KEYNES FUND FINAL REPORT

Name of Principal Investigator: Matthew Elliott

Project title and Round: Free Riding and Group Cooperation

Report Due Date: Report received: Report approved:

Part 1: Summary of work and results

Social and economic networks play a key role in everyday economic activity. A large literature emphasizes the value of networks. For example, networks play a key role in risk sharing or cooperation without commitment more generally (e.g., Townsend (1993), Binzel et al. (2017), Banerjee et al. (2014a,b)). Networks allow friends to exact social punishment, facilitate transfers, and enable information exchange. These features are particularly important in developing economies, as settings with poor institutions lean on social and economic networks as substitutes (see, Breza (2016)).

One such important setting is worker production. When studying this from a networks perspective, a sizeable literature focuses on and the value of information and even perhaps monitoring of friends by their referees (people who gave referrals) (see, for example, Pallais et al. (2016)).

The scope for network effects can vary by context. Different worker settings may have different natural production designs and therefore vary in the extent and nature by which networks play a role. For instance, in certain settings, individuals may each be incentivized at an individual level such as by being paid at a piece rate. In other settings, production may be team based, and any incentive payments might have to depend only on the overall output of the team creating incentives for team members to free-ride on others' efforts. Indeed, in such settings, even without incentive payments, there may be situations that also suffer from free riding in which team members can help each other out to improve overall productivity. Yet in other settings, there may be a hierarchical structure whereby the returns to effort of some depend on others, but not vice versa. Moreover, the actual production design---the organization of workers and the payment structure they face---is a choice of the firm.

In studying how social relationships may be leveraged to enhance productivity, a crucial piece of the analysis is how we model the nature by which workers solve the game. The three natural concepts are Nash equilibria, perfect collusion facilitated by transfers, and favour trading. The overall output, worker attendance, and efforts conditional on attendance all differ under each of these concepts. Further, the role of network relationships play a different role in each of these cases. One may imagine, for instance, that perfect or partial collusion may be more feasible when workers are friends, whereas Nash may be a better description among strangers. On the other hand, it is possible that irrespective of the network arrangements, one of the collusive behaviours could simply be the best description of behaviour. After all, co-workers interact and can develop relationships.

A common theme in the literature, and under each of these concepts, is that the network externalities are *positive*. That is, when friends are hired and then, say, produce in teams, it should generate at least as much output as if the team is comprised of strangers. For example, it might be expected that strangers can mitigate free-riding problems but only by trading favours, whereby one agent works hard to benefit another worker in exchange for that worker's concurrent hard work which benefits the first worker. This might work well in a team production setting where positive externality generating choices can be reciprocated, but poorly in a hierarchical situation where there is scope for one worker doing favours for another, but not for favours being done in return. However, in this case, friends might permit the equilibrium play to evolve from favour trading to full collusion, possibly supported by transfers or by favour being returned outside of the work environment. Thus friendship helps better internalize the positive externalities and reduce free riding. More broadly, the idea is that friends are good because they can sustain cooperation in this sort of setting, and in other settings they add benefits of monitoring and information about others' quality.

Yet, in fact, network relationships among co-workers could be detrimental to productivity. This has not been emphasized in the literature. For instance, having a team of friends might allow for certain collusive behaviour that have negative effects, such as absconding together or collectively applying no effort. This can be true even in settings where individuals do not want to appear to be shrinking relative to members of their network. In this sort of setting, if a firm produces with team incentives such as aggregate output, producing with friends may lead to reduced output. Similarly, when friends are place in a hierarchy, friends may be less demanding of each other than strangers would be.

In this project we focus on several questions. First, do collections of workers behave in a manner best described by Nash equilibria, perfect collusion, or favour-trading? And how do network relationships then affect the answer to this question? Second, what choices of group incentive design---i.e., individual incentives, team aggregative incentives, hierarchical incentives---ought a firm select and how does this vary with whether workers are strangers or networked? Third, what does that tell us in terms of developing a natural descriptive model of behaviour consistent with the data?

