\alpha_2 > 0$, $\alpha_3 > 0$ and $\alpha_4 < 0$. Intuitively we expect $\alpha_1 < 0$ because when i's offer is the same as her network neighbours' and the share of good news events at this offer level increases among this network, the probability of switching goes down.

6.2 Basic model 2: Model for change in input level

This model tests implication 3 (that a sufficiently surprising positive profit outcome from offer $x_{j,1}$ induces individual i to change their offer to the amount $x_{j,1}$ in round 3).

In this model a variable M is defined to capture the social learning opportunity of individual i by considering his exposure to surprising events from which he can learn within his information network.

The basic specification of the model is:

$$\Delta x_i = \beta_1 M_i + (controls)'\beta_2 + v_i \tag{15}$$

M is the empirical analogue of:

$$1\{\pi_{j,t+1} - E_{i,t}(\pi_{j,t+1}(x_{j,t})) > c_{i,t+1}(x_{j,t}))\}(x_{j,t} - x_{i,t}) \tag{16}$$

This captures the anticipated impact that good news observed by i in round 2 should have on his decision to amend his offer in period 3 play. The logical condition is equal to 1 when j's profit in period 2 exceeds i's expectation of j's profits at offer level $x_{j,1}$. When this logical condition holds the social learning variable is weighted by the magnitude and direction by which j's period 1 offer differed from i's. When j's period 2 profit is less than i's expectation at that offer level the logical condition does not hold and M is equal to 0. Under this definition M should be large and positive when i observes a surprisingly high profit experience in his information network, where the offer made by j was substantially higher than the offer made by i. Similarly M will be large and negative when i observes a surprisingly high profit experience by individual j whose period 1 offer was substantially lower than i's period 1 offer.

Empirical analogue of the previous equation:

$$M_i = [1\{\pi_{j,2}(x_{j,1}) - \hat{E}_{i,1}(\pi_{j,2}(x_{j,1})) > 0\}](x_{j,1} - x_{i,1}) \tag{17}$$

As $E_{i,1}(\pi_{j,2}(x_{j,1})$ (i's expectation of j's period 2 profits) is not observed, the estimated value $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ takes its place. My data affords many options for how one might proxy for this expectation.

6.3 Definitions of the social learning variable M

M is broadly calculated as in equation 17. $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$, expected profit at each offer level, is calculated as the median of the profits made by all individuals within i's training group whose offer in period 1 was 'close' to the offer made by j in period 1.

'Close' offers are considered to be all those between R5 below and R5 above j's offer in the primary M definition. However, results are verified to be robust to changes in the definition of closeness in the expected profit definition.

The definition of M in equation 17 is for the case when each individual observes only one other good news event within their information network. As most individuals observe more than one good news observation in their network, M is further defined as a weighted average of the M values derived from each good news event observed.

Eight variations on the definition of M are considered. The variation in the definitions is along three dimensions. Firstly, the definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$, the threshold expectation which determines whether each event is sufficiently profitable to be considered 'good news'. Secondly in the definition of i's information network (i.e. the events that i can potentially learn from). Thirdly in the weighting of each 'good news' event in i's network in determining the overall value of M_i.

$\mathbf{M_1}$ is the defined relative to the 'expected profits' of j based on the median profits of all individuals within i's group in the programme whose offer in the first round was within a range of R5 above or below j's offer. M_1 is then calculated as the mean of all M values of individuals in i's communication network who experienced good news events. M_1 is the standard definition of M used throughout, unless mentioned otherwise.

$\mathbf{M_2}$ is defined as M_1 but with a greater range used in calculating $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$. All offers within R10 above of below j's offer and within j's group in the programme are included.

$\mathbf{M_3}$ is as M_1 and M_2 but with a range of R2 above and below j's offer.

$\mathbf{M_4}$ classifies all offers as either low (R0 to R15), medium (greater than R15 up to and including R30) or high (greater than R30 up to and including R50). $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ is then calculated as the median profit amongst those in j's programme group whose round 1 offer was within the same bracket as j's offer. Except in the definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ M_4 is otherwise calculated in the same way as M_1.

$\mathbf{M_5}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ as used in M_1 but uses an alternative definition of i's information network. Where the previous definitions of M assumed that i could learn from anyone with whom he nominated as a communication link, M_5 restricts i's observed events to those within his network of contacts with whom he said he discussed the programme outside of programme activities. In addition to the standard communication question, respondents were asked whether they talk to the person about programme activities, such as tasks and activities, outside of the times when they were attending the residential programme sessions. Positive responses to this question could reflect a level of relationship where trust game experiences may plausibly have been discussed, given that the trust game was an activity that took place during a residential programme session. It is therefore a reasonable hypothesis that the 'spoke about programme outside programme sessions' network could be highly relevant as an alternative network definition for learning about trust game experiences, hence its inclusion in the M_5 definition of the social learning variable.

$\mathbf{M_6}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ and the same information network as used in M_1 but weights the good news observations differently. Whereas M_1 weights each observation of a good news event equally, M_6 weights each good news observation by the strength of relationship (between 1 and 10, where 1 is a very weak relationship and 10 is a very strong relationship). Strength of relationship is determined by i's answer to the question regarding how close his relationship with individual j is (on the aforementioned scale).

$\mathbf{M_7}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ and the same information network as used in M_1 but weights each good news observation according to the magnitude of the good news. Each good news observation is weighted by $\pi_{j,2}(x_{j,1}) - \hat{E}_{i,1}(\pi_{j,2}(x_{j,1})$, such that the greater a positive surprise the news was the greater it weighs in i's social learning.

$\mathbf{M_8}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ as M_1 but weights all observed good news outcomes (for every participant, not just those in i's actual network) by the predicted values from the network predictive logit model shown in column 6 of Table 11. The reason for this will be discussed in section 7.2.

Table 12 presents summary statistics on the M variable according to the various definitions. Additionally, summary statistics on the number of 'good news' events observed by i under each M definition are presented.

7 Empirical Results

In trust game play 74.3% of participants changed their offer between the first and third rounds. This result is shown in Table 10. The empirical results in this section attempt to explain participants' behavioural changes as the result of a process of social learning.

7.1 Basic Models

7.1.1 Basic Model 1

The key results of the first empirical model, a logit model of predictors of a change in offer, are displayed in Table 13. The dependent variable in all regressions in this table is a dummy variable equal to 1 if the individual changed the amount of their offer between rounds 1 and 3 and equal to 0 if the same offer was made in both rounds. This variable for binary change is regressed on variables characterising the shares of good news and bad news amongst those in the individual's network who made the same offer or alternative offers to the offer that the individual made in the first round. The 'share' variables are defined more precisely as follows: $S(good, x = x_{i,t-1})$, $S(good, x = x_{i,t-1})$, $S(good, x \neq x_{i,t-1})$, $S(bad, x = x_{i,t-1})$ and $S(bad, x \neq x_{i,t-1})$. The empirical model is described fully in the previous section. Summary statistics on the share of good news and bad news experiences observed within individuals' social networks are presented in Table 10 alongside summary statistics on the dependent variable.

According to the theoretical model we would expect the share of good news at the same offer level to have a negative coeffcient, as observations of a suprisingly successful outcome at i's first round offer level should make the individual less inclined to change their offer. Analagously we would expect the share of bad news at the same input level to have a positive coefficient, the share of good news at alternative offer level to have a positive coefficient and the share of bad news at alternative offer levels to have a negative coefficient.

Column 1 of Table 13 shows the results of a basic regression with no additional controls included. The signs on the good news share at the same offer level and the bad news share at different offer levels are both negative, as expected, and the former significant. The remaining two variables of interest both have negative but statistically insignificant coefficients. Marginal effects for all the regression results in Table 13 are presented in Table 20. Throughout this dissertation whereever marginal effects are reported they are calculated as average partial effects, rather than marginal effects at the mean. This is because all specifications for which marginal effects are calculated include dummy variables and it is intuitively easier to interpret average partial effects in this case than to think of the marginal effect at mean values of a dummy variable.

Columns 2 and 3 control for a variety of demographic and socioeconomic characteristics of the individual.Though no control variables are significant, their inclusion changes the signs of some of the variables of interest, though no additional variables become statistically significant.

Controlling for the observable characteristics of indviduals allows for differentiation between confounding factors and the social learning effect. For example, if education levels are correlated with the exposure to good and bad news shares at same or alternative offer levels and if educated individuals are more likely to change their offer, then the social learning effect will not be fully evident, or may be overestimated, depending on the direction of the correlation of education with the news share variables. In addition to this, network characteristics may confound the effect of the news share variables. For example, a particularly popular individual may have a particular pattern of network news exposure and may also be more likely to change their offer, perhaps due to an inherent personality trait that causes them to more readily amend their perception of how trustworthy someone is. The literature demonstrates that in addition to one's links within a network, network structure is also relevant for learning and outcomes (Kohler et al., 2001). This is coherent with Granovetter's (1973) strength of weak ties hyphothesis, which argues that distant links matter for information transition and implies that sparse rather than dense networks should be most efficient for obtaining information (Kohler et al., 2001). To accommodate the relevance of such considerations, columns 4 to 7 of Table 13 control for the various centrality measures discussed in the data section: normalised indegree, normalised outdegree, betweenness and eigenvector centrality.

The various centrality measures control for different aspects of an individual's position and role within the network. Both indegree and outdegree have extremely small coefficients that are not significantly different from zero, either as the only control variables or when all controls are included. This would seem to suggest than an individual's immediate level of connection to a communication network does not directly impact the way in which he changes his offer. The same is true of betweenness, as shown in column 5 of Table 13. Eigenvector centrality has a small and insignificant coefficient when it is the only control and a smaller coefficient when all other controls are included (column 8). The inclusion of centrality controls only minimally affects the results on the news share variables, which are the variables of interest.

Column 8 shows the results of the regression when all the demographic, socioeconomic and centrality controls are included. The results, in terms of sign and significance, are similar to those in the previous specifications. The impact of the share of bad news at different offer levels is to reduce the probability of changing offer. This result is consistent with the social learning hypothesis (though it is not statistically significant). However, the share of bad news at the same offer level also reduces the probability of changing offer, a result that is significant but counter-intuitive in terms of the social learning hypothesis. The impact of good news share at either the same offer level is insignificant and negative, and the good news share at different offer levels insignificant and positive.

It is possible that unobserved characteristics at the level of the training group could be confounding social effects. For example, if one group experienced a group level financial shock, such as through a shared activity that cost them a significant amount of money, this could impact the change in trust game offers of the entire group, which could confound social effects as the network is highly correlated with programme group. Furthermore, there could be a 'social influence' type effect that occurs at the group level, whereby a norm or group average is approached in the behaviour of individuals. This would be a social effect but one that is distinct from the social learning process.

Attempts were made to control for group level fixed effects in the model to account for both of these possibilities. However, in the logit model in question the fixed effects model did not converge and hence results are not presented here. The inclusion of fixed effects would allow for consideration of change in offer over time within groups, such that any unobservable characteristics at the group level that impact change in offer will be partialled out.

The results of the first empirical model are loosely indicative of a social learning process, given the significant negative impact of the share of bad news at alternative offer levels on the probability of changing offer. It is possible that the counterintuitive results in the other variables of interest are the confounding impact of network structure, which is not fully captured in the centrality measures. The second basic empirical model considers a more precise empirical model of social learning, which takes into consideration the impact of network relationships in a more precise manner.

7.1.2 Basic Model 2

The second empirical model defines a social learning variable M. The variable is constructed to incorporate the direction of difference in j's round 1 offer relative to i's, for surprisingly successful events amongst the j's in i's communication network. We would therefore expect, as a very basic result, that those for whom M is positive would increase their offer when they play in round 3, whilst those for whom M is negative would also take into account the information available to them and decrease their offer. The social learning variable has a positive mean for those participants whose trust offer increased between rounds 1 and 3 and a negative mean for those whose offer decreased between rounds 1 and 3 (see Table 14). This is true for all the definitions of the social learning variables that are considered. At a basic level this is consistent with the existence of a process of social learning, since a negative sign on the social learning variable is associated with an average decrease in offer, and vice versa.

This idea is pursued further in regression analysis of the impact of the social learning variable on offer change. M, the social learning variable, is the variable of interest in a regression of the change in trust game offer between rounds 1 and 3 on M and a range of control variables. M is defined in the data section and (unless otherwise specified) the definition of M is as in M_1. Theoretically, if social learning does indeed take place as the model supposes, we would expect M to have a positive coefficient, reflecting individuals accommodating the information gathered from surprisingly successful offers made by others in their networks.

Table 15 presents the key results of this model. The value of the change in offer between trust game rounds for individual i is regressed on i's value of the social learning variable together with a varying set of control variables. Each set of control variables is used three times, with the definition of M changing each time. Demographic, socioeconomic and centrality controls are variously included. Three key measures of M are considered in this table, whilst the remaining M definitions are considered as a robustness check in Table 16. The three M definitions considered here are $M1$, the basic definition, $M6$, the definition weighted by strength of relationship and $M8$, the definition based on predicted likelihood of network relationships existing. The justification for the results based on the predicted network is included later, in a section dealing with network endogeneity.

In each specification, irrespective of the definition of M and regardless of which controls are included, the social learning variable M has a positive and highly significant coefficient. Ignoring for the moment potential endogeneity concerns, the positive and significant coefficient on the social learning variable is strongly indicative of a process of social learning, whereby individuals accommodate the good news experiences of peers in their network and adjust their behaviour in terms of trust game offer in the subsequent round accordingly.

The coefficient on $M6$ (the social learning variable weighted by closeness of relationship) is in all specifications larger than for the other M variables in the equivalent specifications. This is in part due to the smaller mean value of $M6$ (see Table 12), implying that a larger coefficient gives the same social learning impact. The smaller magnitude of $M6$ relative to $M1$ suggests that the good news events at offer levels greatly different from i's do not tend to occur within i's close network of friends. If round 1 offers were not markedly different amongst those in close network relationships (at least amongst those network links experiencing 'good news' events), this increases concerns about the potential endogeneity of the network. The large coefficient on $M6$ also implies that a large difference in round 1 offer, relative to i's offer, made by a close contact who experiences good news is very influential in i's social learning process, since it will impact i's $M6$ value substantially, and $M6$ has a large and significant coefficient implying a large impact on offer updating behaviour.

7.1.3 Robustness to varying the definition of M, the social learning variable

In this section I present robustness checks on the definition of the social learning variable M used in the basic model. The specification of regressions in Table 16 is as in columns 13 to 15 of Table 15 (with all controls included) but varying each time the definition of M for the remaining definitions of M not included in Table 15. Each M definition is explained fully in the data section.

The social learning result is robust to definitions of M that vary through the definition of the expected value of j's profits, incorporate a more restricted and closer learning network or weight the observations in M by either closeness of relationship or by the magnitude of surprise.

7.2 Extended Model: Controlling for potential network endogeneity

There is some possibility that the communication network is endogenous with trust game learning and outcomes. Bala and Goyal (2003) suggest a model of network formation where links are formed on the basis of how the individual perceives they will benefit from the link, thus the network may form endogenously around the intention to seek advice from one's contacts. This is particularly salient in the context of this network as at the point at which first round game play occurred the network was imminently about to form, since it was the first full day of the programme when participants would have been beginning to build network links. It is conceivable, therefore, that individuals sought out those who had done particularly well in the game to form a communication link with them, for example by discussing with them their experiences of the task.

Table 11, later used predicatively in the analysis, shows the results of a logit model that considers factors predicting whether a communication relationship exists between two individuals. The data are used in their dyadic form, whereby an observation exists for each prospective relationship within the network. The data are directed network data so there is one observation for individual i having a relationship with j and another observation for the reverse of that relationship. A dummy variable, equal to one if a communication relationship exists, is regressed on various relative characteristics of the pair, such as the difference between the two individuals' ages and incomes, and combinations such as 'Same Religion' or 'Both Indian', for categorical variables. See the note attached to Table 11 regarding standard errors.

This methodology for identifying predictors of network relationships is drawn from the literature on the formation of social networks. A varied literature uses logit regressions to determine predictors of dyadic relationships and make inferences about network formation. Such models are often used to inform more sophisticated theoretical models of network formation (Bramoulle et al., 2009). The methodology employed here touches on this literature at a very basic level, in the use of logit regressions to determine predictors of dyadic relationships in order to control for potential network endogeneity.

The network formation literature of this sort is based on a choice-theoretic model, which assumes that each individual's utility depends on his network links. A relationship is formed where the benefit of the relationship to the individual exceeds the cost to him of forming the relationship. The model relies on the assumption of the separability of the utility function. Separability implies that the utility derived from the network is equal to the sum of the utilities derived from all links, such that the structure of the network does not affect the utility associated with each link. A logit model estimated by maximum likelihood or a linear probability model can then be used to identify factors that predict whether a relationship exists.

There are three important caveats to this identification strategy. Firstly, there must be sufficient links for identification. If there are few actualised links the researcher should consider whether the entire population was indeed a potential link for everyone else, and natural partitions of the potential relationships, as subsets of the population, may need to be considered. This issue is circumvented in different ways by Mihaly (2007) and Jackson and Rogers (2007), though there is no definitive solution to estimating network formation in a large population (Bramoulle et al., 2009, p. 5). Secondly, there are limitations in the possibility of including individual characteristics in the analysis. Since each observation is a dyad relating to a potential relationship between two individuals, relative values are easily considered. However, there is no clear place in the analysis for individual characteristics. Thirdly, the structure of the error term must be considered. The standard errors will be correlated along at least two dimensions, since the same individual will feature multiple times as individual i and multiple times as individual j in the dyadic relationship combinations. This can be overcome by using individual fixed effects or by a variation on the robust covariance matrix to accommodate two-way clustering in dyadic relationships (Fafchamps and Gubert, 2007). This can be executed in Stata using the command *ngreg* when the adjacency matrix is square (Fafchamps and Gubert, 2007) or by manually computing two-way clustered standard errors otherwise (Thompson, 2011). Alternatively Quadratic Assignment Procedure (QAP), popular in the sociology literature, can be used to calculate p-values (Krackhardt, 1988), although Fafchamps and Gubert (2007) finds the robust covariance matrix to be more efficient, as it does not rely on bootstrapping. With these various caveats considered, the model can be estimated by maximum likelihood using logit regression procedures. It is important to bear in mind that the predictors of relationship formation in this context cannot generally be considered as causal.

In the literature that employs this methodologySantos and Barrett (2010) explore how developing country farmers construct information networks, considering the different types of information flows. In particular they consider the role of different types of characteristics in forming networks, with independent variables grouped into the categories 'identity' and 'interest'. They interpret the result that farmers seek information from different people for different decisions as suggestive that it is more than simply the individual's identity that guides decisions. Variables categorised as 'interest' variables also play a significant role in network formation.

Mayer and Puller (2008) investigate determinants of social network formation on university campuses. They model the social network based on exogenous school environment effects and the effects of endogenous choice arising from choice of friends based on their characteristics. Using Facebook data from ten US universities they measure segmentation of social ties by race, socioeconomic background and ability. They then develop a model of network formation. Preliminary to their network formation model Mayer and Puller (2008) estimate a linear probability model, regressing an indicator of whether two individuals are friends on the relative characteristics of the two individuals including demographic and socioeconomic factors. They find race, gender, relative college year, relative household income and college major (amongst other factors) to be significant. A low R-squared value likely reflects the many unobserved factors that determine friendship formation. The number of common friends is significant when included as an additional regressor in this specification. This is indicative that the network formation process is likely to be more complex than the linear probability model can fully capture. This leads the authors to develop a more complex model of network formation, which they then calibrate to the data.

De Weerdt (2002) investigates determinants of the formation of risk-sharing networks in rural Tanzania. The study questions the convention of treating a set group (for example a village) as representative of risk-sharing networks. De Weerdt considers as predictors of risk-sharing relationships factors relating to information and factors relating to trust, norms and the ability to punish, as well as measures of the heterogeneity of households with respect to the correlation of their income flows. Analysis is at the household level because there are too few links for identification at the individual level. As dependent variables the author considers firstly the number of links between two households; secondly, a 0/1/2 code for no link/unilateral/bilateral links; thirdly, an indicator variable for whether there is at least one link between the two households; and fourthly the geodesic distance between two households. De Weerdt (2002) finds that kinship, geographical proximity, number of common friends, clan membership, religious affiliation and wealth are strong determinants of risk-sharing networks. Fafchamps and Gubert (2007) similarly use logit regressions to explore the determinants of relationships that function as insurance links in a risk-sharing network in the rural Philippines.

Mihaly (2007) considers the relationship between popularity and academic outcomes. She models friendship nomination as a discrete decision predicted by individual and school demographics, then instruments for popularity using individual and school characteristics.

In the logit predictive regressions conducted in this dissertation, demographic characteristics, specific race combinations, socioeconomic indicators, education variables and training group variables are included in various combinations as predictors of the existence of a communication relationship. Table 11 present the results of the the logit predictive model. I report primarily on the results of column 6 of Table 11, with all controls included, as this is the specification used predicatively in the analysis. The average partial effects for the regressions in Table 11 are presented in Table 18.

Difference in age positively predicts the existence of a communication link, such that a greater age difference increases the probability that two individuals have ever communicated. This is contrary to homophily in networks, which suggests that links in social networks have a higher propensity to form between similar individuals. Additionally there is an increased likelihood of a link existing if two individuals are of the same home language, same religion or same gender, in accordance with the theory of homophily. Various combinations of race are significantly more likely to have communicated, relative to a black-black pair, however the race variables are difficult to interpret given the very small number of non-black individuals in the programme. Difference in income has a negative coefficient, implying that a greater difference in the income of two individuals is associated with a decreased likelihood of two individuals having ever communicated. The difference in self-perceived relative socio-economic status aged 15 (as captured by the 'ladder' question) is associated with a decreased probability of a communication link existing. The opposite is true for the difference in current self-perceived socio-economic status. The further apart two individuals are in terms of current self-perceived socio-economic status, the more likely they are to have communicated. This, along with difference in age being associated with increased likelihood of a relationship existing, is interesting because it counters the idea of homophily and may imply connections made strategically, where there is an advantage to having links with individuals different to oneself. The greater the difference in years of education, the less likely two individuals are to have communicated. Being in the same programme group or the same programme region both strongly predict the likelihood of a link existing.

These results cannot be interpreted causally, since there are many potentially omitted variables and confounding factors. However, the concern of reverse causality (where individuals share similar characteristics as a result of a network link) is reduced due to the fact that baseline data on characteristics were gathered prior to the formation of the network.

Returning to the issue of possible network endogeneity of the social learning result, predictive values from the logit specification of communication links are used to form a predicted network. The predicted network, which is arguably exogenous, is then used to construct the $M8$ definition of the social learning variable. The intention here is that by capturing the impact of learning from those that the individual was exogenously likely to have communicated, rather than their actual network (which may be endogenous with outcome behaviours) the $M8$ definition of the social learning variable will not be spuriously correlated with behavioural updating due to the endogeneity of network structure with outcomes or other relevant determining factors.

In general it is important to exercise caution that the explanatory variables used in the predictive model are not themselves endogenous with the network. Certain variables used predicatively could be influenced by the network itself. However, this is a lesser concern here as the baseline survey, from which the characteristics in the predictive regression were drawn, was conducted before the network was formed. Therefore, a variable that might otherwise be endogenous with the network, such as wealth, which may increase as a result of social capital or lucrative business and employment links in the network, is unlikely to be endogenous with the network in this case.

The revised version of the social learning variable, $M8$, is calculated as the weighted average of M across all individuals, weighted by the predicted likelihood of a link existing. The social learning result remains robust to the use of the predicted network, suggesting that the social learning result is not a product of network endogeneity. These results are displayed in columns 3, 6, 9, 12 and 15 of Table 15.

Additionally, I consider the validity of the network predictors used in determining the predicted network by running logit regressions of alternative network relationships. Columns 1 to 4 of Table 17 report regressions predicting the existence of the following relationship types: participants the individual has communicated with about the programme outside of the programme sessions, participants who have approached the individual about working together on a community project, participants who have approached the individual seeking assistance on a community project and participants whom the individual has asked to lend him money. The average partial effects for the regressions in Table 17 are presented in Table 19. The results of the predictive models of alternative network relationships can be compared with the results of the predicted communication relationship in column 6 of Table 11 to see whether the factors predicting communication relationships are in general predictive of various levels of social interaction. A very similar set of factors significantly predict relationships where two individuals have talked about the programme outside of programme activities, though fewer variables are significant in this regression. Approaching to work on a project or seeking assistance on a project are also very similarly determined in the predictive models. Same home language and same programme group strongly predict the existence of these relationship types. A relationship in which the individual has asked to borrow money is also significantly predicted by home language and presence in the same programme group. However, due to the very few dyads for which the value of the dependent variable is equal to 1 for the borrowing relationship this result should in any case be treated with caution. Overall, although fewer factors significantly predict the relationships of the types in Table 17, there is nothing to strongly contradict the coefficients on predictors used to form the predicted network for $M8$. The factors identified as predicting a communication relationship appear plausible as predictors of relationships more generally, including other types of network link that may be relevant for learning. This adds weight to the validity of the $M8$ social learning variable that controls for potential network endogeneity.

The persistent significance of the social learning variable when it is defined to accommodate a predicted network that is plausibly exogenous strengthens the social learning result by suggesting that the results of Table 15 are not spuriously generated by the endogeneity of network structure with behaviour and outcomes.

8 Interpretation of the Social Learning Results

The result that individuals appear to learn from the experiences of their peers regarding the trustworthiness of an anonymous partner in the game, and update their behaviour accordingly, is significant on multiple levels. The social learning result furthers evidence in the social effects literature of the existence of social learning processes, which has specific bearing for policy that aims to have a social multiplier effect through influencing preferences of a subgroup. The result also speaks more directly to the literature on trust and brings forth positive possibilities given the relevance of trust in the economics literature more broadly.

That individuals update their priors about the trustworthiness of strangers by learning from their network contacts is an interesting result in its own right. As discussed in the literature review, beliefs, for example in models of statistical discrimination in the labour market, are often assumed either to be static or to involve no clear evolutionary process empirically. The key findings here demonstrate the updating of beliefs in terms of trustworthiness according to the observed experiences of peers in a social network.

The result also contributes to the literature on social effects by further demonstrating the relevance of social interactions (rather than solely individual and household characteristics) for outcomes. By identifying the significant impact of the variation in exposure to good and bad news at various offer levels on change in offers, I demonstrate the relevance of a specific type of social interaction for outcomes, in line with Moffitt (2001). In the sociology literature, social networks play an important role in building trust (see, for example, Granovetter (1985) and Coleman (1988)) and are also widely used for trust-intensive exchanges (Karlan et al., 2009, p. 1312). This research contributes in demonstrating a social learning process where an updating process occurs in the perception of trustworthiness of others, and behaviour is adjusted accordingly. A general implication of results of social effects for policy is the potential to implement policies with intended 'social multiplier' effects, which can increase the efficiency of policy impact. If policy is targeted at a subset of the population, with the intention of influencing preferences, priors and ultimately outcomes of others in the population, then evidence of social effects in this arena is highly relevant. The social learning result supports the possibility of potential policy that operates through social effects to have an impact on outcomes beyond the target group through social interactions.

Finally, the finding that the perceived trustworthiness of others can be learnt has implications more broadly, given the relevance of trust for a host of economic outcomes. Trust and pro-social behaviour are important economic considerations for various reasons. Trust is key to many of the informal mechanisms that drive economic development. It has been linked to investment (Bottazzi et al., 2011), economic growth (Zak and Knack, 2001), government performance and lack of corruption (Porta et al., 1996), and international trade and flows of finance (Sapienza et al., 2006). At a micro level trust is important where there are problems of asymmetric information, such as moral hazard, for example for repayment in group lending situations (though personal trust rather than general trust seems to the driving factor (Cassar et al., 2007)) and additionally for informal insurance (Chandrasekhar et al., 2011) and job search (Karlan et al., 2009).

Defining social capital as interpersonal networks and the factors required for individuals to make credible promises within these, Dasgupta et al. (2009) posit trust as a key feature of social capital. They extend this to incorporate the mutual enforcement of agreements. Under this definition 'mutual trust is the basis of cooperation' (Dasgupta et al., 2009, p. 5). Social capital, intrinsically linked to trust in this definition, has been linked in the literature to a variety of economic outcomes. A concrete example of this mechanism is the framing of transactions with 'social collateral', where an individual complies with an agreement to avoid losing a valuable relationship if they renege on the agreement, for example the repayment of a loan. Karlan et al. (2009) extend this empirically to real-world social networks, where detailed network data allow for consideration of 'trust intermediaries' who enable the upholding of agreements between two individuals who are not directly connected, through the virtue of the social collateral at stake in their social connections. Additionally, trust amongst strangers has been shown to facilitate collective action, which could be relevant for greater efficiency in the provision of public goods (Putnam et al., 1994) (Coleman, 1988).

The result of the malleability of perceived trustworthiness and willingness to trust presents possibilities for influencing economic development through taking advantage of social learning processes in beliefs that underlie economic processes. If a subset of individuals in a network, ideally opinion leaders, experience trustworthy agents (borrowers, government, job application referees etc) then there is the possibility for a social learning process to generate increased levels of trust more widely within the network. Increased levels of trust have the possibility to positively influence economic development in the ways discussed above. An important caveat to the above conclusions is that the evidence from which individuals learn must promote trust. The network must face positive experiences of trustworthiness in order for this outcome to hold. If individuals learn from their peers the apparent lack of trustworthiness of others then the results of this may be the converse. The learning of trust is therefore a result to be engineered in policy in appropriate contexts, rather than a solution, or necessarily even a positive finding from a policy perspective, in its own right.

The results contribute to the literature through the empirical consideration of social effects in a complete social network, in such a way as to identify a social learning process. Additionally this research is unusual in its consideration of preferences in the context of the consideration of social effects, through the use of preference data elicited through experimental games. However, there are a great many dimensions to networks and preferences beyond the scope of this dissertation, for example the specific role of network structural characteristics in the evolution of preferences.

As discussed in the empirical challenges section, there are significant challenges to controlling for all possible confounding social effects of other types, such as social influence and mimicry. Therefore it is possible that the effect observed, at least in part, may be due to other social effects. This is a necessary consideration with regards to policy implications that endeavour to exploit social effects.

An important caveat when considering the validity of these results is the possibility of mean reversion. There is a possibility of mean reversion resulting in the spurious positive correlation of the M variable with change in offer, because i's round 1 offer appears in both the dependent variable and the social learning variable (see Conley and Udry (2010, p. 66) for an analogous example). Similarly the censoring of offers at R0 and R50 could also spuriously lead to the appearance of social learning. However, due to limited sample size, this is a potential confound that cannot be easily tested for in my data.

9 Conclusions

In the context of a real world social network in a youth leadership programme in South Africa, I demonstrate empirical evidence of a social learning process in the perceptions of the trustworthiness of others. This result is found in game play of an experimental trust game across two rounds. By identifying the significant impact of the share of bad news at alternative offer levels, reducing the probability that an individual will change their offer in the second round of play, I allude to the possibility of social learning process. Subsequently, I define a social learning variable that accommodates the relative offers made by those who experienced surprisingly positive outcomes in round 1 trust game play, within the individual network. Using this I show a positive and significant relationship between the information available for social learning, captured in the social learning variable, and the change in participants' offers in the final round of game play. The social learning result is robust to multiple alternative definitions of the social learning variable, as well as to the use of a predicted network measure which controls for potential network endogeneity. Though some endogeneity and identification concerns persist, such as concerns regarding censoring of offer values and the possibility of mean reversion, the result is suggestive of a social learning process. The social learning result contributes to the literature that attempts to identify social effects empirically.

Table 1: Descriptive Statistics of Participants

	Western Cape		Kwazulu Natal		Gauteng		Total	
	No.	%	No.	%	No.	%	No.	%
Race								
African	38	82.6%	47	92.2%	56	96.6%	141	91.0%
Coloured	6	13.0%	3	5.9%	0	0.0%	9	5.8%
Indian	0	0.0%	0	0.0%	1	1.7%	1	0.6%
White	1	2.2%	0	0.0%	0	0.0%	1	0.6%
Other	1	2.2%	1	2.0%	1	1.7%	3	1.9%
Gender (1=Female)								
Males	20	44.4%	32	62.7%	32	55.2%	84	54.5%
Females	25	55.6%	19	37.3%	26	44.8%	70	45.5%
Home Language								
IsiNdebele	0	0.0%	0	0.0%	2	3.4%	2	1.3%
IsiXhosa	27	58.7%	14	27.5%	0	0.0%	41	26.3%
IsiZulu	9	19.6%	17	33.3%	20	33.9%	46	29.5%
Sepedi	0	0.0%	0	0.0%	12	20.3%	12	7.7%
Sesotho	1	2.2%	8	15.7%	10	16.9%	19	12.2%
Setswana	0	0.0%	7	13.7%	6	10.2%	13	8.3%
SiSwati	0	0.0%	0	0.0%	1	1.7%	1	0.6%
Tshivenda	0	0.0%	0	0.0%	1	1.7%	1	0.6%
IsiTsonga	0	0.0%	0	0.0%	3	5.1%	3	1.9%
Afrikaans	4	8.7%	4	7.8%	0	0.0%	8	5.1%
English	5	10.9%	1	2.0%	3	5.1%	9	5.8%
Other (specify)	0	0.0%	0	0.0%	1	1.7%	1	0.6%
Employed								
Unemployed	7	15.2%	12	23.1%	8	13.8%	27	17.3%
Employed	39	84.8%	40	76.9%	50	86.2%	129	82.7%
Doing one year ago								
Don't know	0	0.0%	0	0.0%	1	1.7%	1	0.6%
Full time scholar or student	12	26.1%	14	26.9%	13	22.0%	39	24.8%
Homemakers	0	0.0%	1	1.9%	0	0.0%	1	0.6%
Other	0	0.0%	1	1.9%	3	5.1%	4	2.5%
Self-Employed	0	0.0%	2	3.8%	2	3.4%	4	2.5%
Unemployed and actively searching for a job	4	8.7%	6	11.5%	10	16.9%	20	12.7%
Unemployed but not searching for a job	0	0.0%	0	0.0%	1	1.7%	1	0.6%
Volunteer	2	4.3%	5	9.6%	8	13.6%	15	9.6%
Working for Pay	28	60.9%	23	44.2%	21	35.6%	72	45.9%
Facebook Account								
No facebook account	2	4.4%	8	15.7%	2	3.6%	12	7.9%
Facebook account	43	95.6%	43	84.3%	54	96.4%	140	92.1%
Other social media								
Do not use other social media	11	24.4%	11	21.6%	15	26.8%	37	24.3%
Use other social media	34	75.6%	40	78.4%	41	73.2%	115	75.7%

1. Descriptive statistics on binary variable characteristics of participants. The sample here is limited to those for whom trust game data were available across both rounds. This reflects the broad sample used in the analysis. Variables are discussed further in the Baseline Data section.
2. Columns are reflective of the three provincial groups, with the final column presenting statistics on the entire set of participants across all groups.
3. Employment is a broad definition, including self-employment, casual work or help in a family business.

Table 2: Descriptive Statistics of the Participants

Variable	Mean	Std. Dev.	Min.	Max.	N
Age (years)	24.671	3.39	19	40	155
Gender (1=Female)	0.455	0.5	0	1	154
Years of Education	12.727	1.411	9	18	154
Income (Monthly)	2975.519	4678.127	0	33300	154
Income (Monthly) IST	6.124	3.843	0	11.106	154
Reservation Wage	3139.849	3043.312	0	18000	152
Mother's years of education	10.173	3.443	1	16	139
Socioeconomic status (ladder) now	2.904	1.036	1	6	157
Socioeconomic status (ladder) aged 15	2.414	1.068	1	6	157
Expect wallet returned by stranger	0.223	0.418	0	1	148

1. Statistics calculated only the estimation sample, which includes only those for whom both round 1 and round 3 trust game data were available.
2. Montly income includes gross income from all earned sources, including regular employment, self-employment and casual work. Income (Monthly) IST is the Inverse Sine Transformation of income, which transforms the skewed income distribution whilst avoiding the problem that a standard log transformtion faces with zero values for income. The IST transformed income is the variable used in the regression analysis.
3. Monetary values (income measures and reservation wage) are given in South African Rand.
4. Socioeconomic status questions relate to self-perceived position on a 6 step socio-economic ladder, where 1 is the poorest people and 6 the richest.
5. 'Expect wallet returned by stranger' is equal to 1 if the individual considered it likel or very likely that a lost wallet containing R200 would be returned with the money in it.

Table 3: Summary Statistics of Trust Game

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
A's offer in round 1	19.36	12.77	0	50	158
A's offer in round 3	18.25	13.15	0	50	158
Change in A's offer between rounds 1 and 3	-1.11	14.48	-45	45	158
A's profit in the trust game round 1	-0.85	14.23	-30	60	158
A's profit in the trust game round 3	0.13	16.66	-50	70	158
How much A expected B to return to them, Round 1	29.69	26.56	0	160	157
How much A expected B to return to them, Round 3	27.11	22.44	0	120	158
Profit A expected, Round 1	10.4	21.62	-49	130	157
Profit A expected, Round 3	8.85	18.7	-30	80	158

1. Statistics on trust game variables are all given in South African Rand.
2. Player A's profit is calculated as the amount he chose to send to player B plus the amount he received back from player B.

Table 4: Direction of change of trust game offers between rounds 1 and 3

Direction of Change in Offer	No.	%
No change	41	25.9%
Increase	51	32.3%
decrease	66	41.8%

1. This table shows summary statistics of the change in trust game offer between rounds. The first column shows the number of individuals who kept the same offer in both rounds, increased their offer in round 3 and decreased their offer in round 3, respectively. The second column shows the percentage of the group who behaved in each of these ways. These data are presented only for the estimation sample, those for whom trust game data were available from both rounds 1 and 3.

Table 5: Summary statistics on sample and those not in sample

	(1) In sample mean/sd	(2) Attrition mean/sd
Age (years)	24.67	24.61
	3.39	3.01
Years of Education	12.73	12.83
	1.41	1.83
Income (Monthly) IST	6.12	7.59
	3.84	3.30
A's offer in round 1	19.36	21.05
	12.77	12.69
How much A expected	29.69	29.10
B to return to the 1	26.56	22.54
Employed	0.83	0.92
	0.38	0.28
Observations	158	62

1. Mean and standard deviation (below) presented.
2. Column 1 shows summary statistics for those in the estimation sample (with trust game data from both rounds) and column 2 for those who are not in the sample (primarily those who did not return to the group in round 3).
3. Income (Monthly) IST is earned income from all earned sources, transformed by the Inverse Sine Transformation to transform the skewed distribution of income whilst dealing successfully with zero income values.
4. 'How much A expected B to return is the amount, in Rand, that player A wrote that they expected to receive back from player B in the trust game in round 1.
5. Employed is a broad definition, including regular employment, self-employment and casual work.

Table 6: Network Densities: Communication (directed)

Group	Western Cape	n	Gauteng	n	Kwazulu Natal	n
All three groups	0.319	59	0.290	68	0.336	60
Group 1	0.948	18	0.800	15	0.927	23
Group 2	0.924	15	0.908	20	0.977	12
Group 3	0.795	26	0.584	33	0.777	25

1. Density is the proportion of potential links that could exist within a network which actually exist.
2. Densities shown here are for the directed network, where a link in each direction is possible between two individuals.

3. Densities are given at both the group level and the province level.

Table 7: Network Densities: Communication (reciprocal)

Group	Western Cape	n	Gauteng	n	Kwazulu Natal	n
All three groups	0.268	59	0.272	68	0.216	60
Group 1	0.895	18	0.862	15	0.676	23
Group 2	0.867	15	0.855	20	0.832	12
Group 3	0.652	26	0.603	33	0.407	25

1. Density is the proportion of potential links that could exist within a network which actually exist.
2. Densities shown here are for the reciprocal network, where a link between two individuals only exists if both indivduals nominate each other, and therefore have a reciprocal relationship. There is no possibility for a link in one direction that is not reciprocated.
3. Densities are given at both the group level and the province level.

Table 8: Summary Statistics of Network Centrality Measures

Variable	Mean	Std. Dev.	Min.	Max.	N
Degree	16.362	5.124	2	27	152
Outdegree	19.546	6.096	1	41	152
Indegree	19.322	5.316	3	33	152
Normalised Indegree	31.48	8.82	5.172	50.847	152
Normalised Outdegree	31.828	9.966	1.695	67.241	152
Betweenness	35.195	61.13	0	322.991	152
Eigenvector Centrality	0.095	0.088	0.001	0.256	152

1. Degree is the number of direct network links an individual has.
2. Outdegree is the number of outward links (i.e. nominated communication links) an individual has.
3. Indegree is the number of inward links an individual has, that is, how many people nominated them as someone with whom they have a communication link.
4. Normalised indegree and normalised outdegree are measures that are scaled to be comparable regardless of network size. Normalised degree is calculated as degree multiplied by $\frac{100}{n-1}$.
5. Betweenness captures the extent to which an individual lies on the shortest paths between other individuals, and thus reflects the centrality of the individual in terms of network fragmentation if they were to leave the network.
6. Eigenvector centrality captures the extent to which the individual is connected to other well-connected individuals.
7. See Section 4.5 for further definitions and discussion of these measures.