We study these questions through a field experiment with a three group design structures: individual incentives, team aggregate incentives, and hierarchical incentives. In this context we develop theoretical predictions for the Nash equilibria, perfect collusion, and favour trading by looking at how worker productivity, attendance, and total output are predicted to change across these three treatments. Our choice of group design structures both reflect stylized versions of real-world designs but also serve to have diverging predictions under each solution concept and therefore aid in identification. After considering the case of all workers being strangers, we then model workers as mutually linked in a network---i.e., as friends. So we look at the role network relationships plays by studying each of these cases wherein now individuals are mutually friends.

The field experiment is therefore a simple three by two design. We vary both the group design structure (by changing how individuals are compensated for their outputs and their team mates' outputs) as well as whether individuals are jointly producing with strangers or other members of their network neighbourhood (their friends). As the design choices exactly correspond to the theory, the treatment arms allow us to both distinguish between Nash equilibria, perfect collusion, and partial collusion solutions both among strangers and friends.

Every subject is a member of a group of three individuals. The first arm randomizes groups to one of three payment schemes (see Figure 1): (1) individual incentives, whereby individuals receive piece rate for their output; (2) team incentives, whereby all members of their team are payed piece rate for the aggregate output; and (3) hierarchical incentives, whereby we construct a hierarchy of payments where worker A gets paid as a function of B's and C's outputs as well and B gets paid as a function of C's output as well and C gets only an individual incentive. The second arm randomizes whether the three individuals in the group are strangers or whether they are friends.

To implement this, we study a sample of 420 workers who conduct village census surveys.[1] These individuals began by showing up with at least two other (mutual) friends. The pool of individuals were then split up into groups of three. For half the population the individuals were matched to strangers whereas for the remaining half the individuals were matched with two of their friends. The groups were then assigned to one of three payment schemes: (1) individual incentive, (2) team incentive, or (3) hierarchical incentive. Individual incentive simply pays each individual for the number of census enumerations they conducted over the period. Team incentive pays each

---

[1] The individuals were hired for a temporary job in between searching for usual work, and are typically enumerators in the labor market. They were told that they were not going to be hired by this organization at any point in the future.

individual the same amount as a function of the team aggregate production. Finally, hierarchical incentives pay A as a function of B and C's productivities, B as a function of C's productivity, and C individually.

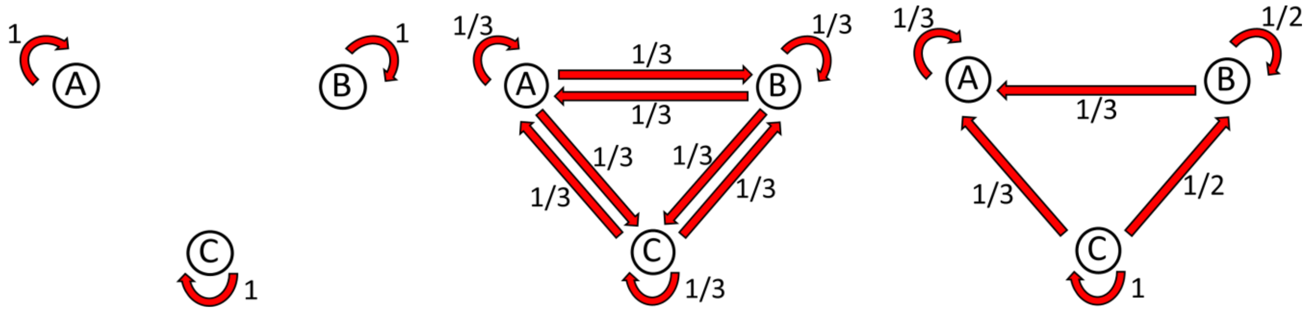

Figure 1: The three payment treatments. A red arrow from B to A with weight x, represents that A's incentive payment depends on B's output with weight x.