Table 9: Networks Statistics from the Literature

Network	**Type**	n	m	z	l	α	C_1	C_2	r
Film actors	undirected	449913	25516482	113.43	3.48	2.3	0.20	0.78	0.208
Company directors	undirected	7673	55392	14.44	4.60	-	0.59	0.88	0.276
Math coauthorship	undirected	253339	496489	3.92	7.57	-	0.15	0.34	0.120
Physics coauthorship	undirected	52909	245300	9.27	6.19	-	0.45	0.56	0.363
Biology coauthorship	undirected	1520251	11803064	15.53	4.92	-	0.088	0.60	0.127
Telephone call graph	undirected	47000000	80000000	3.16		2.1			
Email messages	directed	59912	86300	1.44	4.95	1.5/2		0.16	
Email address books	directed	16881	57029	3.38	5.22	-	0.17	0.13	0.092
Student relationships	undirected	573	477	1.66	16.01	-	0.005	0.001	-0.029
Sexual contacts	undirected	2810				3.2			
Western Cape	directed	59	1091	18.49	1.9		0.709	0.731	
KwaZulu Natal	directed	60	1189	19.82	1.9		0.561	0.609	
Gauteng	directed	68	1319	19.40	1.9		0.672	0.689	

1. Rows 1 to 10 replicate data reported by Newman (2003, p.182) of statistics on social network in the literature.The programme networks discussed here are listed in the last three columns for comparison to social netowrks in the literature.
2. The statistics presented in this table are: total number of vertices n, which is the number of people in the network; total number of edges m, the number of relationship links; mean degree z, the average number of connections an individual has; mean vertex-vertex distance l, the average distance between individuals in the network; exponent α of the degree distribution if the distribution follows a power law (or "-" if not. In/ out-degree exponents are given for directed graphs); clustering coefficient $C^{(1)}$, from equation 18; clustering coefficient $C^{(2)}$ from equation 19, which both reflect the extent to which the network bunches into smaller groups; and degree correlation coefficient r.

$$C^{(1)} = \frac{3 \times number\ of\ triangles\ in\ the\ network}{number\ of\ connected\ triples\ of\ vertices} \tag{18}$$

$$C^{(2)} = \frac{1}{n} \sum_i C_i \tag{19}$$

Table 10: Summary statistics of the 'good news' variables for the logit specification

Variable	Mean	Std. Dev.	Min.	Max.	N
Dummy equal to 1 if offer changed between rounds	0.743	0.438	0	1	152
Share of good news events at same offer	0.058	0.066	0	0.5	152
Share of bad news events at same offer	0.076	0.08	0	0.571	152
Share of good news events at different offer	0.269	0.118	0	0.538	152
Share of bad news events at different offer	0.314	0.111	0	0.556	152

1. The 'dummy equal to 1 if offer changed between rounds' variable is the dependent variable in the first empirical model.
2. The 'share' variables are the variables of interest on the right hand side of the first empirical model, and capture the proportion of events in an individual's network that were good news at the same offer level as the individual, bad news at the sam offer level as the individual, and so forth.

ss

Table 11: Logit model: Factors predicting whether two participants have ever communicated

	(1) Demographic	(2) Race	(3) Socioeconomic	(4) Education	(5) Activate Training	(6) All
Ever Communicated						
Difference in Age (Years)	0.004					0.033***
	(0.00)					(0.01)
Same Home Language	0.902***					0.931***
	(0.04)					(0.09)
Same Religion	0.205***					0.330***
	(0.04)					(0.08)
Same Gender	0.109**					0.159*
	(0.03)					(0.08)
Same Race	-0.099*					
	(0.04)					
Both Coloured		0.863***				-0.021
		(0.17)				(0.18)
Both Indian		2.096*				8.716***
		(0.82)				(1.37)
Both White		2.319***				0.528
		(0.67)				(0.42)
Black-Coloured		-0.180***				0.259*
		(0.05)				(0.11)
Black-Indian		0.087				0.737**
		(0.11)				(0.28)
Black-White		-0.268*				0.007
		(0.11)				(0.21)
Coloured-Indian		-0.256				1.866*
		(0.43)				(0.73)
Coloured-White		0.997***				-0.267
		(0.25)				(0.29)
Indian-White		-0.543				
		(1.04)				
Both Employed			0.090*			0.061
			(0.04)			(0.09)
Both Unemployed			0.404***			0.330
			(0.11)			(0.27)
Difference in Income (Monthly)			-0.002			-0.022**
			(0.00)			(0.01)
Difference in Reservation Wage			0.000			0.000
			(0.00)			(0.00)
Ladder difference age 15			-0.027*			-0.073*
			(0.01)			(0.03)
Ladder difference now			0.035**			0.089**
			(0.01)			(0.03)
Dif in years of education				-0.015		-0.056**
				(0.01)		(0.02)
Dif in Mother's years of educ				-0.007		-0.009
				(0.00)		(0.01)
Same Group					4.071***	4.292***
					(0.06)	(0.09)
Same Training Province					6.360***	6.532***
					(0.14)	(0.20)
o.Indian-White						0.000
						(.)
Constant	-2.372***	-2.096***	-2.174***	-2.099***	-9.111***	-9.866***
	(0.04)	(0.02)	(0.04)	(0.02)	(0.16)	(0.24)
Observations	36669	39194	34261	30190	35530	21014

1. Robust standard errors in parentheses.
2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
3. The dependent variable in this table is a dummy variable equal to one if an individual identifies as have a communication relationship with another indvidial and equal to zero otherwise.
4. The Indian-White race combination was omitted by Stata from the specification in column 6 because it predicts failure perfectly.
5. Note: It would be preferable to use two-way clustered standard errors (as discssed in the section reviewing network formation papers). However, the use of the *ngreg* command in Stata would require a reduced sample size for the data to have the necessary symmetric dyadic data structure. As the logit results here are primarily used for prediction the standard errors are not of great relevance, so I considered it better to use the robust standard errors and maintain the larger sample size.

Table 12: Number of individuals who influence each person in the various definitions of M

Variable	Mean	Std. Dev.	Min.	Max.	N
M1	0.518	4.983	-16.25	12.813	152
M2	0.51	4.781	-15	12.813	152
M4	0.49	4.886	-16.25	12.813	152
M3	0.011	4.901	-16.25	12.045	152
M5	0.645	5.79	-25	19.286	152
M6	0.043	0.28	-0.684	1.5	152
M7	0.231	1.254	-3.677	4.129	152
M8	1.239	13.004	-30.171	20.874	152
Number of good news events in network(using M1)	6.862	3.373	1	16	152
Number of good news events in network(using M2)	6.507	3.095	1	15	152
Number of good news events in network(using M4)	6.664	3.257	1	16	152
Number of good news events in network(using M3)	7.184	3.25	1	16	152
Number of good news events in network (using M5)	4.382	2.694	0	14	152
Number of good news events in network (using M6)	6.862	3.373	1	16	152
Number of good news events in network (using M7)	6.862	3.373	1	16	152

1. This table presents summary statistics on the social learning variable M.

2. Summary statistics are also given for the number of good news events that an individual observes within their social network when each definition of M is used. This captures the number of peers whose experience the individual will incorporate in the learning process.

ss

3. $\mathbf{M_1}$ is the defined relative to the 'expected profits' of j based on the median profits of all individuals within i's training group whose offer in the first round was within a range of R5 above of below j's offer. M_1 is then calculated as the mean of all M values of individuals in the network who experienced good news events. M_1 is the standard definition of M used throughout, unless mentioned otherwise.

   $\mathbf{M_2}$ is defined as M_1 but with a greater range used in calculating $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$, with all offers within R10 above of below j's offer and within his training network included.

   $\mathbf{M_3}$ is as M_1 and M_2 but with a R2 range above and below j's offer.

   $\mathbf{M_4}$ classifies all offers as either low (R0 to R15), medium (greater than R15 up to and including R30) or high (greater than R30 up to and including R50). $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ is then calculated as the median of all offers within j's training group whose round 1 offer was within the same bracket as j's offer.

   $\mathbf{M_5}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ as used in M_1 but uses a different definition of i's information network. Where the previous definitions of M assumed that i could learn from anyone with whom he nominated as ever having communicated with, this definition restricts i's observed events to those within his network of contacts with whom he said he discussed the programme outside of training activities.

   $\mathbf{M_6}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ and the same information network as used in M_1 but weights the good news observations differently. Whereas M_1 weights each observation of a good news event equally, M_6 weights each good news observation by the strength of relationship (between 1 and 10) that i said that he has with the individual j that he observes.

   $\mathbf{M_7}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ and the same information network as used in M_1 but weights each good news observation according to the magnitude of the good news. Each good news observation is weighted by $\pi_{j,2}(x_{j,1}) - \hat{E}_{i,1}(\pi_{j,2}(x_{j,1})$, such that the greater a positive surprise the news was, the greater it weighs in i's social learning.

   $\mathbf{M_8}$ uses the same definition of $\hat{E}_{i,1}(\pi_{j,2}(x_{j,1}))$ as M_1 but weights all observed good news outcomes (for every participant, not just those in i's actual network) by the predicted values from the network predictive logit model shown in column 6 of Table 11.

ss

Table 13: Logit model predicting whether a change in trust game offer occurs between rounds 1 and 3.

	(1) No controls	(2) Demographic	(3) Socioeconomic	(4) Centrality1	(5) Centrality2	(6) Centrality3	(7) Centrality4	(8) All controls
Dummy equal to 1 i e								
Share of good news	-7.112*	-9.181*	-6.628	-6.947*	-7.071*	-6.797*	-6.792	-4.262
events at same offer	(-2.13)	(-2.52)	(-1.46)	(-2.03)	(-2.11)	(-1.99)	(-1.94)	(-0.67)
Share of bad news	1.460	-2.246	-0.154	1.502	1.487	1.823	1.741	-4.413
events at same offer	(0.52)	(-0.67)	(-0.05)	(0.52)	(0.53)	(0.63)	(0.58)	(-0.88)
Share of good news	1.453	0.699	2.256	1.238	1.442	1.438	1.286	3.416
events at differen r	(0.86)	(0.37)	(1.03)	(0.71)	(0.85)	(0.85)	(0.73)	(1.11)
Share of bad news	-1.701	-3.698	-3.430	-1.837	-1.703	-1.456	-1.634	-4.652
events at differen r	(-0.84)	(-1.55)	(-1.42)	(-0.90)	(-0.84)	(-0.70)	(-0.75)	(-1.51)
Age (years)		0.132						0.143
		(1.74)						(1.50)
Gender (1=Female)		-0.493						-0.876
		(-1.11)						(-1.56)
Coloured		1.599						2.279
		(0.99)						(1.30)
o.Indian		0.000						0.000
		(.)						(.)
o.White		0.000						0.000
		(.)						(.)
isiNdebele		-1.187						-1.216
		(-0.72)						(-0.67)
isiZulu		-0.020						0.205
		(-0.04)						(0.31)
Sepedi		0.697						
		(0.61)						
Sesotho		-0.393						-0.492
		(-0.56)						(-0.57)
Setswana		-0.519						-0.559
		(-0.66)						(-0.61)
o.SiSwati		0.000						0.000
		(.)						(.)
o.Tshivenda		0.000						0.000
		(.)						(.)
isiTsonga		-0.306						
		(-0.23)						
Afrikaans		0.222						0.283
		(0.13)						(0.16)
English		-0.190						-0.010
		(-0.18)						(-0.01)
Years of Education			0.126					0.109
			(0.82)					(0.58)
Mother's years of			-0.004					0.006
education			(-0.07)					(0.07)
Income (Monthly) IST			0.070					0.117
			(1.13)					(1.51)
Expect wallet			0.288					0.214
returned by stranger			(0.47)					(0.28)
Share of network who			1.593					3.816
expect wallet retu d			(0.77)					(1.42)
Normalised Indegree				0.005			0.003	0.011
				(0.20)			(0.10)	(0.25)
Normalised Outdegree				0.008			0.007	-0.032
				(0.34)			(0.26)	(-0.85)
Betweenness					0.000		-0.001	-0.005
					(0.12)		(-0.14)	(-0.95)
Eigenvector						1.348	0.972	5.481
Centrality						(0.58)	(0.35)	(1.22)
o.Sepedi								0.000
								(.)
o.isiTsonga								0.000
								(.)
Constant	1.562	-0.068	-0.370	1.241	1.548	1.319	1.164	-3.462
	(1.48)	(-0.03)	(-0.16)	(1.00)	(1.45)	(1.15)	(0.91)	(-0.94)
Observations	152	142	124	152	152	152	152	109

1. T-statistics (based on robust standard errors) in parentheses.
2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
3. DV:Dummy variable equal to 1 if the individual changed their offer between rounds and 0 if the same offer was made in both rounds.
4. The results estimate Basic Model 1.

Table 14: Summary statistics of M by direction of change in A offer between rounds 1 and 3)

	(1) No change *mean/sd*	(2) Increase *mean/sd*	(3) Decrease *mean/sd*
M1	−0.65	3.01	−0.74
	5.06	4.01	4.96
M2	−0.65	2.91	−0.67
	4.82	3.94	4.71
M4	−0.67	2.95	−0.74
	5.00	3.98	4.80
M3	−0.99	2.61	−1.43
	5.30	3.66	4.75
M5	−0.47	3.30	−0.77
	5.85	4.28	6.13
M6	−0.02	0.19	−0.03
	0.24	0.32	0.23
M7	0.08	0.75	−0.08
	1.25	1.01	1.32
M8	−0.88	9.01	−3.62
	13.80	9.66	12.08
Observations	39	50	63

1. This table shows summary statistics on each of the definitions of the social learning variable M grouped by whether the individual increased their offer, decreased their offer or made the same offer in both rounds. The results here demonstrate that the correlation of change in offer with the M variable in the manner predicted by the model.
2. Mean values with standard deviations given below.

Table 15: Social learning regressions for change in trust game offer. DV: Change in A offer in trust game between rounds 1 and 3

	(1) A	(2) A	(3) A	(4) B	(5) B	(6) B	(7) C	(8) C	(9) C	(10) D	(11) D	(12) D	(13) E	(14) E	(15) E
M1	1.320***			1.425***			1.118***			1.263***			1.098***		
	(6.19)			(5.73)			(4.79)			(5.85)			(3.87)		
M6		22.074***			25.845***			17.723***			21.953***			20.451***	
		(5.75)			(5.45)			(4.27)			(5.59)			(3.63)	
M8			0.616***			0.699***			0.536***			0.598***			0.544***
			(8.07)			(7.95)			(6.42)			(7.58)			(5.43)
Age (years)				0.265	0.294	0.253							0.425	0.379	0.390
				(0.75)	(0.82)	(0.78)							(1.09)	(0.97)	(1.07)
Gender (1=Female)				-2.104	-0.410	-3.335							-1.436	0.218	-1.810
				(-0.84)	(-0.17)	(-1.46)							(-0.52)	(0.08)	(-0.70)
Coloured				5.366	5.234	5.063							4.325	4.394	4.335
				(0.59)	(0.57)	(0.61)							(0.49)	(0.49)	(0.52)
Indian				-7.691	-6.253	-1.882							-0.854	0.140	1.678
				(-0.52)	(-0.42)	(-0.14)							(-0.06)	(0.01)	(0.12)
White				11.591	14.093	11.485							5.067	7.181	7.285
				(0.78)	(0.94)	(0.84)							(0.35)	(0.49)	(0.53)
isiNdebele				-4.699	-0.302	-7.032							-8.102	-3.643	-8.732
				(-0.47)	(-0.03)	(-0.77)							(-0.80)	(-0.36)	(-0.91)
isiZulu				1.966	1.029	4.041							3.495	2.601	4.333
				(0.65)	(0.34)	(1.44)							(1.08)	(0.80)	(1.42)
Sepedi				-1.974	-1.944	-2.277							-6.199	-5.875	-6.881
				(-0.41)	(-0.40)	(-0.52)							(-1.08)	(-1.01)	(-1.28)
Sesotho				4.753	3.281	5.388							-1.140	-2.411	0.097
				(1.23)	(0.85)	(1.52)							(-0.28)	(-0.58)	(0.03)
Setswana				2.825	1.952	2.369							0.415	-0.248	1.206
				(0.61)	(0.42)	(0.56)							(0.09)	(-0.05)	(0.27)
SiSwati				6.124	4.708	3.638									
				(0.44)	(0.33)	(0.28)									
Tshivenda				4.351	3.036	8.882							-6.818	-4.642	0.485
				(0.31)	(0.22)	(0.70)							(-0.47)	(-0.31)	(0.04)
isiTsonga				4.503	4.892	5.838							-1.924	-2.402	3.055
				(0.55)	(0.60)	(0.78)							(-0.14)	(-0.17)	(0.24)
Afrikaans				-2.610	-3.712	-2.321							-4.585	-5.789	-4.294
				(-0.27)	(-0.39)	(-0.27)							(-0.48)	(-0.61)	(-0.48)
English				-4.473	-6.618	-3.384							-4.470	-6.214	-3.834
				(-0.71)	(-1.05)	(-0.59)							(-0.67)	(-0.93)	(-0.61)
Years of Education							0.835	1.029	0.726				0.956	0.987	0.956
							(1.00)	(1.22)	(0.94)				(1.01)	(1.03)	(1.08)
Mother's years of education							-0.049	-0.086	-0.047				0.207	0.317	0.136
							(-0.15)	(-0.25)	(-0.15)				(0.50)	(0.75)	(0.36)
Income (Monthly) IST							-0.248	-0.122	-0.285				-0.396	-0.283	-0.400
							(-0.80)	(-0.38)	(-0.98)				(-1.08)	(-0.76)	(-1.15)
Expect wallet returned by stranger							-1.446	-1.533	-0.822				-2.548	-2.324	-2.005
							(-0.49)	(-0.50)	(-0.29)				(-0.74)	(-0.67)	(-0.62)
Share of network who expect wallet retu d							-13.988	-12.562	-10.981				-12.635	-12.466	-8.120
							(-1.72)	(-1.52)	(-1.45)				(-1.29)	(-1.26)	(-0.89)
Normalised Indegree										0.024	0.023	0.089	-0.097	-0.112	-0.060
										(0.16)	(0.16)	(0.66)	(-0.50)	(-0.57)	(-0.34)
Normalised Outdegree										-0.014	0.080	0.019	-0.077	0.032	-0.009
										(-0.10)	(0.58)	(0.15)	(-0.46)	(0.18)	(-0.05)
Eigenvector Centrality										6.608	4.633	7.593	0.428	-2.388	4.813
										(0.47)	(0.33)	(0.58)	(0.02)	(-0.13)	(0.28)
Betweenness										-0.043*	-0.055*	-0.036	-0.021	-0.031	-0.017
										(-1.98)	(-2.53)	(-1.78)	(-0.80)	(-1.19)	(-0.71)
o.SiSwati													0.000	0.000	0.000
													(.)	(.)	(.)
Constant	-1.671	-1.941	-1.750	-8.610	-9.675	-8.732	-6.430	-9.949	-5.948	-1.059	-3.716	-4.553	-13.236	-17.363	-17.407
	(-1.57)	(-1.79)	(-1.76)	(-0.95)	(-1.05)	(-1.05)	(-0.62)	(-0.96)	(-0.62)	(-0.20)	(-0.70)	(-0.92)	(-0.79)	(-1.02)	(-1.10)
Observations	152	152	152	146	146	146	124	124	124	152	152	152	120	120	120

1. t-statistics in parentheses.
2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

s

Table 16: Robustness checks of social learning result to alternative definitions of the social learning variable, M.

	(1) inputchange	(2) inputchange	(3) inputchange	(4) inputchange	(5) inputchange
M2	1.119***				
	(3.79)				
M3		1.208***			
		(4.32)			
M4			1.088***		
			(3.78)		
M5				1.070***	
				(4.48)	
M7					4.171***
					(3.51)
Age (years)	0.441	0.391	0.420	0.336	0.363
	(1.13)	(1.03)	(1.08)	(0.89)	(0.92)
Gender	-1.290	-1.446	-1.239	-1.895	-1.552
(1=Female)	(-0.46)	(-0.53)	(-0.45)	(-0.70)	(-0.54)
Coloured	4.348	3.754	4.445	5.693	3.015
	(0.49)	(0.43)	(0.50)	(0.66)	(0.33)
Indian	-0.644	-0.605	-0.640	1.150	2.202
	(-0.04)	(-0.04)	(-0.04)	(0.08)	(0.15)
White	4.999	5.588	4.991	19.671	6.159
	(0.34)	(0.39)	(0.34)	(1.35)	(0.42)
isiNdebele	-8.015	-5.903	-7.845	-6.674	-5.285
	(-0.79)	(-0.59)	(-0.77)	(-0.67)	(-0.51)
isiZulu	3.492	4.597	3.480	3.677	3.402
	(1.07)	(1.42)	(1.07)	(1.16)	(1.03)
Sepedi	-6.134	-5.514	-6.220	-3.530	-3.988
	(-1.07)	(-0.98)	(-1.08)	(-0.62)	(-0.68)
Sesotho	-1.334	-0.165	-1.364	-1.569	0.379
	(-0.32)	(-0.04)	(-0.33)	(-0.39)	(0.09)
Setswana	0.333	0.777	0.261	0.700	0.711
	(0.07)	(0.16)	(0.05)	(0.15)	(0.15)
o.SiSwati	0.000	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)	(.)
Tshivenda	-7.334	-6.264	-7.240	-5.761	-6.018
	(-0.50)	(-0.44)	(-0.50)	(-0.40)	(-0.41)
isiTsonga	-2.000	1.463	-1.860	-2.482	1.464
	(-0.15)	(0.11)	(-0.14)	(-0.19)	(0.11)
Afrikaans	-4.925	-4.436	-4.831	-4.850	-4.461
	(-0.52)	(-0.47)	(-0.51)	(-0.52)	(-0.46)
English	-4.788	-4.020	-4.699	-5.999	-6.217
	(-0.71)	(-0.61)	(-0.70)	(-0.93)	(-0.93)
Years of	1.009	1.027	0.991	0.872	1.225
Education	(1.06)	(1.11)	(1.04)	(0.94)	(1.29)
Mother's years	0.204	0.231	0.191	0.266	0.126
of education	(0.50)	(0.57)	(0.47)	(0.66)	(0.31)
Income (Monthly)	-0.393	-0.434	-0.383	-0.247	-0.389
IST	(-1.06)	(-1.20)	(-1.04)	(-0.68)	(-1.04)
Expect wallet	-2.537	-2.253	-2.599	-2.818	-2.640
returned by st r	(-0.74)	(-0.67)	(-0.75)	(-0.85)	(-0.76)
Share of network	-13.339	-11.564	-13.037	-11.745	-8.102
who expect wal d	(-1.36)	(-1.20)	(-1.33)	(-1.23)	(-0.82)
Normalised	-0.089	-0.087	-0.098	-0.076	-0.066
Indegree	(-0.46)	(-0.46)	(-0.51)	(-0.41)	(-0.34)
Normalised	-0.090	-0.061	-0.083	-0.115	-0.113
Outdegree	(-0.54)	(-0.37)	(-0.50)	(-0.71)	(-0.67)
Eigenvector	-0.631	-4.771	0.334	2.323	-5.734
Centrality	(-0.03)	(-0.27)	(0.02)	(0.13)	(-0.31)
Betweenness	-0.020	-0.015	-0.020	-0.024	-0.014
	(-0.79)	(-0.58)	(-0.80)	(-0.96)	(-0.55)
Constant	-13.899	-14.132	-13.157	-11.222	-15.270
	(-0.82)	(-0.85)	(-0.78)	(-0.68)	(-0.90)
Observations	120	120	120	120	120

1. t-statistics in parentheses.
2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 17: Logit model: Factors predicting various alternative network relationships

	(1) Talk Activate	(2) Community Project	(3) Project Assistance	(4) Borrow Money
main				
Difference in Age	0.034***	0.008	-0.011	-0.041
(Years)	(0.01)	(0.01)	(0.01)	(0.02)
Same Home Language	0.283***	0.543***	0.473***	0.491*
	(0.08)	(0.10)	(0.11)	(0.24)
Same Religion	-0.026	-0.079	-0.118	-0.187
	(0.08)	(0.10)	(0.11)	(0.23)
Same Gender	-0.130	0.170	0.103	0.321
	(0.07)	(0.09)	(0.10)	(0.22)
Both Coloured	-0.507	-0.200	0.007	0.412
	(0.29)	(0.35)	(0.34)	(0.62)
o.Both Indian	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)
Both White	-1.656			
	(1.08)			
Black-Coloured	0.019	-0.056	0.063	0.316
	(0.11)	(0.14)	(0.15)	(0.32)
Black-Indian	0.070	0.645*	0.665*	-0.810
	(0.25)	(0.26)	(0.27)	(1.02)
Black-White	-0.693**	-1.362**	-1.442**	-0.930
	(0.23)	(0.46)	(0.52)	(1.02)
Coloured-Indian	1.734	2.862*		
	(1.09)	(1.12)		
Coloured-White	-0.207	0.017	0.208	
	(0.37)	(0.45)	(0.46)	
o.Indian-White	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)
Both Employed	-0.068	-0.101	0.063	0.297
	(0.08)	(0.11)	(0.12)	(0.28)
Both Unemployed	-0.613*	-0.647	-0.196	-0.620
	(0.27)	(0.35)	(0.34)	(1.03)
Difference in Income	-0.006	0.001	0.010	-0.024
(Monthly)	(0.01)	(0.01)	(0.01)	(0.02)
Difference in	-0.000	0.000	0.000***	-0.000
Reservation Wage	(0.00)	(0.00)	(0.00)	(0.00)
Ladder difference	-0.050	0.056	0.060	-0.113
age 15	(0.03)	(0.04)	(0.04)	(0.07)
Ladder difference	0.110***	0.026	-0.008	0.033
now	(0.03)	(0.04)	(0.04)	(0.09)
Dif in years of	-0.006	-0.029	-0.034	0.027
education	(0.02)	(0.02)	(0.02)	(0.05)
Dif in Mother's	-0.003	-0.009	-0.010	-0.016
years of educ	(0.01)	(0.01)	(0.01)	(0.02)
Same Group	3.453***	3.071***	3.166***	2.518***
	(0.09)	(0.14)	(0.15)	(0.32)
o.Same Training	0.000	0.000	0.000	0.000
Province	(.)	(.)	(.)	(.)
o.Both White		0.000	0.000	0.000
		(.)	(.)	(.)
o.Coloured-Indian			0.000	0.000
			(.)	(.)
o.Coloured-White				0.000
				(.)
Constant	-3.235***	-4.430***	-4.763***	-6.471***
	(0.11)	(0.15)	(0.17)	(0.34)
Observations	7016	7007	7000	6939

1. Robust standard errors in paratheses
2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
3. DV: Dummy variable equal to 1 if the relationship between two individuals exists. The relationship type in each column is indicated by the column heading.
4. Same Province variable and Both Indian variable predict failure

Table 18: Logit model: Factors predicting whether two participants have ever communicated (Marginal Effects)

	(1) Demographic	(2) Race	(3) Socioecono c	(4) Education	(5) Activate T g	(6) All
Difference in Age)	0.000					0.001
Same Home Language	0.119					0.032
Same Religion	0.021					0.011
Same Gender	0.011					0.005
Same Race	-0.009					
Both Coloured		0.116				-0.001
Both Indian		0.390				0.586
Both White		0.446				0.017
Black-Coloured		-0.016				0.008
Black-Indian		0.009				0.024
Black-White		-0.023				0.000
Coloured-Indian		-0.023				0.073
Coloured-White		0.141				-0.008
Indian-White		-0.043				
Both Employed			0.009			0.002
Both Unemployed			0.046			0.011
Difference in Inco)			-0.000			-0.001
Difference in Rese e			0.000			0.000
Ladder difference 15			-0.003			-0.002
Ladder difference w			0.003			0.003
Dif in years of ed n				-0.001		-0.002
Dif in Mother's ye c				-0.001		-0.000
Same Group					0.246	0.252
Same Training Prov e					0.316	0.320
Observations	36669	39194	34261	30190	35530	21014

1. This table presents the marginal effects for the logit regressions in Table 11.
2. Due to the specification containing dummy variables, the marginal effects reported are average partial effects
3. The dependent variable in this regression is a dummy variable equal to 1 if an individual identifies as having a communication relationship with a particular individual and equal to zero otherwise.
4. Data are used in their dyadic form (i.e. one observation for each potential relationship).

Table 19: Logit model: Factors predicting various alternative network relationships (Marginal Effects)

	(1) Talk Activ e	(2) Community t	(3) Project As e	(4) Borrow Money
Difference in Age)	0.004	0.001	-0.001	-0.001
Same Home Language	0.031	0.043	0.033	0.008
Same Religion	-0.003	-0.005	-0.007	-0.002
Same Gender	-0.014	0.012	0.007	0.005
Both Coloured	-0.053	-0.013	0.000	0.006
Both White	-0.144			
Black-Coloured	0.002	-0.004	0.004	0.005
Black-Indian	0.008	0.053	0.050	-0.007
Black-White	-0.071	-0.059	-0.054	-0.007
Coloured-Indian	0.207	0.328		
Coloured-White	-0.022	0.001	0.014	
Both Employed	-0.007	-0.007	0.004	0.004
Both Unemployed	-0.063	-0.036	-0.011	-0.006
Difference in Inco)	-0.001	0.000	0.001	-0.000
Difference in Rese e	-0.000	0.000	0.000	-0.000
Ladder difference 15	-0.005	0.004	0.004	-0.001
Ladder difference w	0.012	0.002	-0.001	0.000
Dif in years of ed n	-0.001	-0.002	-0.002	0.000
Dif in Mother's ye c	-0.000	-0.001	-0.001	-0.000
Same Group	0.475	0.355	0.348	0.106
Observations	7016	7007	7000	6939

1. This table presents the marginal effects for the logit regressions in Table 17.
2. Due to the specification containing dummy variables, the marginal effects reported are average partial effects
3. The dependent variable in this regression is a dummy variable equal to 1 if an individual identifies as having a relationship (of a particular type) with an individual and equal to zero otherwise. Relationship types are indicated in column headings.
4. Data are used in their dyadic form (i.e. one observation for each potential relationship).

Table 20: Logit model predicting whether a change in trust game offer occurs between rounds 1 and 3 (Marginal Effects)

	(1) No controls	(2) Demographic	(3) Socioecono c	(4) Centrality1	(5) Centrality2	(6) Centrality3	(7) Centrality4	(8) All controls
Share of good news m	-1.286	-1.476	-1.126	-1.255	-1.278	-1.229	-1.228	-0.624
Share of bad news e	0.264	-0.361	-0.026	0.271	0.269	0.330	0.315	-0.646
Share of good news f	0.263	0.112	0.383	0.224	0.261	0.260	0.233	0.500
Share of bad news f	-0.307	-0.595	-0.583	-0.332	-0.308	-0.263	-0.296	-0.681
Age (years)		0.021						0.021
Gender (1=Female)		-0.087						-0.145
Coloured		0.172						0.202
isiNdebele		-0.230						-0.208
isiZulu		-0.003						0.028
Sepedi		0.090						
Sesotho		-0.068						-0.079
Setswana		-0.092						-0.089
isiTsonga		-0.052						
Afrikaans		0.033						0.038
English		-0.031						-0.001
Years of Education			0.021					0.016
Mother's years of n			-0.001					0.001
Income (Monthly) IST			0.012					0.017
Expect wallet retu g			0.046					0.030
Share of network w l			0.271					0.559
Normalised Indegree				0.001			0.001	0.002
Normalised Outdegree				0.001			0.001	-0.005
Betweenness					0.000		-0.000	-0.001
Eigenvector Centra y						0.244	0.176	0.803
Observations	152	142	124	152	152	152	152	109

1. This table presents the marginal effects for the logit regressions in Table 13.
2. Due to the specification containing dummy variables, the marginal effects reported are average partial effects
3. The dependent variable in this regression is a dummy variable equal to 1 if an individual changed their offer in the trust game between rounds and equal to 0 if they made the same offer in both rounds.

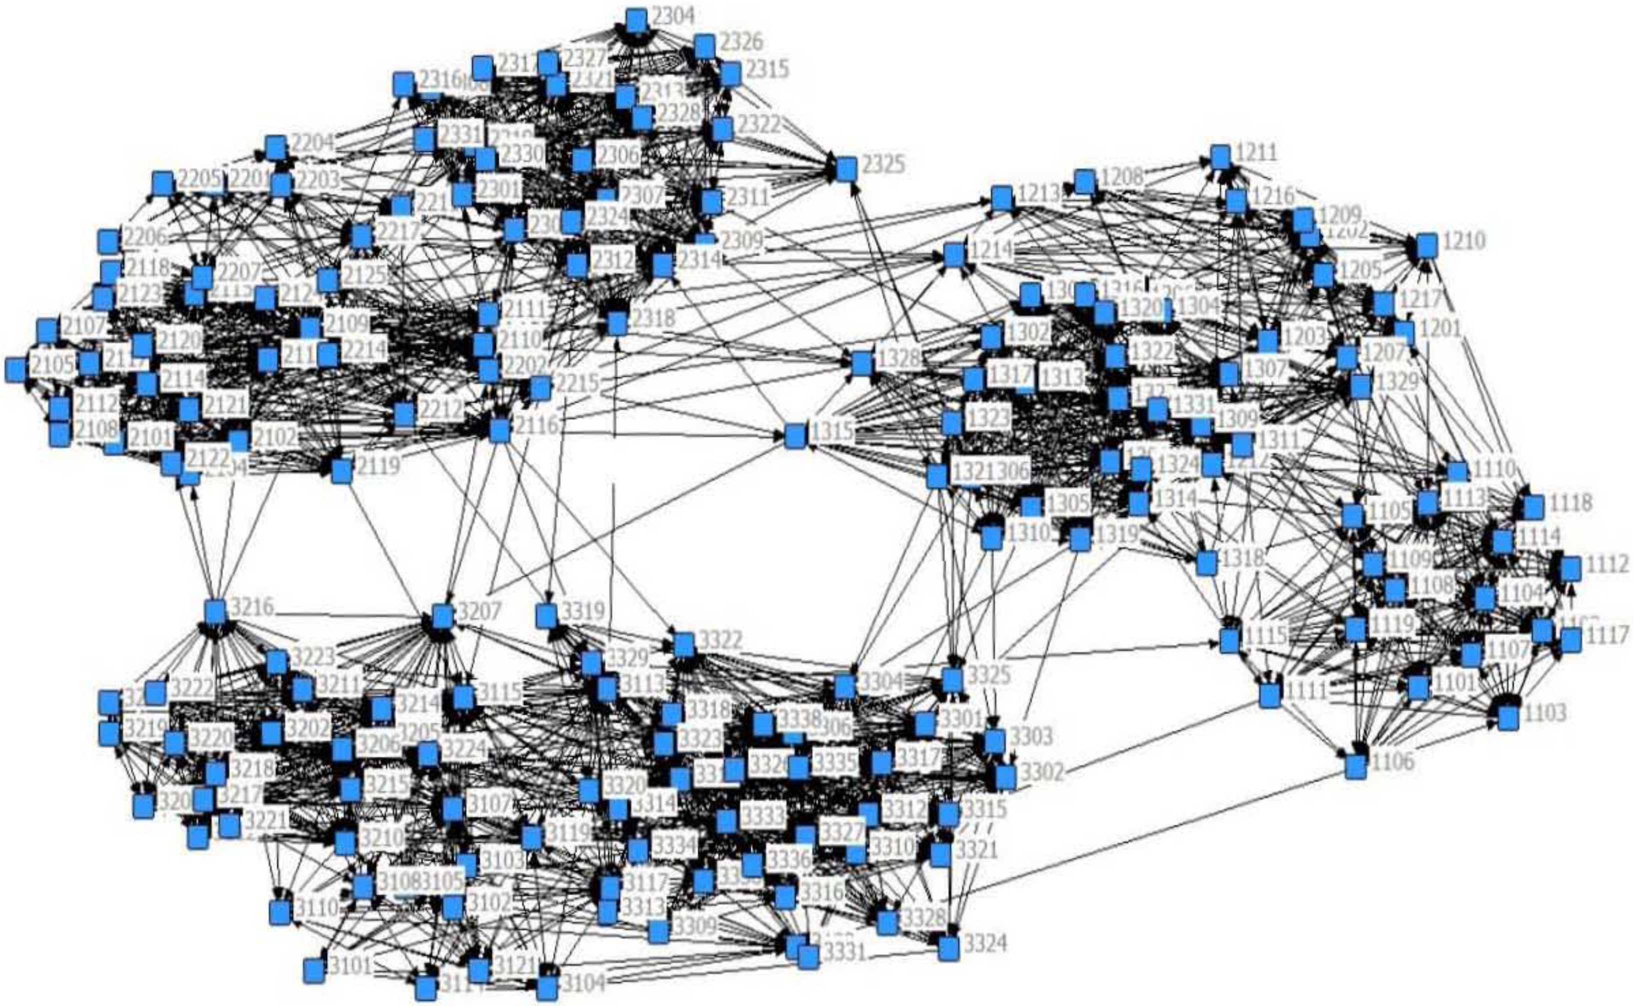

Figure 1: Combined communication network including participants from all three regions

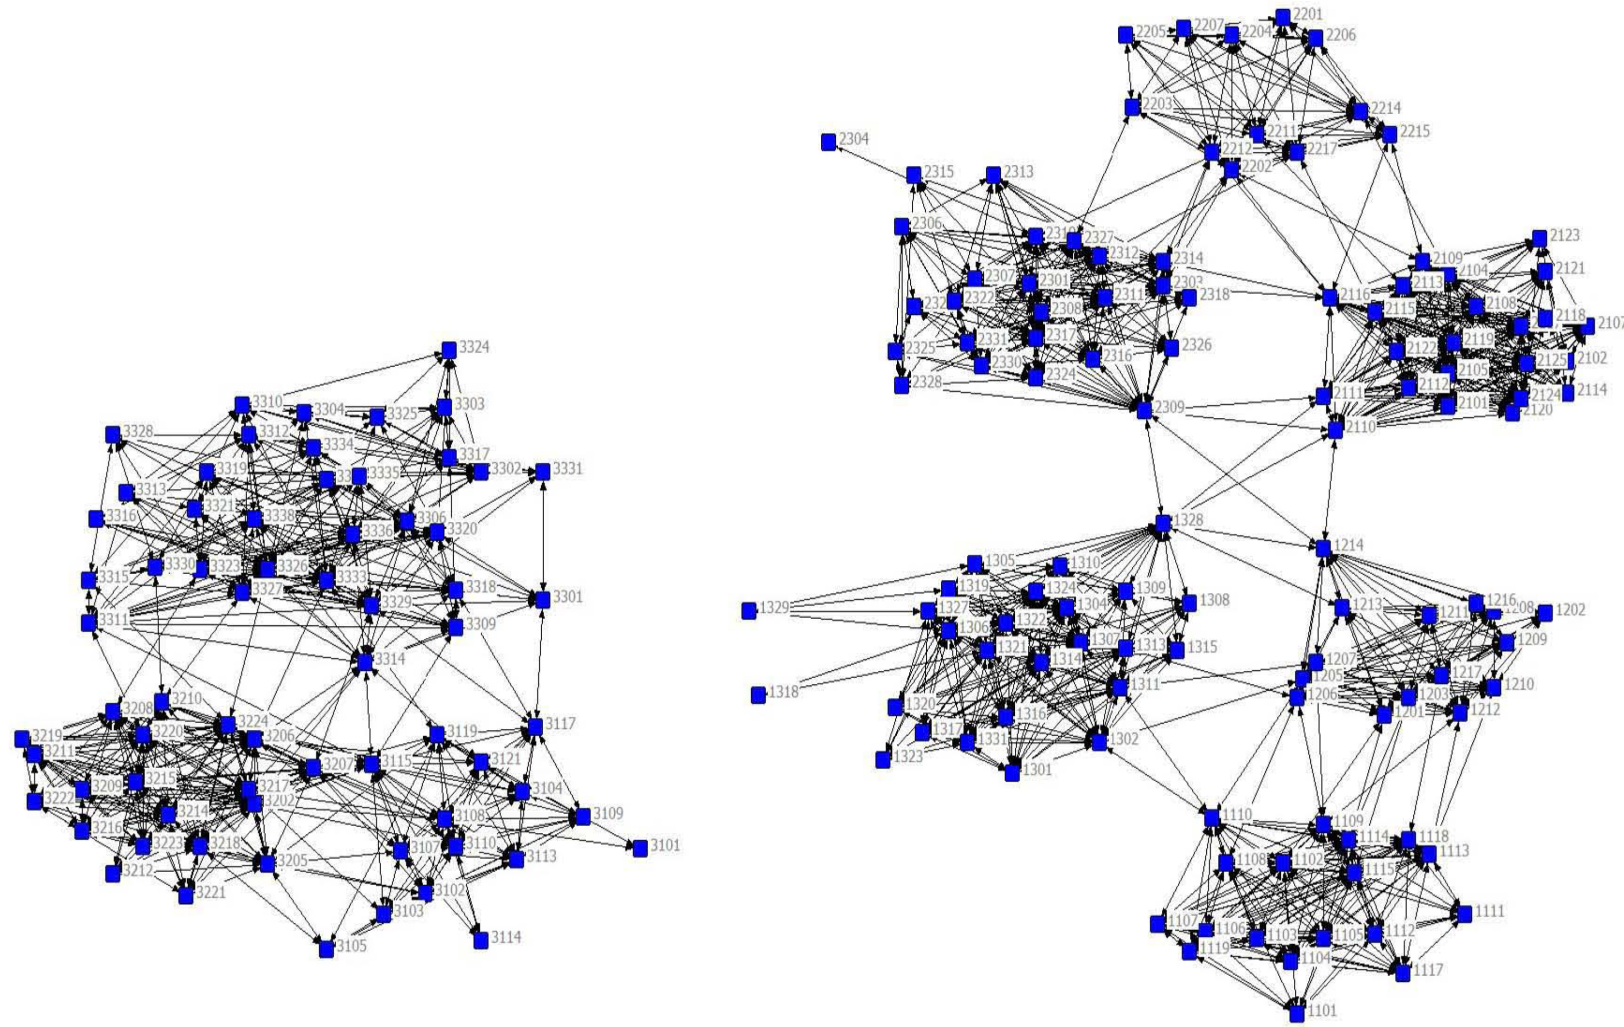

Figure 2: Combined communication network including participants from all three regions: Reciprocal relationships