Our main outcomes are three-fold. First, we look at total production, which is the key outcome variable from the perspective of the firm. This is just the total number of census enumeration surveys done over the period. Second, we look at attendance. Workers can choose not to attend one or even several or all days. The theory predicts some attrition in equilibrium, differential by group incentive design and varying by network relationships. Third, we look at production conditional on attendance. While this is a selected sample and should be considered with the obvious caution (it is non-causal), cutting the data this way helps us study how effort of those that show up varies and whether that is consistent with the theory. We also instrument for attendance by using the distance from workers' homes to the village they were assigned to survey.

We develop the theoretical predictions for our three theories by first considering effort choices conditioning on all three team members turning up. In this case we get the following orderings:

Output conditional on attendance:

- Nash equilibrium: Individual = Hierarchy C > Hierarchy B > Hierarchy A = Team
- Favour Trading: Individual = Team = Hierarchy C > Hierarchy B > Hierarchy A.
- Full Collusion: Hierarchy C > Individual = Team > Hierarchy B > Hierarchy C.

In the experiment the team were sent home if any one team member did not turn up. We thus model attendance by setting an agents' expected payoff from attending equal to their expected payoff conditional on all agents attending weighted by the probability that both other team members attend. This creates a system of equations with multiple fixed points. We study the largest fixed point.

Although we don't have functional forms to work with, we are still able to partially order the theoretical attendance predictions across markets.

The Nash predictions on attendance are:

- Individual > Team
- Hierarchy A, Hierarchy B, Hierarchy C > Team.
- Hierarchy B > Hierarchy C.

The prefect collusion predictions on attendance are:

- Hierarchy A=Hierarchy B=Hierarchy C > Individual > Team

The favour trading predictions on attendance are:

- Individual = Team

Predictions on overall output are then found by considering which attendance predictions (partial orderings) go in the same direction as the output predictions.

In comparison to these theoretical predictions the results from our experiment are as follows (see the regression analysis reported in Table 1 in the appendix):

Overall output

- For Friends, output is lower for the team treatment than any other (for all pairwise comparisons).
- For Strangers, output is higher for team than for hierarchy.

Attendance:

- For Friends, attendance is lower for the team treatment than any other (for all pairwise comparisons).
- For Strangers, attendance is higher for the team treatment than any other (for all pairwise comparisons).
- For the team treatment, friends attend less often than strangers. On the other hand, friends attend at a higher rate for all other treatments.

Output conditional on attendance (instrumenting):

- For Friends, output given attendance is lower for the team treatment than for individual and hierarchy position B.
- For Strangers, output given attendance is highest for individual in all pairwise comparisons.
- Friends produce less in the individual treatment than strangers, but more in the team treatment compared to strangers.

The above summarises all the statistically significant relationships we find. When a pairwise comparison is not made there is no statistical evidence to reject that the outcomes variable was the same for the two treatments.

If we focus first on just the friend treatments, then the results are broadly consistent with Nash equilibrium play. The Nash equilibrium predicts that attendance should be lowest for team, as is found. It also predicts that overall output should be lower for the team treatment that for any other, as is also found. Finally it predicts that output conditional on attendance should be lower for the team treatment than for the individual treatment and the Hierarchy treatments A and B. Again our findings broadly support this. We therefore conclude that the Nash equilibrium predictions are consistent with the friends treatment. It is straightforward to show that, in comparison, predictions made by both the favour trading and prefect collusion theories are violated.

The results for the stranger treatments are harder to reconcile with the theory. All the theories predict that attendance is weakly lower for the team treatment than for any other treatment, while for strangers we find that attendance is higher for the team treatment than for any other treatment.

The comparison between friend and stranger treatments is also a bit puzzling. For example, it is odd that in the team treatment where the group have to work to together and their payoffs are highly interdependent, attendance is higher for strangers than for friends, while for all other treatments attendance by friends is higher. It is also then a little strange that friends produce more than strangers in the team treatment, but less than strangers in the individual treatment.