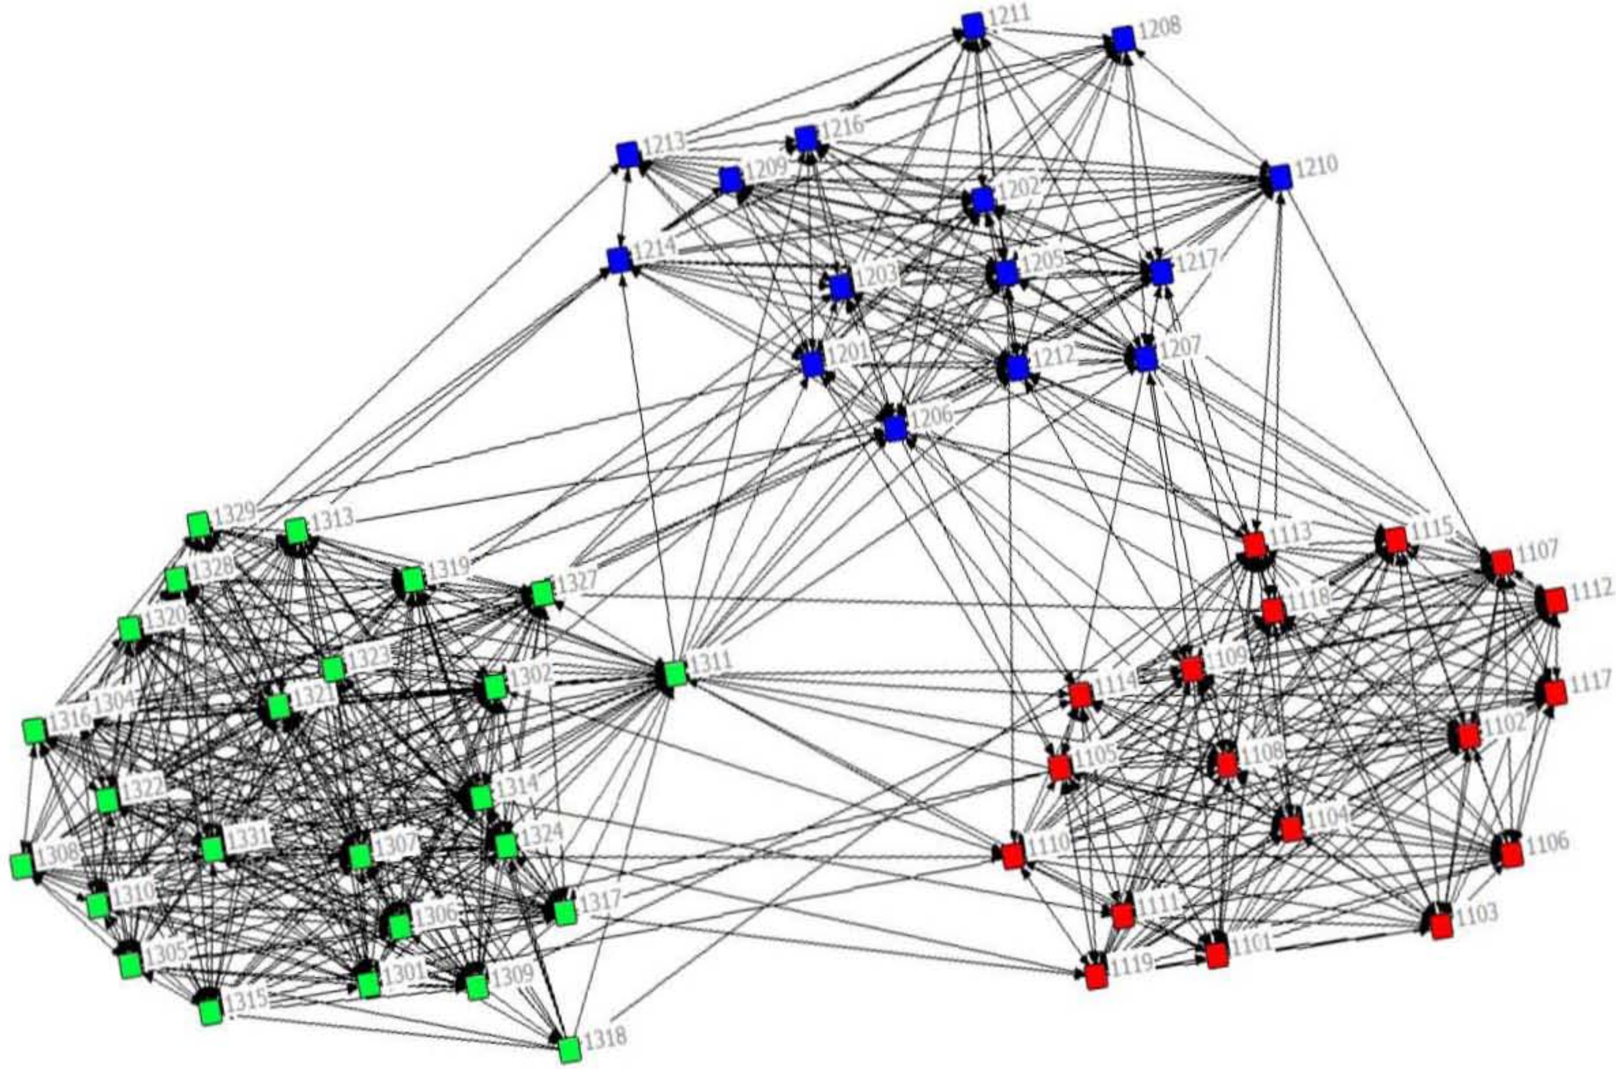

Figure 3: Western Cape communication network

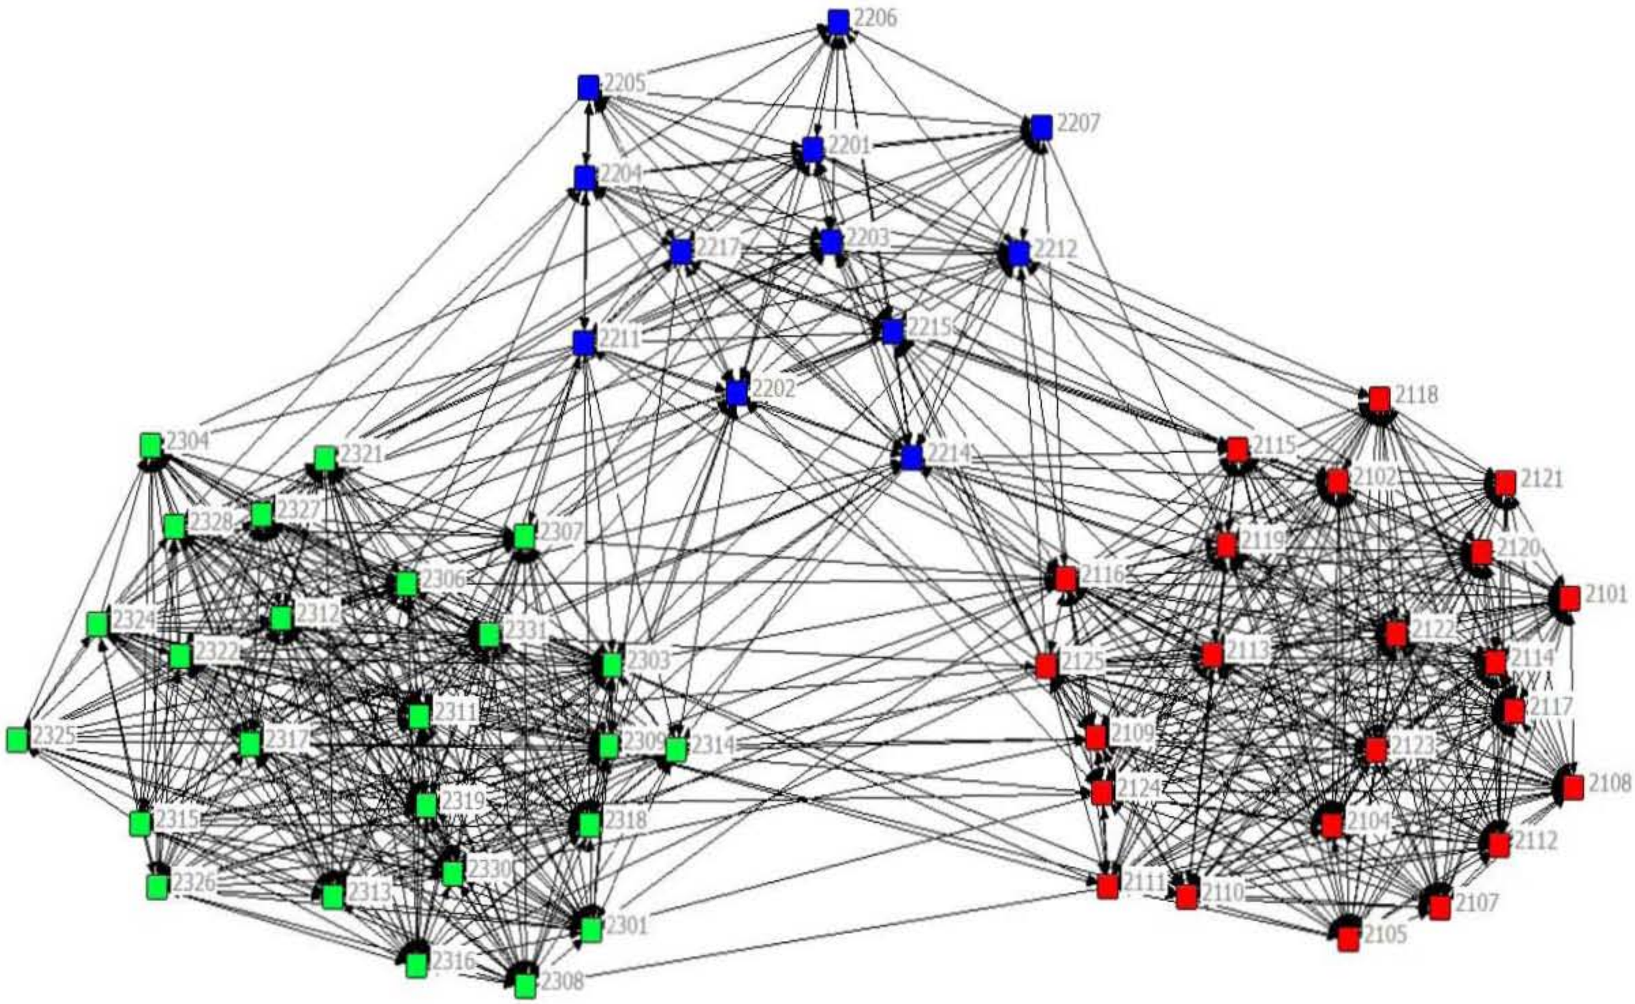

Figure 4: Kwazulu Natal communication network

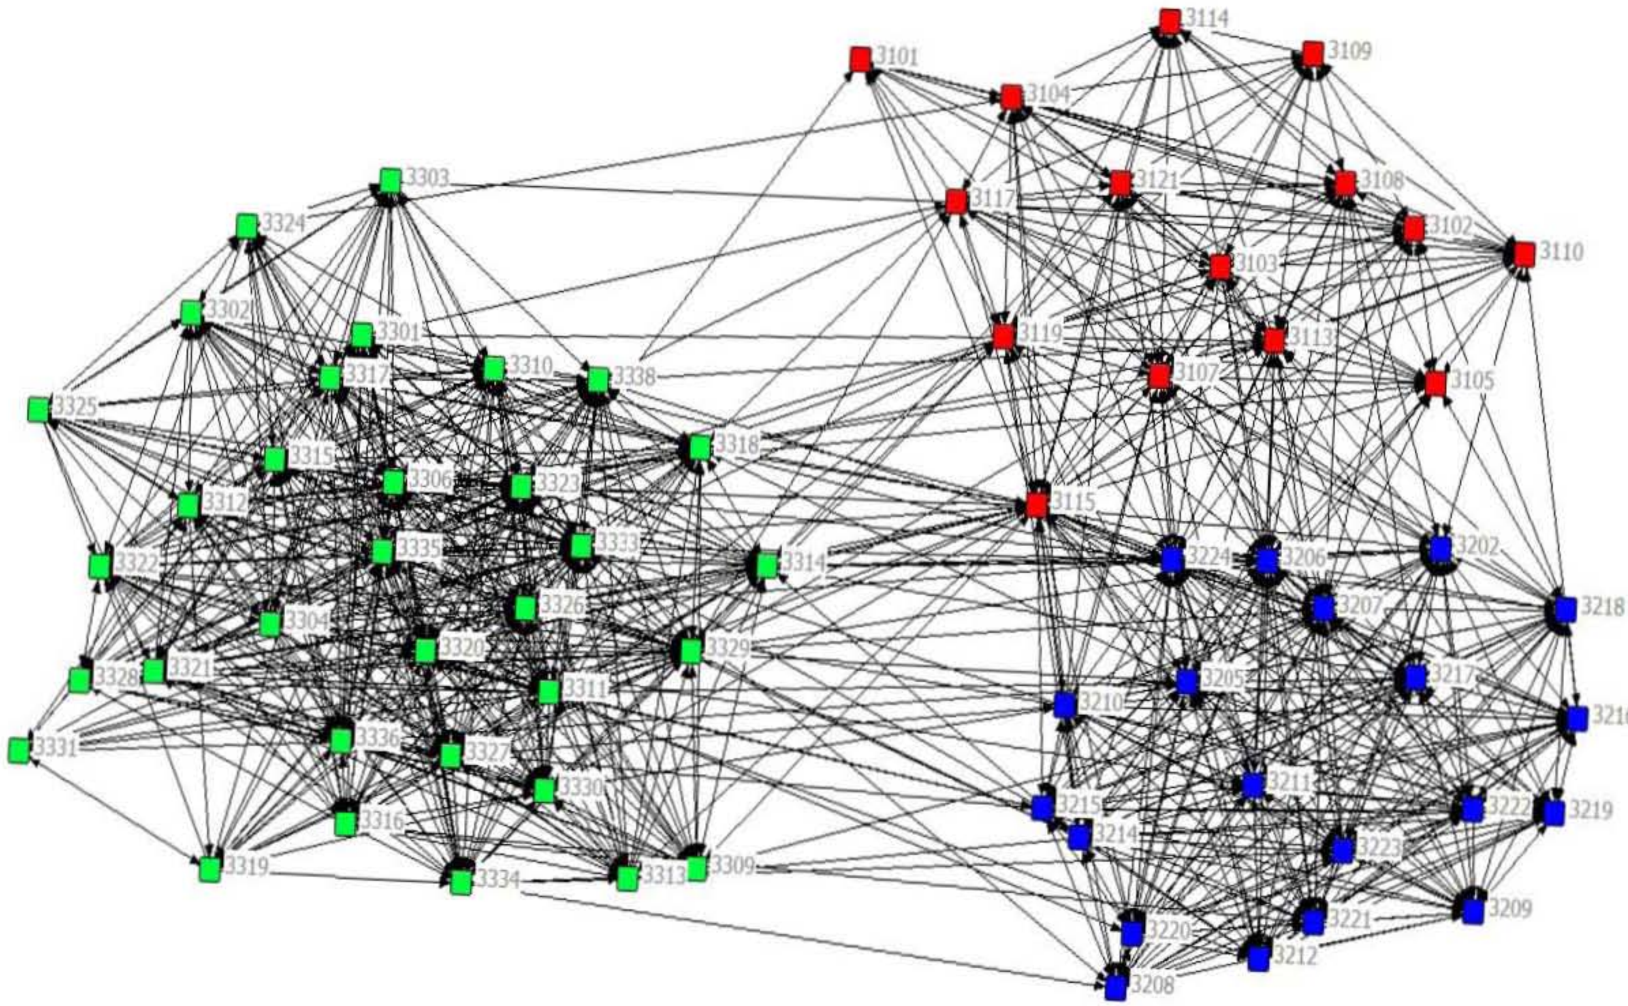

Figure 5: Gauteng communication network

African	Blue
White	Green
Coloured	Red
Indian	Yellow
Missing	Black

Figure 6: Key: Race

Female	Red	Circle
Male	Blue	Triangle

Figure 7: Key: Gender

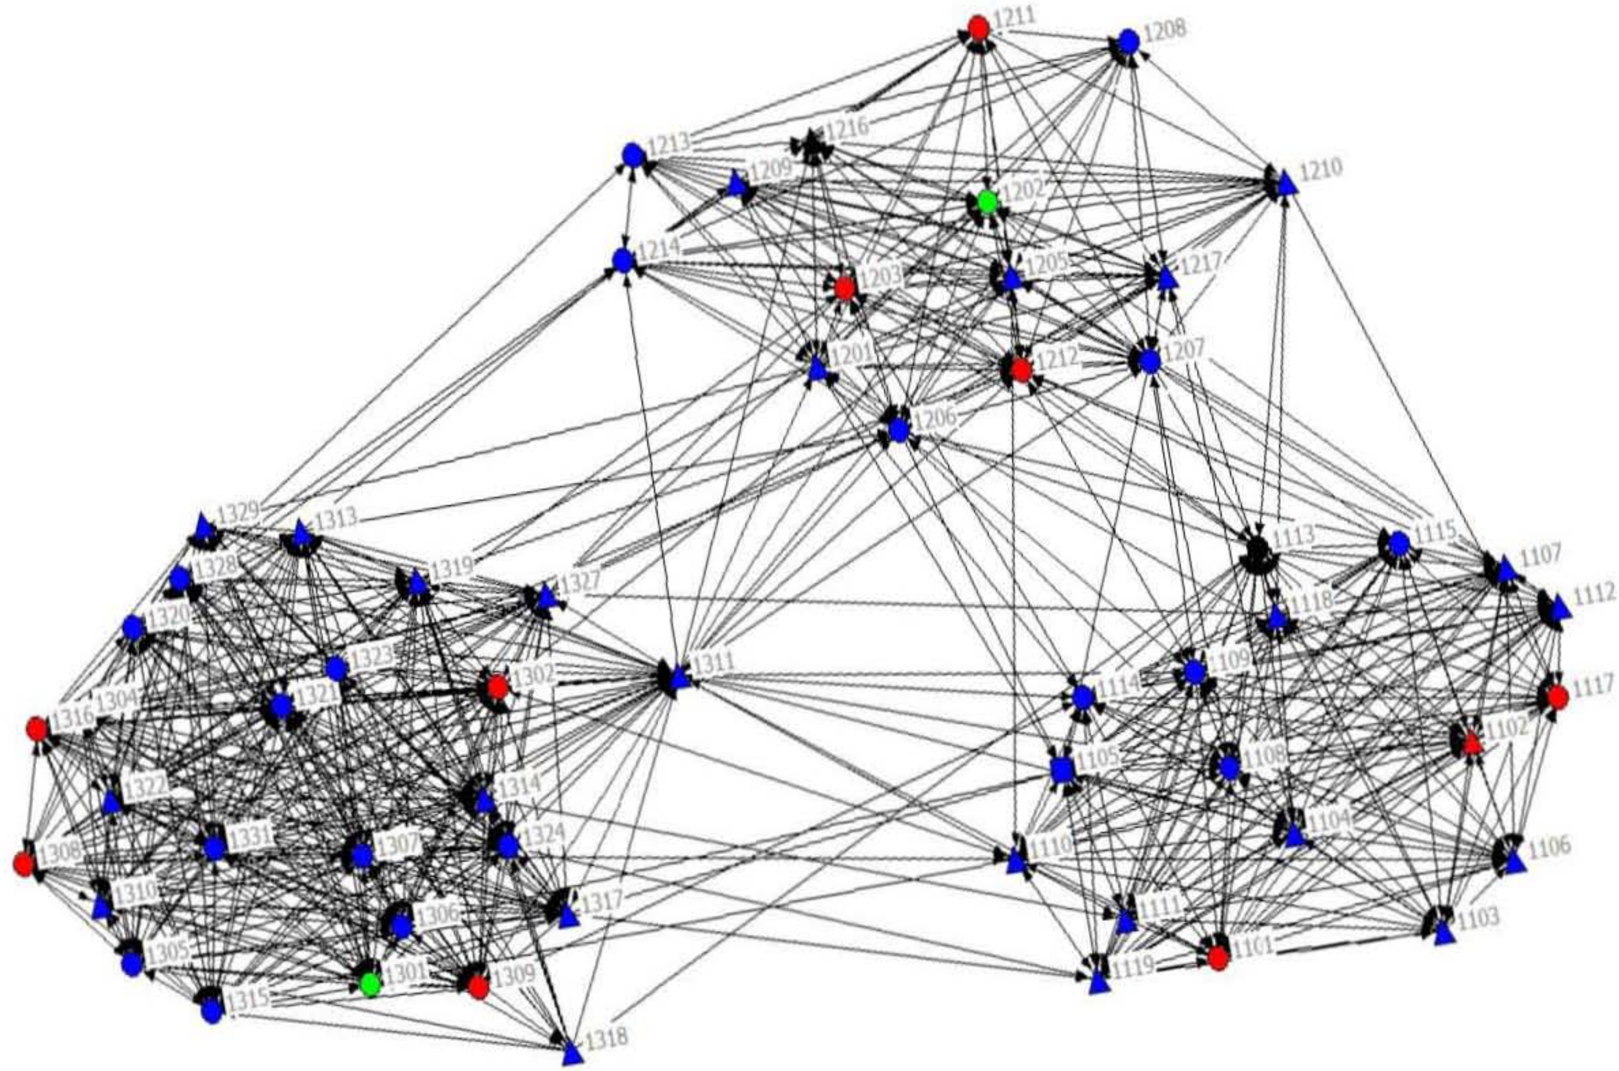

Figure 8: Western Cape network by race and gender.

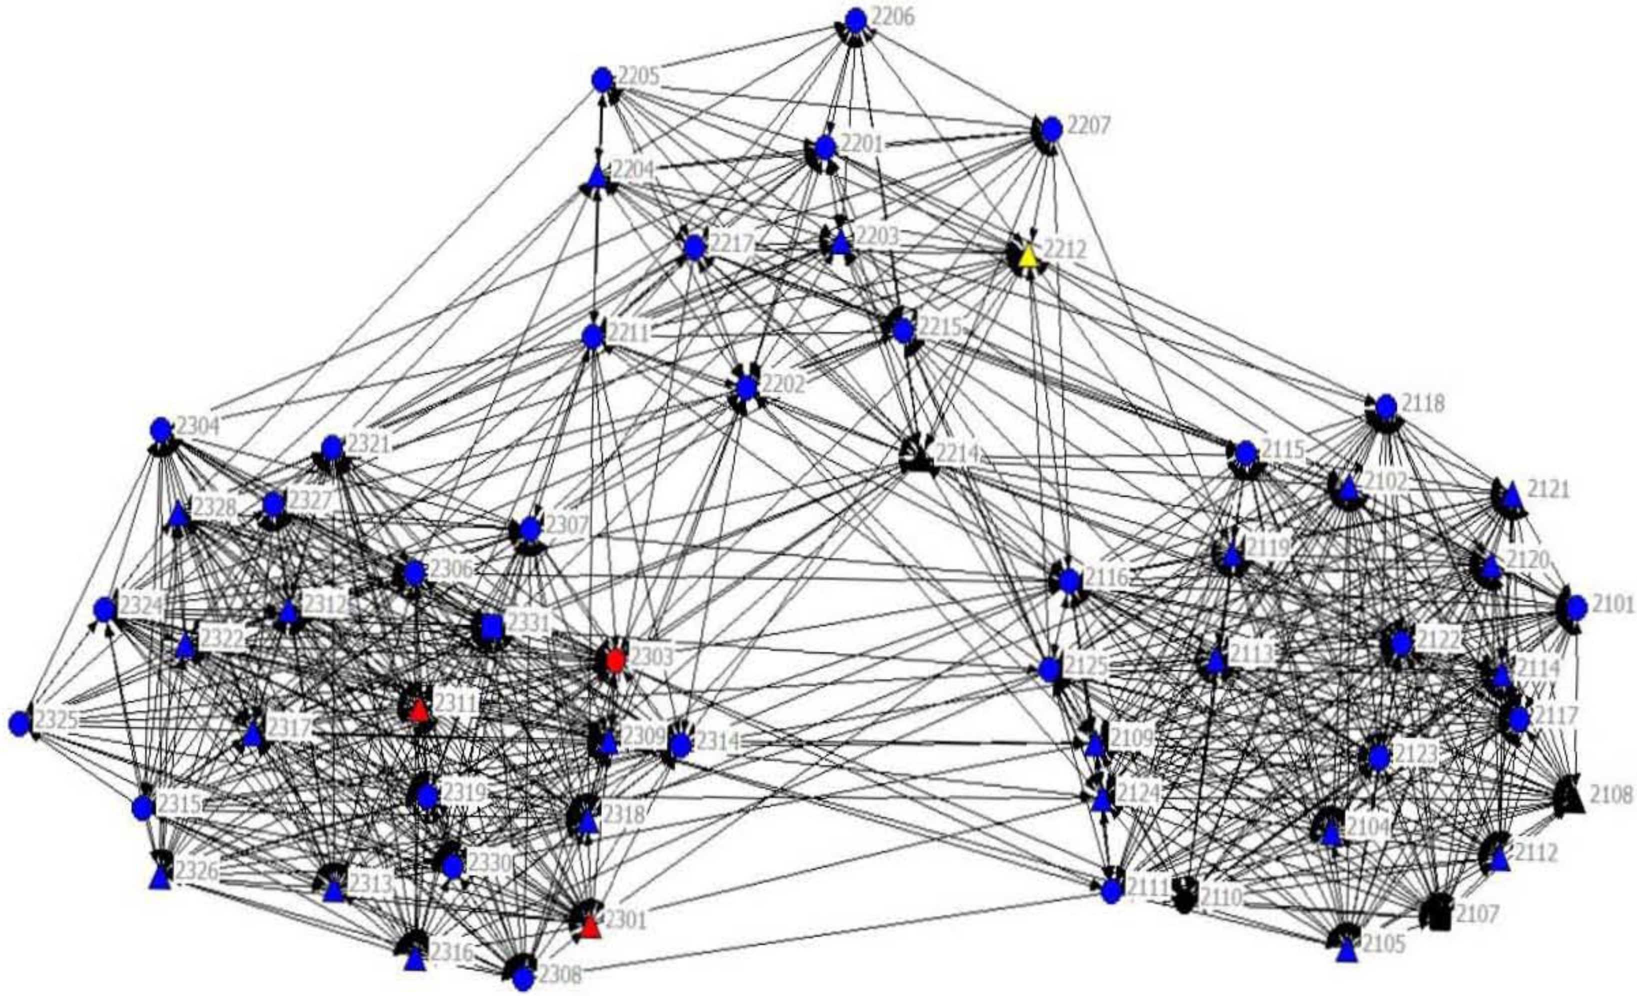

Figure 9: Kwazulu Natal network by race and gender.

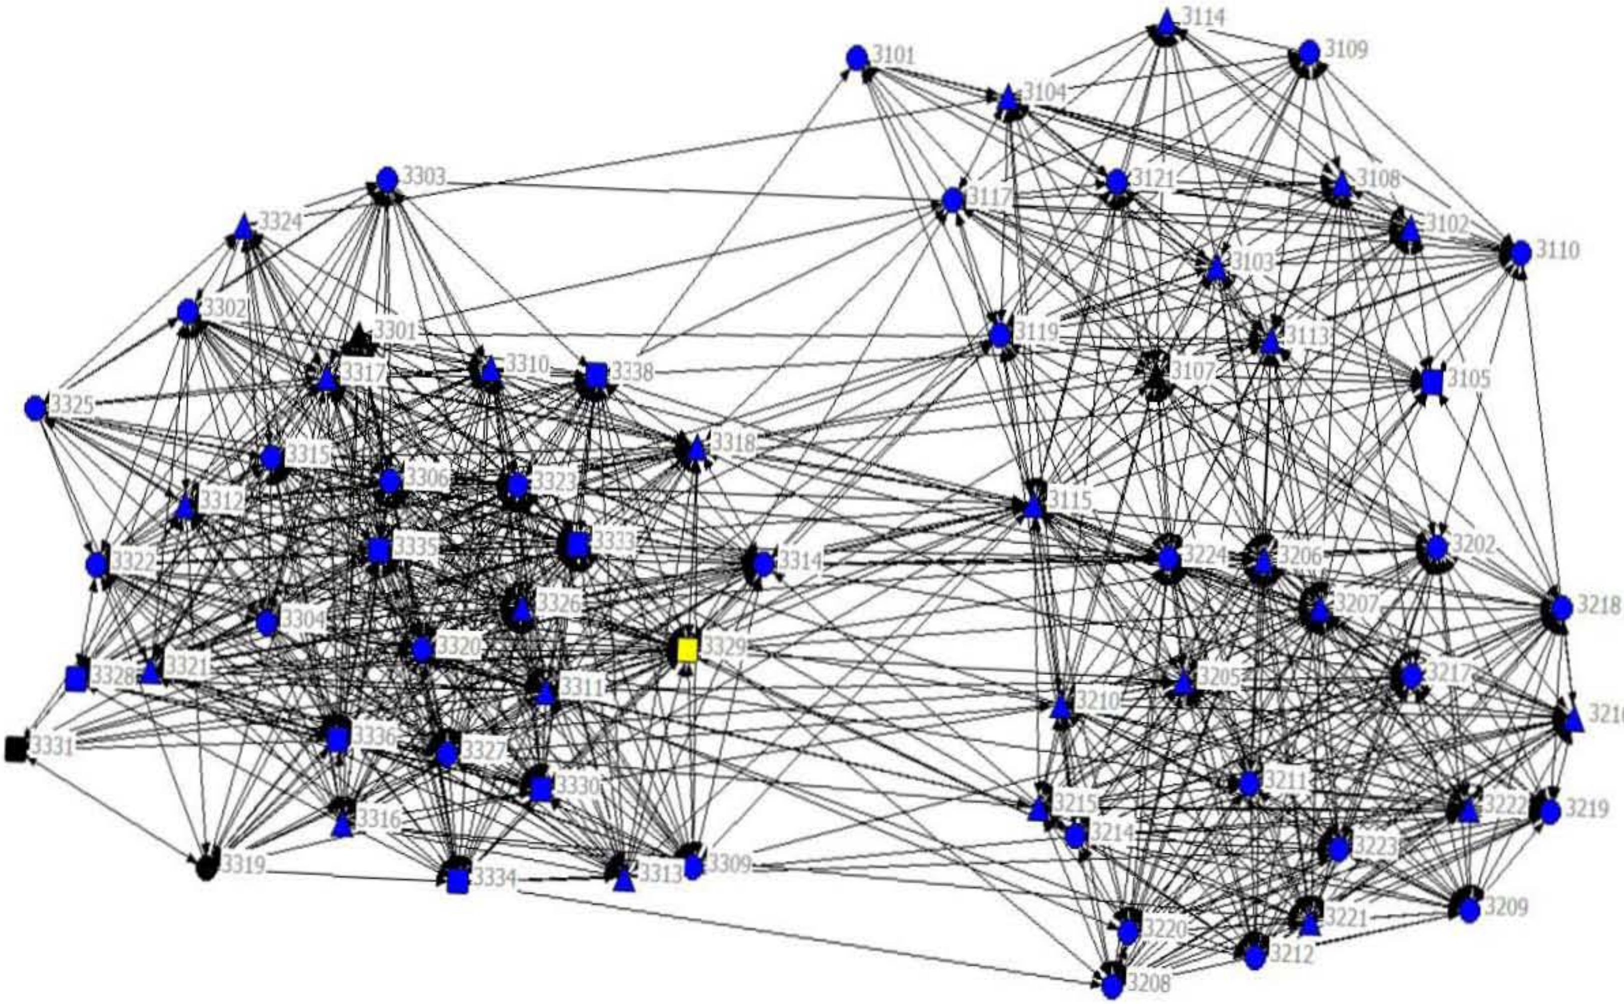

Figure 10: Gauteng network by race and gender.

Afrikaans	Red
English	Green
isiNdebele	Turquoise
isiTsonga	Grey
Sepedi	Orange
Setswana	Navy Blue
Siswati	Brown
Sotho	Pink
Tshivenda	Purple
Xhosa	Blue
Zulu	Yellow
Other	Black

Figure 11: Key: Home Languages

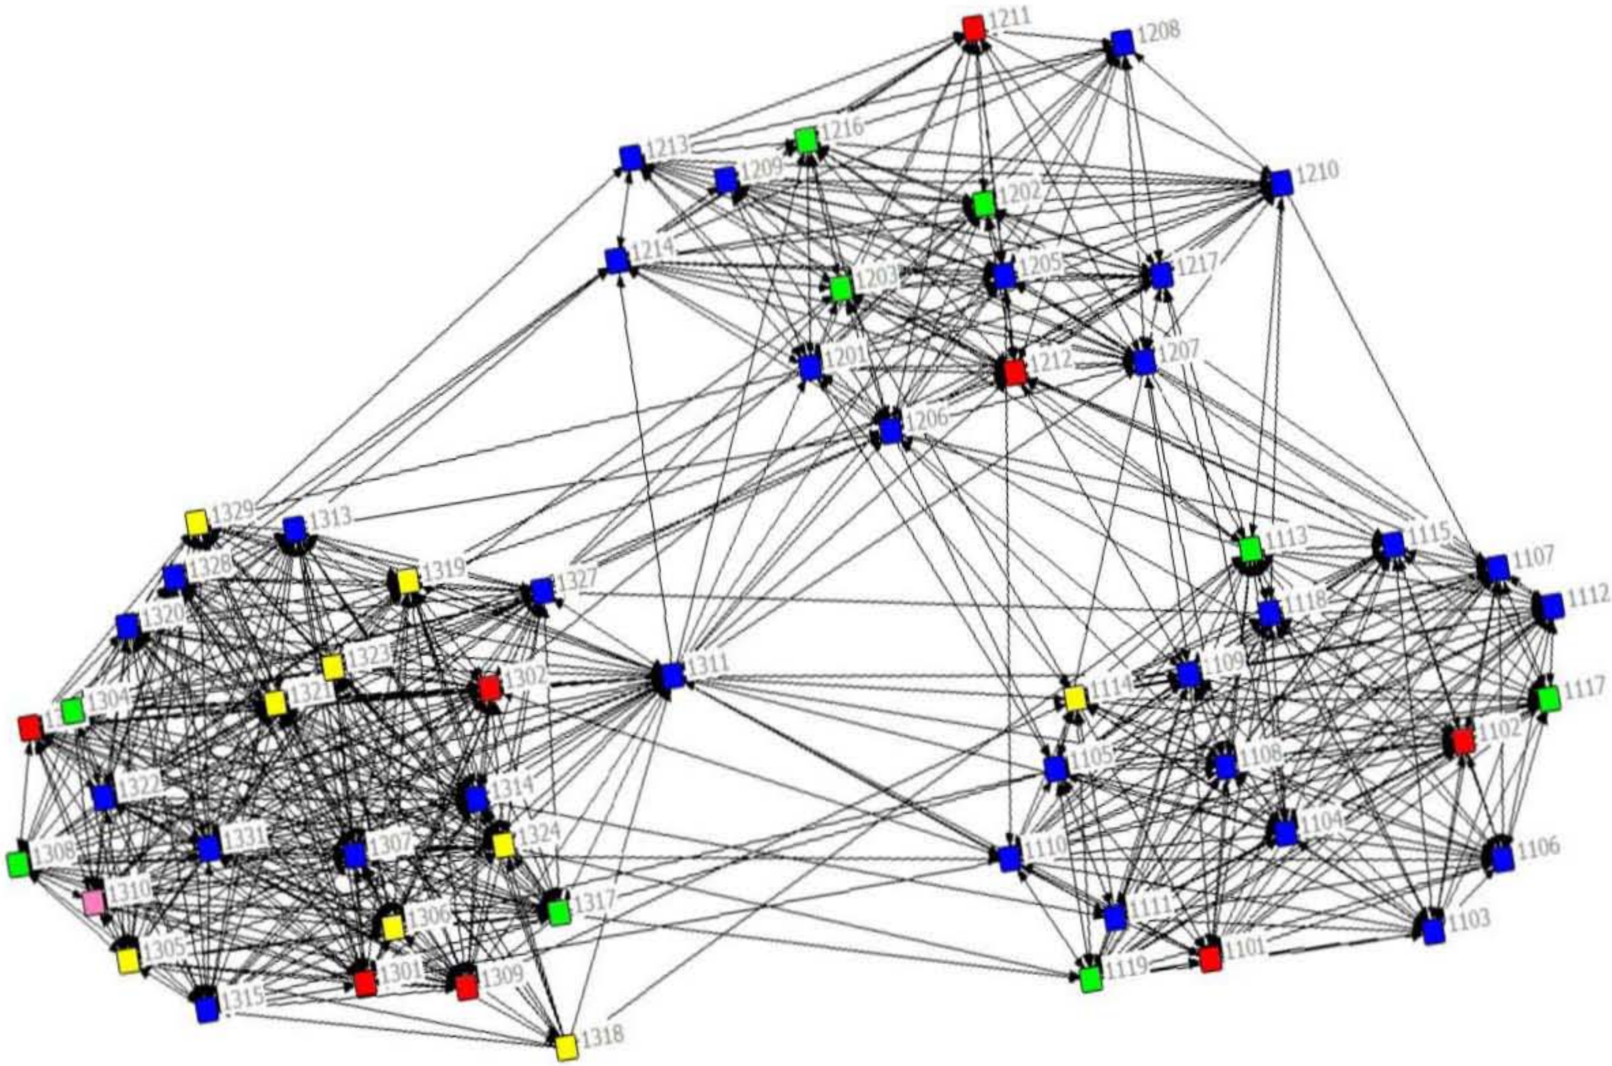

Figure 12: Western Cape network by home language.

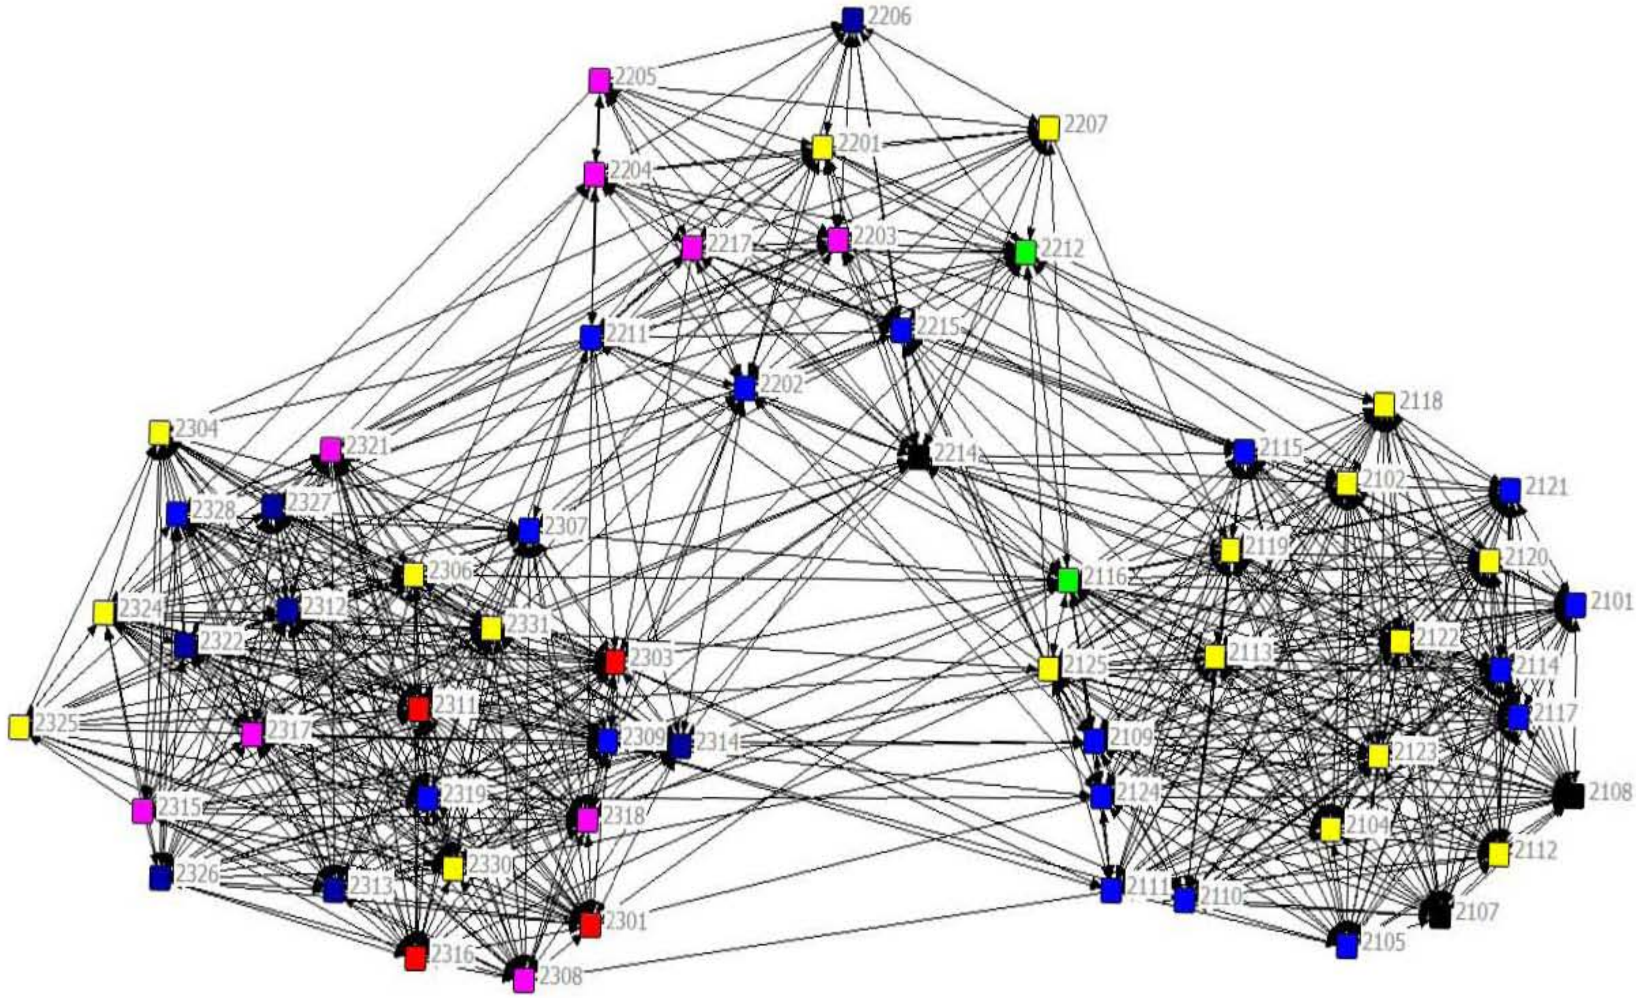

Figure 13: Kwazulu Natal network by home language.

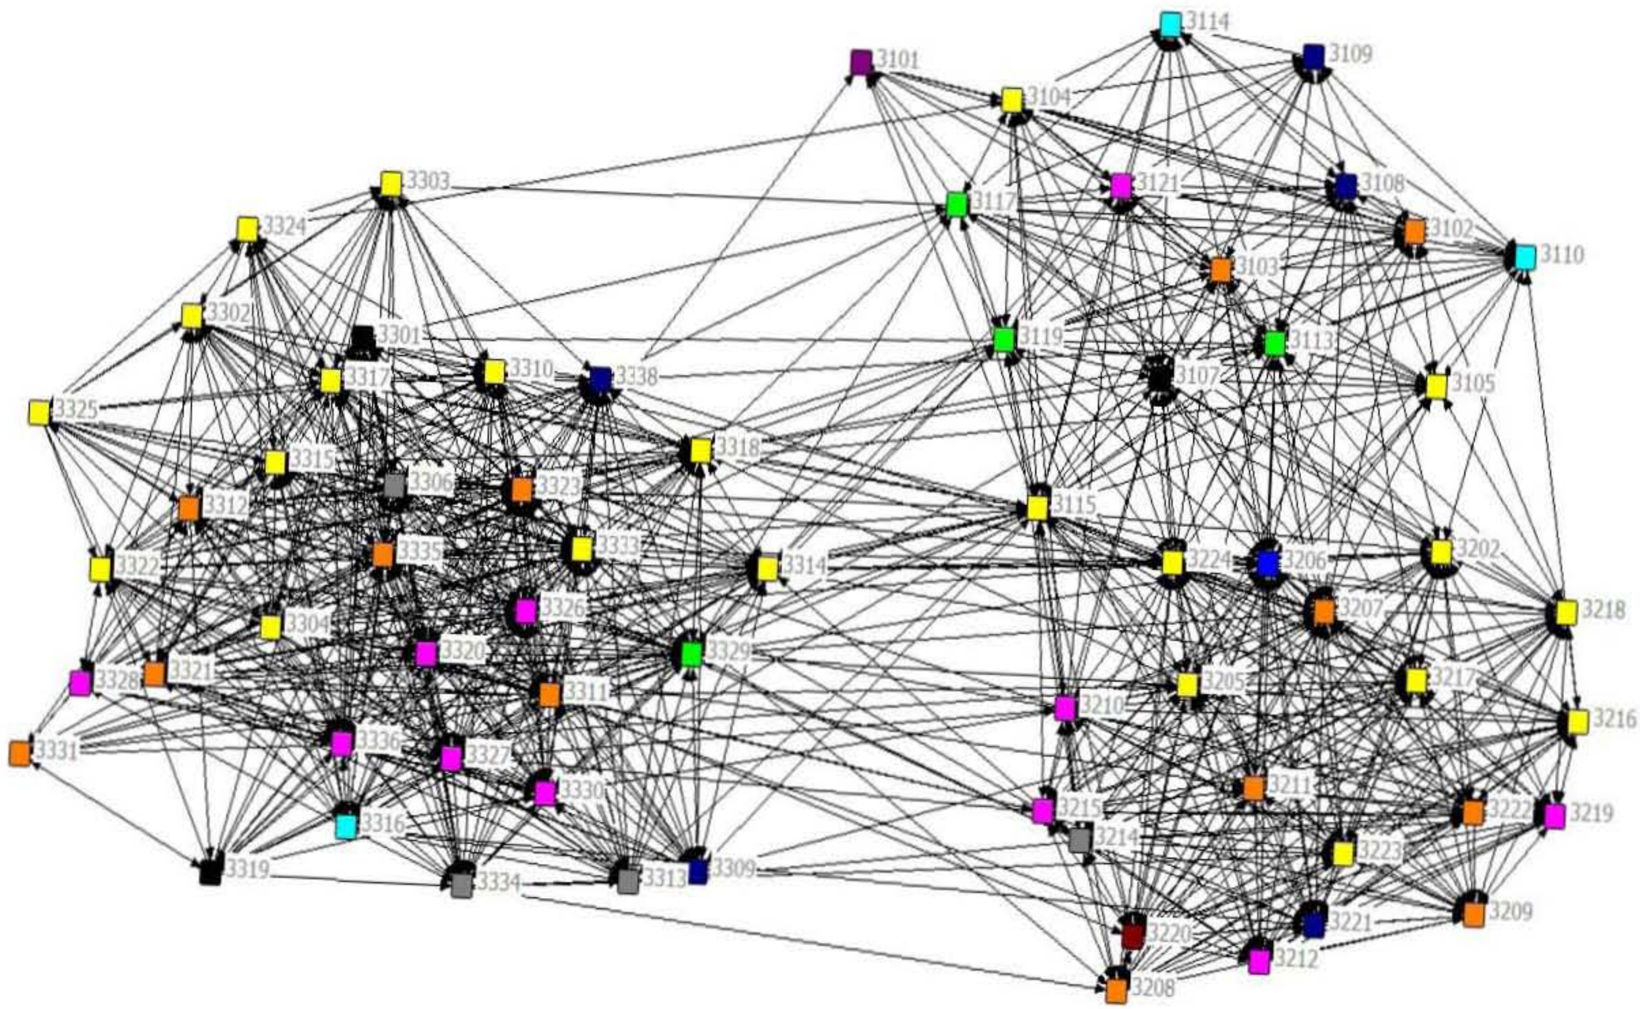

Figure 14: Gauteng network by home language.