To help us dig a little deeper into our results, and better understand what is going on, we conducted an end line survey in which we asked participants about their perceived effort levels of others, and about the use of transfers within their team (see Tables 2 and 3 in the appendix). Given our other results, these too were surprising. We found that transfers are used more by friends in the team treatment than in the individual treatment, and there is also some evidence they are used more by friends in the team treatment than by strangers. Such a use of transfers by friends is at odds with Nash theory that otherwise seems to fit well. There is also some evidence that friends rate the productivity of their team mates lower. In Hierarchy A, B and C treatments, friends rate the effort of others team members lower on average than strangers do while there were no statistically significant effects in the other treatments.

Our findings suggest forces are at play in our setting that are inconsistent with the vast majority of the previous literature. Indeed, as well as not being able to explain our results with the theories we have considered, we also cannot explain our findings with inequality aversion, altruism or a standard signalling model.

Perhaps the best explanation we have at present for our findings is a social capital accumulation story. Strangers have the opportunity to form new social capital, to form new productive relationships, and so in treatments that require them to work together the possibility of doing this incentivises participation. At the same time, friends risk falling out and loosing social capital when they work together in an environment where their pay depends on their friends' efforts. This can explain friends being more reluctant to participate in the team treatment than the individual treatment, while strangers are keener to participate in the team treatment than the individual treatment. Of course, this would not make sense if failing to show up imposes a bigger loss on friends in team treatment than showing up and being perceived to not work as hard as you should. However, it seems reasonable to suppose that friends are able to coordinate their decisions to not show up, mitigating such risk.

Our results were not what we were expecting, but we think they are interesting and believe that they provide some useful policy guidance, or at least words of warning, while also advancing what is known about the main questions we set out to address. They suggest that leveraging social relationships to overcome free-riding problem in team production environments will be difficult. While social relationships play an important role in risk-sharing for the same population and mitigate free riding problems in this setting, it is not straightforward to transplant this into a production / working environment. We do not find that friends are systematically more productive than strangers in team based environments. Conditional on showing up there is some evidence they work harder for each other, but friends are also less willing to participate in such settings. This suggests that while it may be possible to coax higher effort out of friends, higher average wages will also be required to incentivise attendance. Further, in environments where output is not team based but mainly individual, so in situations where piece rates may be appropriate, the situation is reversed. Our evidence suggests that strangers typically work harder, perhaps because they are less distracted, but at the same time they are less inclined to participate and so can be expected to require higher wages to take such jobs. Finally, we find that the standard theories are not able to provide a coherent picture of workers' participation and effort choices. It seems that the situation is broader than is captured by these theories. We advance a social capital accumulation story that can help partially reconcile our observations, but there may be other equally appealing stories.

Part 2: Impact and outputs

We are only just finished with the data analysis part of the project. We have not yet written a completed draft that is ready for public consumption. When we do, we will make it available through several open access repositories, including the Cambridge working paper series. We'll then publish the paper in an economics journal and are confident we will publish it in a top field or general interest journal.

Dissemination of the work has only just started. So far we have only presented the paper once, in an internal Keynes fund workshop. We will continue to disseminate the work widely, both in the US and Europe through seminars and workshops.

The project involved a team of principal investigators, Professor Attila Ambrus from Duke, Professor Arun Chandrasekhar from Stanford and Dr. Matthew Elliott from Cambridge. Dr. Elliott has separate ongoing work with both Professor Ambrus and separately with Professor Chandrasekhar.

The current project was also enabled by a collaboration with a team in India who helped pilot and then implement the experiment. This was organized through IFMR in India and the Abdul Latif Jameel Poverty Action Lab. The ongoing work involving Professor Chandrasekhar and Dr. Elliott is also with IFMR and the Abdul Latif Jameel Poverty Action Lab.

Part 3: Any possible future plans

Our next immediate steps are to finish writing the paper, disseminate the work and get the paper published well. The new project being undertaken by Professor Chandrasekhar and Dr. Elliott, as referred to in the previous section, builds on our work in this project by studying inter business relationships as opposed to intra business relationships. The focus in that project is on business relationships between producers and their suppliers. The idea of this work is to better understand the functions that these relationships perform, whether markets that are intermediated by relationships work well and whether there are high returns to increasing social capital in these contexts. Our results in the current project helps to motivate our new work. First, this current project suggests that business relationships might play a more nuanced role than theory suggests. Second, our social capital accumulation story is consistent with investments in business relationships being important.

References

Banerjee, A., A. Chandrasekhar, E. Duflo and M. Jackson (2014a): "Exposure to formal loans: Changes in social networks," MIT and Stanford working paper.

Banerjee, A., E. Breza, E. Duflo and C. Kinnan (2014b): "Does microfinance foster business growth? The importance of entrepreneurial heterogeneity," mimeo MIT.

Binzel, C., E. Field and R. Pande (2017): "Does the Arrival of a Formal Financial Institution Alter Informal Sharing Arrangements? Experimental Evidence from Village India," mimeo Harvard University.

Breza, E. (2016): "Field experiments, social networks, and development." In *The Oxford handbook of the economics of networks*. Oxford University Press.

Pallais, Amanda, and Emily Glassberg Sands (2016): "Why the referential treatment? Evidence from field experiments on referrals." *Journal of Political Economy* 124.6, 1793-1828.

Appendix A: Regression Tables

Table 1: Daily Productivity and Attendance

VARIABLES	(1) Attendance	(2) Daily Productivity	(3) Daily Productivity Heckman
Team	0.089	0.284	-2.329
	(0.042)	(1.065)	(0.645)
	[0.033]	[0.790]	[0.000]
H.A	-0.041	-2.355	-1.666
	(0.043)	(1.117)	(0.674)
	[0.339]	[0.035]	[0.013]
H.B	-0.037	-2.514	-2.039
	(0.043)	(1.111)	(0.675)
	[0.385]	[0.024]	[0.003]
H.C	-0.028	-1.805	-1.403
	(0.044)	(1.156)	(0.679)
	[0.527]	[0.119]	[0.039]
Friend Treatment	0.119	1.484	-1.672
	(0.036)	(0.969)	(0.549)
	[0.001]	[0.126]	[0.002]
Friend Treatment × Team	-0.183	-3.147	1.429
	(0.051)	(1.317)	(0.809)
	[0.000]	[0.017]	[0.077]
Friend Treatment × H.A	0.107	3.894	1.504
	(0.056)	(1.467)	(0.900)
	[0.055]	[0.008]	[0.095]
Friend Treatment × H.B	0.095	4.006	2.183
	(0.055)	(1.445)	(0.901)
	[0.085]	[0.006]	[0.015]
Friend Treatment × H.C	0.101	3.264	1.136
	(0.055)	(1.498)	(0.902)
	[0.068]	[0.029]	[0.208]
Observations	2,095	2,095	2,065
Stranger, Indiv Mean	0.713	16.646	23.363
Team + Friend × Team p-val	0.001	0.000	0.066
Friend + Friend × Team p-val	0.075	0.057	0.685
H.A + Friend × H.A p-val	0.049	0.095	0.782
Friend + Friend × H.A p-val	0.000	0.000	0.814
H.B + Friend × H.B p-val	0.084	0.100	0.806
Friend + Friend × H.B p-val	0.000	0.000	0.476
H.C + Friend × H.C p-val	0.028	0.123	0.650
Friend + Friend × H.C p-val	0.000	0.000	0.457

NOTES:
In Col2 we replace productivity with 0 for those who did not show-up.
In Col3 we run a heckman two-step model.
Across all the regressions we control for day, phase and degree fixed effects.