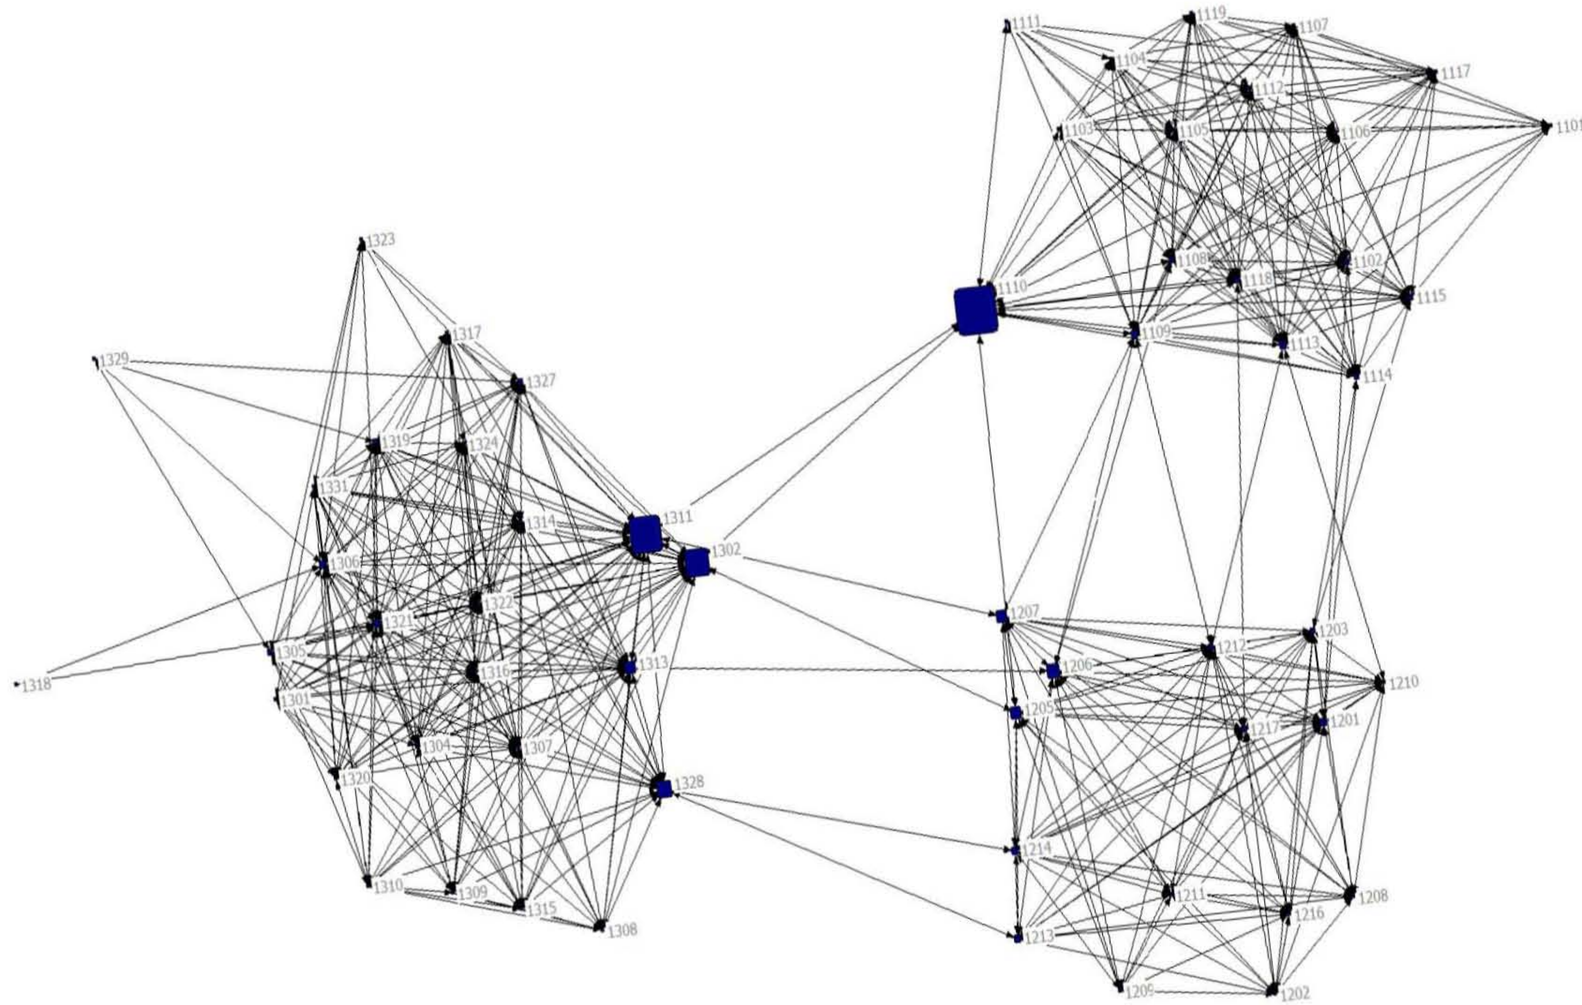

Figure 15: Western Cape: Betweenness

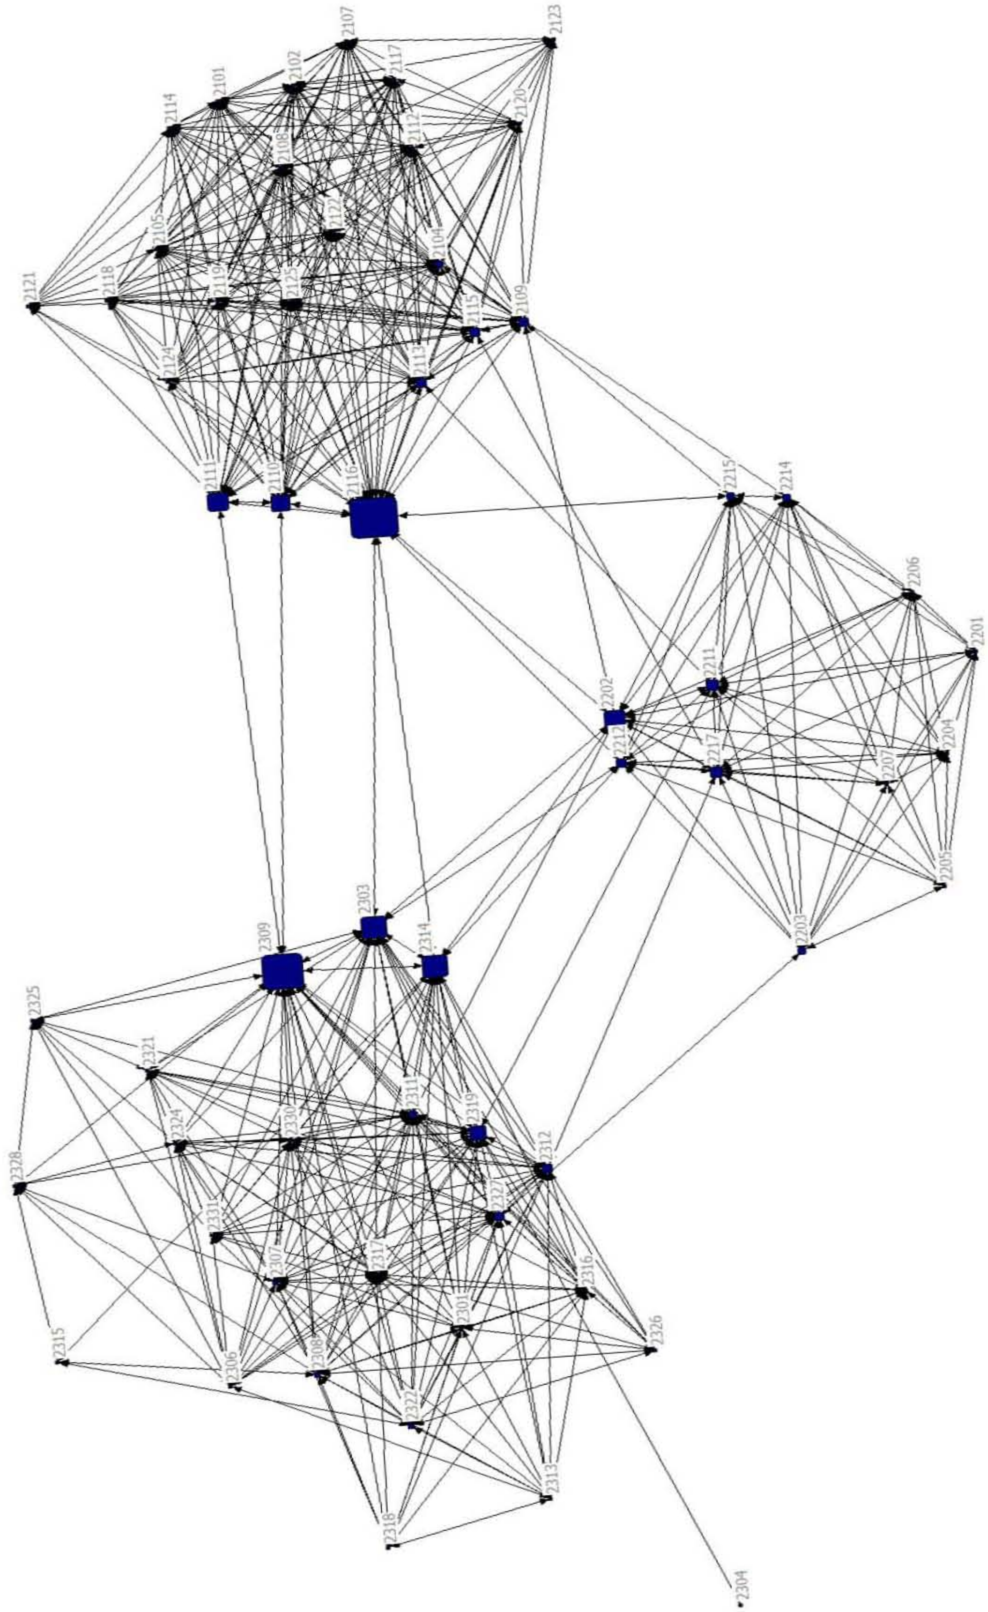

Figure 16: Kwazulu Natal: Betweenness

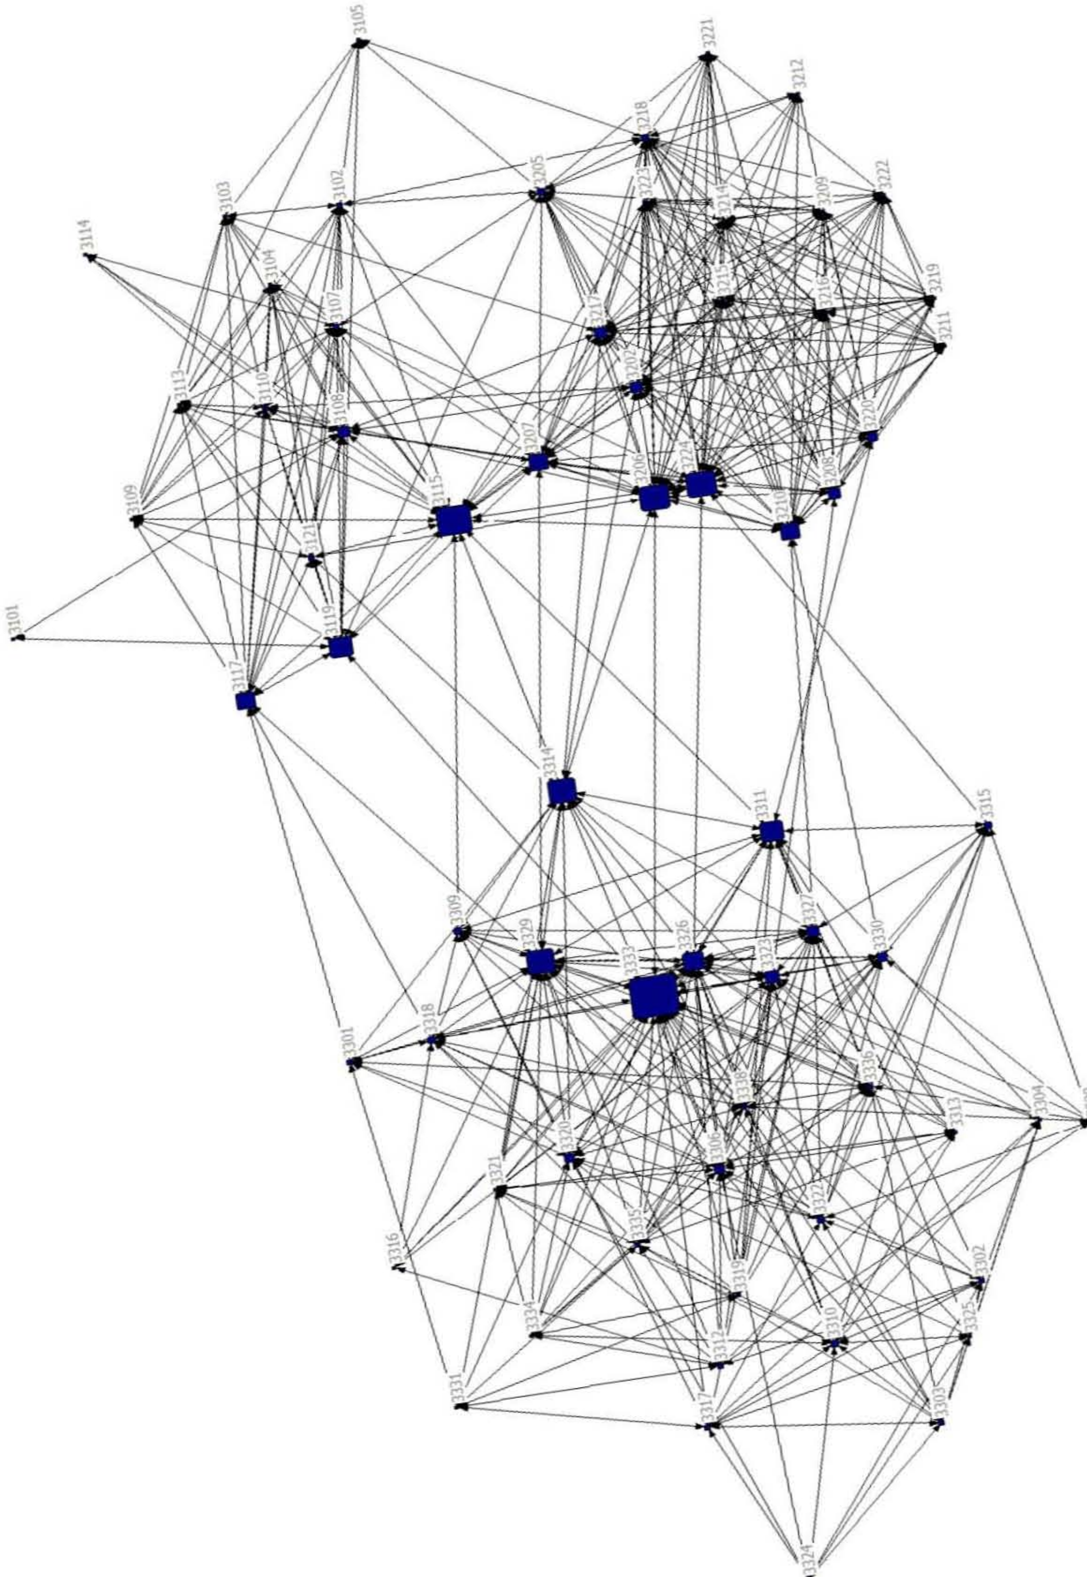

Figure 17: Gauteng: Betweenness

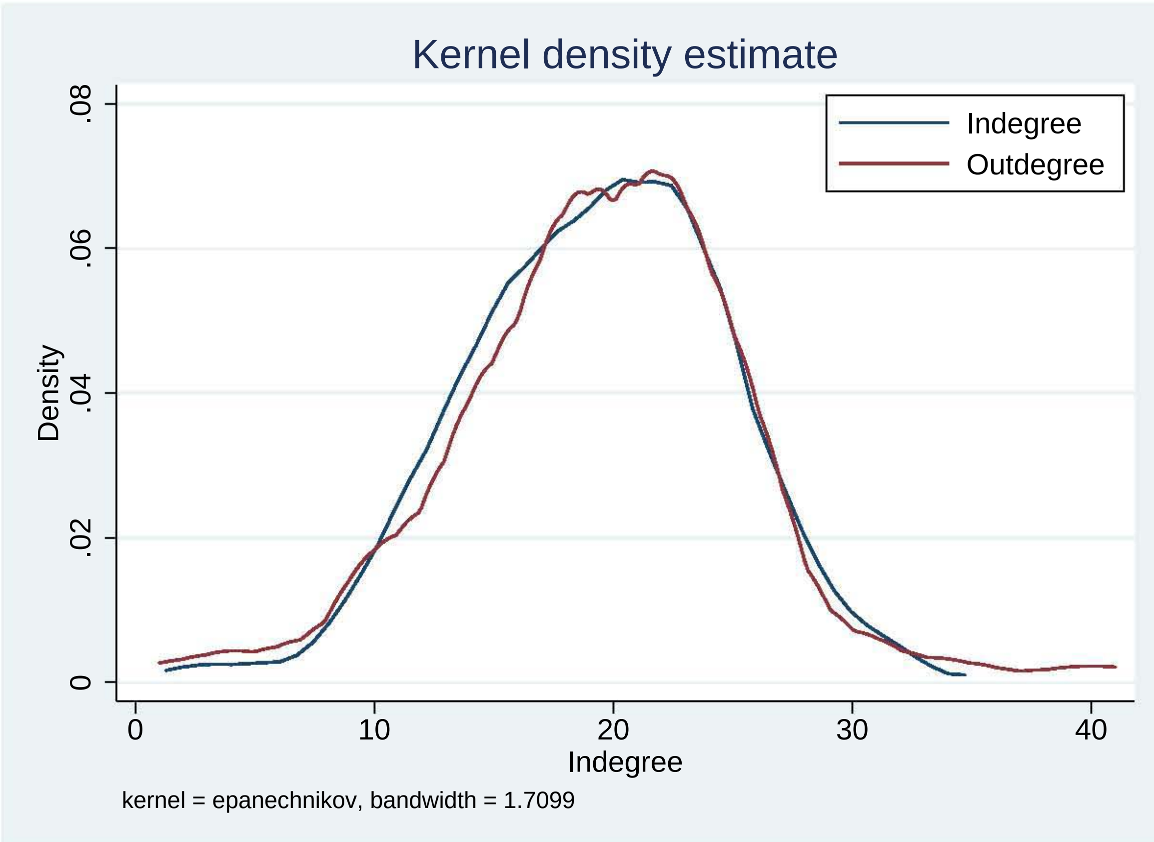

Figure 18: Degree Distribution

1. Kernel density distributions for both the Indegree and Outdegree of participants in the sample.
2. Indegree is the number of individuals who nominated the person as someone with whom they had communicated.
3. Outdegree is the number of individuals that the person nominated as having ever communicated with.