Table 2: Transfers

VARIABLES	(1) Transfers	(2) Transfers_heckit	(3) Transfers	(4) Transfers_heckit
Team	-0.071	-0.075	30.411	30.300
	(0.098)	(0.111)	(51.725)	(181.925)
	[0.467]	[0.500]	[0.557]	[0.868]
H.A	-0.082	-0.079	20.723	22.675
	(0.098)	(0.117)	(64.267)	(192.337)
	[0.404]	[0.502]	[0.747]	[0.906]
H.B	-0.037	-0.023	21.974	30.530
	(0.095)	(0.118)	(59.014)	(193.272)
	[0.694]	[0.848]	[0.710]	[0.874]
H.C	-0.027	-0.018	-46.151	-43.993
	(0.110)	(0.121)	(88.729)	(198.378)
	[0.807]	[0.882]	[0.603]	[0.824]
Friend Treatment	0.146	0.142	208.351	205.437
	(0.099)	(0.097)	(115.676)	(158.721)
	[0.141]	[0.141]	[0.073]	[0.196]
Friend Treatment × Team	0.176	0.187	205.403	209.913
	(0.137)	(0.141)	(252.185)	(231.109)
	[0.200]	[0.184]	[0.416]	[0.364]
Friend Treatment × H.A	-0.076	-0.059	-115.201	-107.462
	(0.151)	(0.157)	(137.038)	(257.171)
	[0.614]	[0.704]	[0.401]	[0.676]
Friend Treatment × H.B	0.026	0.034	-32.488	-27.707
	(0.159)	(0.158)	(125.052)	(259.377)
	[0.873]	[0.830]	[0.795]	[0.915]
Friend Treatment × H.C	0.042	0.059	26.181	37.012
	(0.166)	(0.161)	(111.209)	(263.555)
	[0.799]	[0.713]	[0.814]	[0.888]
Observations	307	413	307	413
Stranger, Indiv Mean	0.212	0.212	29.697	29.697
Team + Friend × Team p-val	0.291	0.195	0.307	0.091
Friend + Friend × Team p-val	0.001	0.001	0.091	0.014
H.A + Friend × H.A p-val	0.177	0.177	0.318	0.614
Friend + Friend × H.A p-val	0.577	0.508	0.147	0.633
H.B + Friend × H.B p-val	0.925	0.913	0.916	0.987
Friend + Friend × H.B p-val	0.200	0.165	0.035	0.394
H.C + Friend × H.C p-val	0.902	0.692	0.836	0.967
Friend + Friend × H.C p-val	0.197	0.121	0.036	0.255

NOTES:
In col1 and col2 the dependent variable is weather a transfer was made or not.
In col3 and col4, the dependent variable is amount transfered
In col2 and col4 we use heckman two-step model.
Across all the regressions we control for phase and degree fixed effects.

Table 3: Effort

VARIABLES	(1) Effort Rating	(2) Effort Rating Heckman
Team	-0.182	-0.198
	(0.244)	(0.315)
	[0.457]	[0.531]
H.A	-0.294	-0.297
	(0.351)	(0.334)
	[0.402]	[0.373]
H.B	-0.442	-0.455
	(0.420)	(0.335)
	[0.293]	[0.174]
H.C	-0.477	-0.456
	(0.346)	(0.344)
	[0.170]	[0.185]
Friend Treatment	-0.018	0.015
	(0.253)	(0.274)
	[0.944]	[0.958]
Friend Treatment × Team	0.316	0.307
	(0.362)	(0.400)
	[0.383]	[0.444]
Friend Treatment × H.A	-0.436	-0.353
	(0.454)	(0.445)
	[0.338]	[0.428]
Friend Treatment × H.B	-0.220	-0.255
	(0.498)	(0.449)
	[0.659]	[0.570]
Friend Treatment × H.C	-0.214	-0.238
	(0.449)	(0.457)
	[0.634]	[0.602]
Observations	307	413
Stranger, Indiv Mean	8.788	8.788
Team + Friend × Team p-val	0.608	0.658
Friend + Friend × Team p-val	0.270	0.275
H.A + Friend × H.A p-val	0.010	0.025
Friend + Friend × H.A p-val	0.233	0.341
H.B + Friend × H.B p-val	0.017	0.016
Friend + Friend × H.B p-val	0.606	0.506
H.C + Friend × H.C p-val	0.014	0.019
Friend + Friend × H.C p-val	0.537	0.545

NOTES:
The dependent variable is the mean effort rating
In col2 we use heckman two-step model
Across all the regressions we control for phase and degree fixed effects.