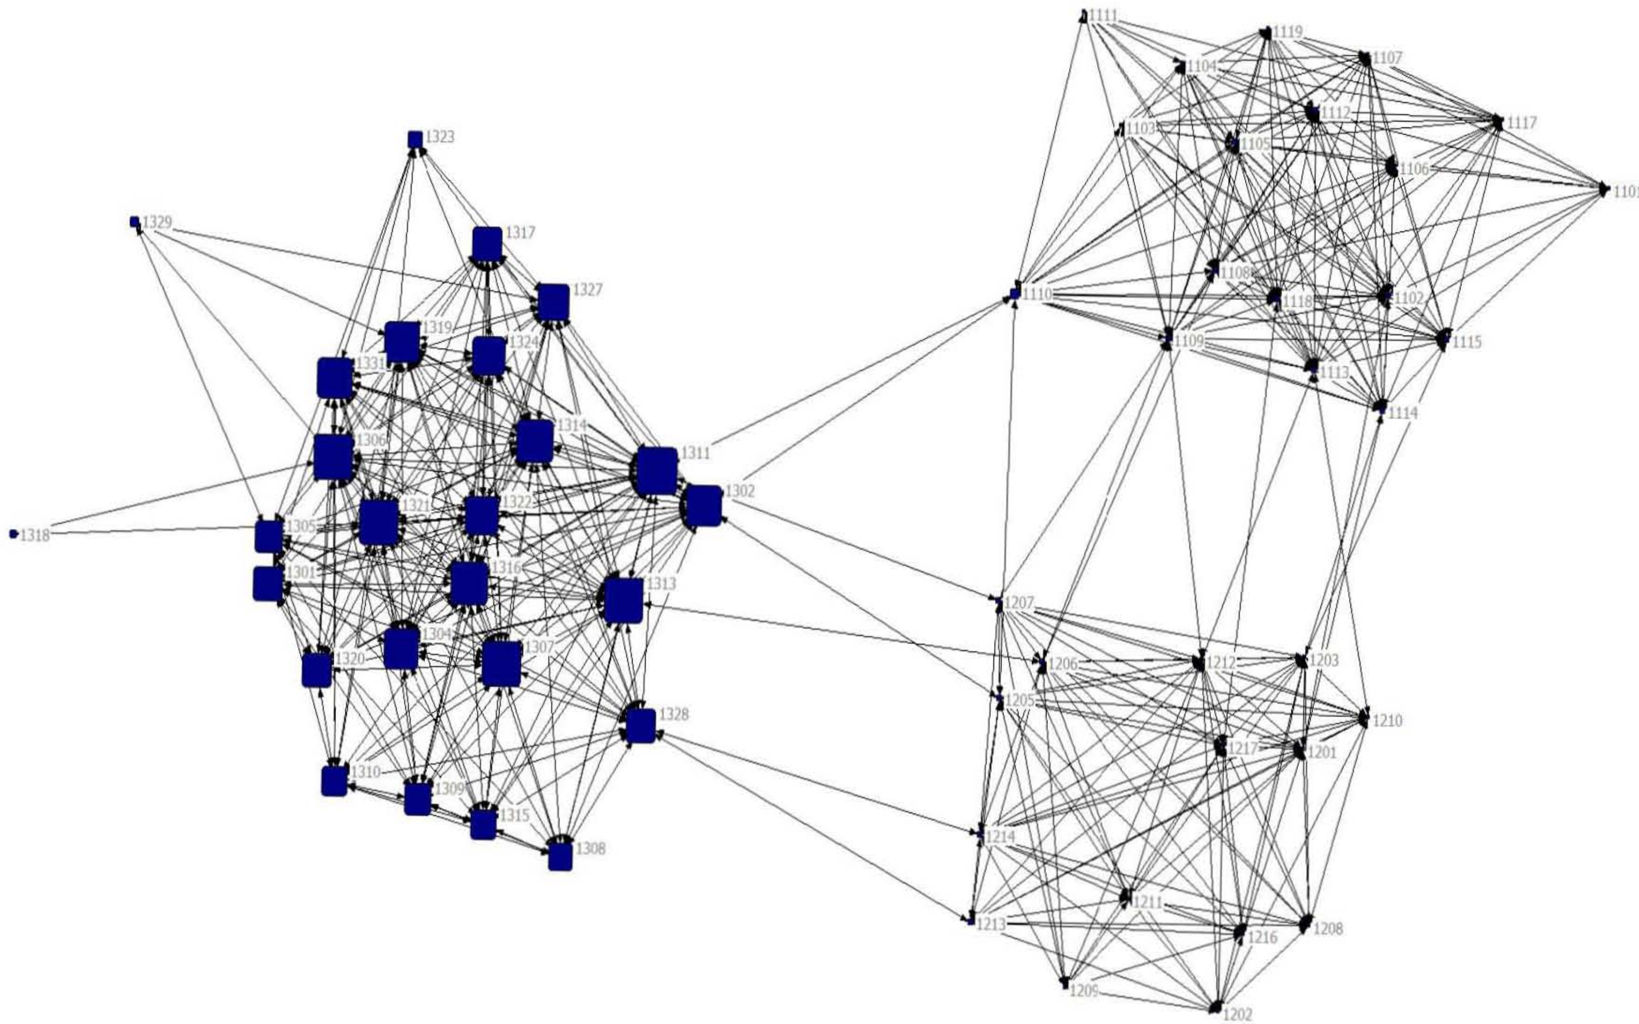

Figure 19: Western Cape: Eigenvector Centrality

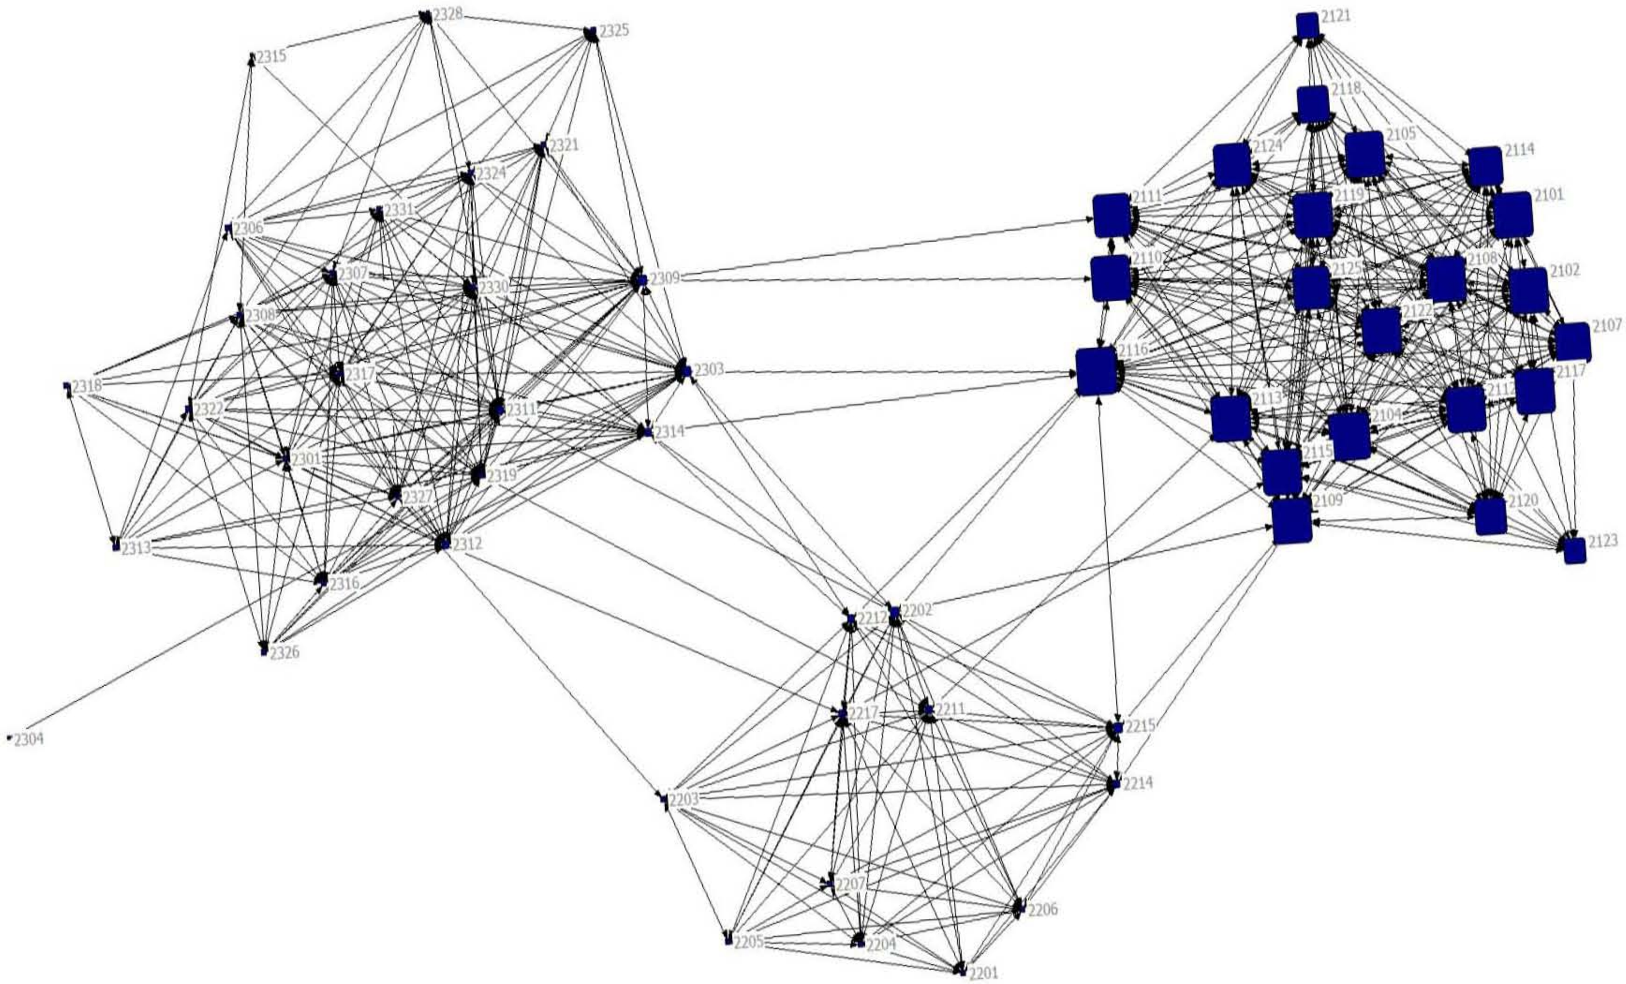

Figure 20: Kwazulu Natal: Eigenvector Centrality

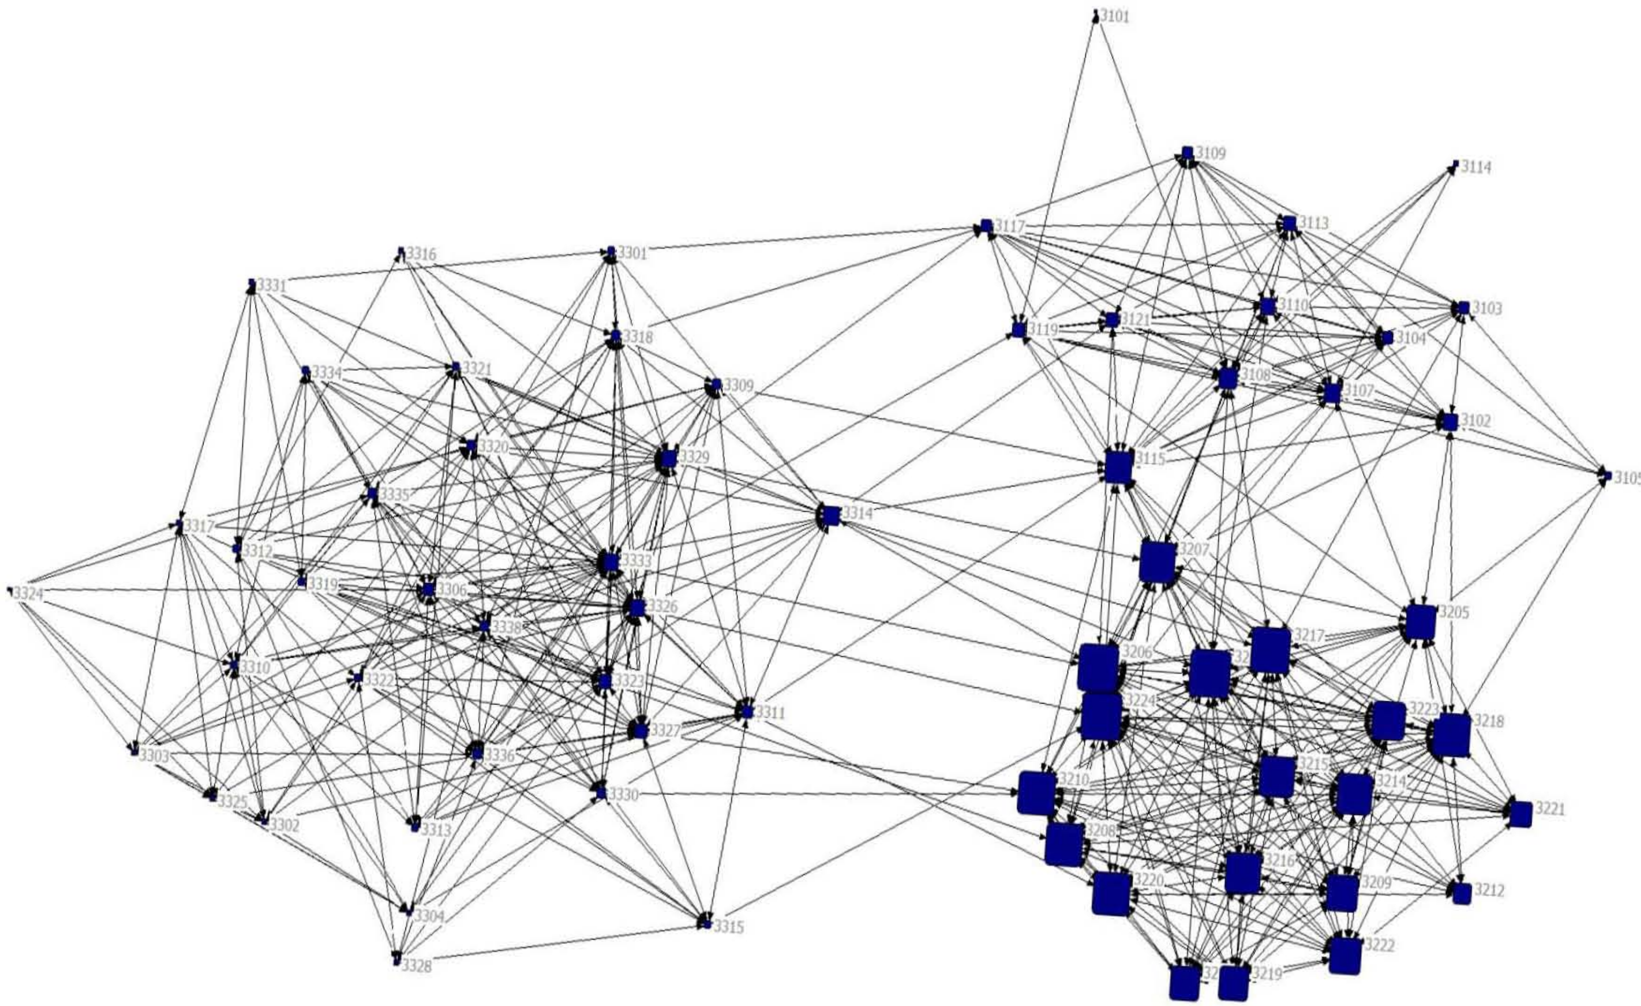

Figure 21: Gauteng: Eigenvector Centrality

References

C. Alexander, M. Piazza, D. Mekos, and T. Valente. Peers, schools, and adolescent cigarette smoking. *Journal of adolescent health*, 29(1):22–30, 2001.

K.J. Arrow. What has economics to say about racial discrimination? *The Journal of Economic Perspectives*, 12(2):91–100, 1998.

V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2003.

O. Bandiera and I. Rasul. Social networks and technology adoption in northern mozambique. *The Economic Journal*, 116(1957):869–902, 2006.

A.V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.

J. Berg, J. Dickhaut, and K. McCabe. Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1):122–142, 1995.

T. Besley and A. Case. Diffusion as a learning process: Evidence from hyv cotton. Technical report, 1994.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, pages 992–1026, 1992.

L. Bottazzi, M. Da Rin, and T.F. Hellmann. The importance of trust for investment: Evidence from venture capital. Technical report, National Bureau of Economic Research, 2011.

S. Bowles, S.N. Durlauf, and K. Hoff. *Poverty Traps*. Princeton University Press, 2011.

Y. Bramoulle, H. Djebbari, and B. Fortin. Identification of peer effects through social networks. *Journal of Econometrics*, 150(1):41–55, May 2009. doi: 10.1016/j.jeconom.2008.12.021.

A. Cassar, L. Crowley, and B. Wydick. The effect of social capital on group loan repayment: evidence from field experiments*. *The Economic Journal*, 117(517):F85–F106, 2007.

C.P. Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, 2003.

A.G. Chandrasekhar, C. Kinnan, and H. Larreguy. Informal insurance, social networks, and savings access: Evidence from a lab experiment in the field. Technical report, MIT Working Paper, 2011.

J. Coleman, E. Katz, and H. Menzel. The diffusion of an innovation among physicians. *Sociometry*, pages 253–270, 1957.

J.S. Coleman. Social capital in the creation of human capital. *American Journal of Sociology*, pages 95–120, 1988.

T.G. Conley and C.R. Udry. Learning about a new technology: Pineapple in ghana. *The American Economic Review*, 100(1):35–69, 2010.

E. Costenbader and T.W. Valente. The stability of centrality measures when networks are sampled. *Social Networks*, 25(4):283–307, 2003.

J.C. Cox. How to identify trust and reciprocity. *Games and Economic Behavior*, 46 (2):260–281, 2004.

P. Dasgupta, S. Goyal, K.G. Maler, R. Putnam, and I. Serageldin. A matter of trust: Social capital and economic development. 2009.

J. De Weerdt. *Risk-sharing and endogenous network formation*. Number 2002/57. WIDER Discussion Papers//World Institute for Development Economics (UNU-WIDER), 2002.

D. Di Cagno and E. Sciubba. Trust, trustworthiness and social networks: Playing a trust game when networks are formed in the lab. *Journal of Economic Behavior and Organization*, 75(2):156–167, 2010.

E. Duflo, M. Kremer, and J. Robinson. Nudging farmers to use fertilizer: theory and experimental evidence from kenya. Technical report, National Bureau of Economic Research, 2009.

S.N. Durlauf. Groups, social influences and inequality: a memberships theory perspective on poverty traps. *Wisconsin Madison-Social Systems Working papers*, 2002.

M. Fafchamps and F. Gubert. The formation of risk sharing networks. *Journal of Development Economics*, 83(2):326–350, July 2007.

A.D. Foster and M.R. Rosenzweig. Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of political Economy*, pages 1176–1209, 1995.

O. Frank. Network sampling and model fitting. *Models and Methods in Social Network Analysis*, pages 31–56, 2005.

A. Gaviria and S. Raphael. School-based peer effects and juvenile behavior. *Review of Economics and Statistics*, 83(2):257–268, 2001.

M. Granovetter. Economic action and social structure: the problem of embeddedness. *American Journal of Sociology*, pages 481–510, 1985.

M.S. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.

H. Hogset and C.B. Barrett. Social learning, social influence, and projection bias: A caution on inferences based on proxy reporting of peer behavior. *Economic Development and Cultural Change*, 58(3):563–589, 2010.

M.O. Jackson and B.W. Rogers. Meeting strangers and friends of friends: How random are social networks? *The American Economic Review*, pages 890–915, 2007.

D. Karlan, M. Mobius, T. Rosenblat, and A. Szeidl. Trust and social collateral. *The Quarterly Journal of Economics*, 124(3):1307–1361, 2009.

H.P. Kohler, J.R. Behrman, and S.C. Watkins. The density of social networks and fertility decisions: Evidence from south nyanza district, kenya. *Demography*, 38(1): 43–58, 2001.

D. Krackhardt. Predicting with networks: Nonparametric multiple regression analysis of dyadic data. *Social Networks*, 10(4):359–381, 1988.

M. Kremer and E. Miguel. The illusion of sustainability. *The Quarterly Journal of Economics*, 122(3):1007–1065, 2007.

S. Leider and M.M. Mobius. Directed altruism and enforced reciprocity in social networks. *The Quarterly Journal of Economics*, 124(4):1815–1851, 2009.

A. Lev-On, A. Chavez, and C. Bicchieri. Group and dyadic communication in trust games. *Rationality and Society*, 22(1):37–54, 2010.

L.S.O. Liverpool-Tasie and A. Winter-Nelson. Social learning and farm technology in ethiopia: Impacts by technology, network type, and poverty status. *The Journal of Development Studies*, 48(10):1505–1521, 2012.

A. Maertens. Who cares what others think (or do)? social learning, social pressures and imitation in cotton farming in india. *Job Market Paper, University of Pittsburg*, 2010.

C.F. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542, 1993.

P.V. Marsden and N.E. Friedkin. Network studies of social influence. *Sociological Methods and Research*, 22(1):127–151, 1993.

A. Mayer and S.L. Puller. The old boy (and girl) network: Social network formation on university campuses. *Journal of Public Economics*, 92(1-2):329–347, February 2008. doi: 10.1016/j.jpubeco.2007.09.001.

K. Mihaly. Too popular for school? friendship formation and academic achievement, 2007.

R.A. Moffitt. Policy interventions, low-level equilibria, and social interactions. *Social dynamics*, pages 45–82, 2001.

C.M. Moser and C.B. Barrett. The disappointing adoption dynamics of a yield-increasing, low external-input technology: the case of sri in madagascar. *Agricultural Systems*, 76(3):1085–1100, 2003.

K. Munshi. Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1):185–213, 2004.

MEJ Newman. The structure and function of complex networks. *Structure*, 2003.

A.J. O'Malley and P.V. Marsden. The analysis of social networks. *Health Services and Outcomes Research Methodology*, 8(4):222–269, 2008.

R.L. Porta, F. Lopez-De-Silane, A. Shleifer, and R.W. Vishny. Trust in large organizations. Technical report, National Bureau of Economic Research, 1996.

R.D. Putnam. *Bowling alone: The collapse and revival of American community*. Simon & Schuster, 2001.

R.D. Putnam, R. Leonardi, and R.Y. Nanetti. *Making democracy work: Civic traditions in modern Italy*. Princeton university press, 1994.

U. Rajan. Trembles in the bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 82(1):248–266, 1998.

P.A. Samuelson. Consumption theory in terms of revealed preference. *Economica*, 15 (60):243–253, 1948.

P. Santos and C.B. Barrett. Identity, interest and information search in a dynamic rural economy. *World Development*, 38(12):1788–1796, December 2010.

P. Sapienza, L. Zingales, and L. Guiso. Does culture affect economic outcomes? Technical report, National Bureau of Economic Research, 2006.

S.B. Thompson. Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics*, 99(1):1–10, 2011.

T.W. Valente. Social network thresholds in the diffusion of innovations. *Social Networks*, 18(1):69–89, 1996.

K. Van den Broeck and S. Dercon. Information flows and social externalities in a tanzanian banana growing village. *The Journal of Development Studies*, 47(2):231–252, 2011.

S. Wasserman and K. Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.

P.J. Zak and S. Knack. Trust and growth. *The Economic Journal*, 111(470):295–321, 2001.