Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field

Reshmaan Hussam (Yale), Natalia Rigol (MIT), and Benjamin Roth (MIT)*

June 24, 2017

Abstract

The impacts of cash grants and access to credit are known to vary widely, but progress on targeting these services to high-ability, reliable entrepreneurs is so far limited. This paper reports on a field experiment in Maharashtra, India that assesses (1) whether community members have information about one another that can be used to identify high-ability microentrepreneurs, (2) whether organic incentives for community members to misreport their information obscure its value, and (3) whether simple techniques from mechanism design can be used to realign incentives for truthful reporting. We asked 1,380 respondents to rank their entrepreneur peers on various metrics of business profitability and growth and entrepreneur characteristics. We also randomly distributed cash grants of about $100 to measure their marginal return to capital.

We find that the information provided by community members is predictive of many key business and household characteristics including marginal return to capital. While on average the marginal return to capital is modest, preliminary estimates suggest that entrepreneurs given a community rank one standard deviation above the mean enjoy an 8.8% monthly marginal return to capital and those ranked two standard deviations above the mean enjoy a 13.9% monthly return. When respondents are told their reports influence the distribution of grants, we find a considerable degree of misreporting in favor of family members and close friends, which substantially diminishes the value of reports. Finally, we find that monetary incentives for accuracy, eliciting reports in public, and cross-reporting techniques motivated by implementation theory all significantly improve the accuracy of reports.

*We would like to first thank our team for their tireless work on this project, especially our research manager Sitaram Mukherjee, our project assistant Prasenjit Samanta, and our field manager Sayan Bhattacharjee. We are also very grateful to Rohan Parakh for excellent advice, research assistance, and field management. This research was possible with funding from the Asian Development Bank, PEDL, and Schultz Fund at MIT. We received valuable feedback about this project from Rohini Pande, Benjamin Olken, Abhijit Banerjee, Esther Duflo, Arielle Bernhardt and the participants of the MIT development lunch.

1 Introduction

Identifying high-ability microentrepreneurs is essential to generating sustainable economic growth among the poor. Empirical studies from Sri Lanka, Ghana, and Mexico in which microenterprises are randomly assigned to receive small cash grants find average marginal returns to capital between 4.5% and 25% per month (De Mel, McKenzie, and Woodruff (2008), Fafchamps, McKenzie, Quinn and Woodruff (2014), McKenzie and Woodruff (2008)). But these estimates of average marginal returns mask substantial heterogeneity amongst firms. For instance, while the average marginal return to capital among male business owners in Sri Lanka was 4.5% per month, the quantile treatment effect ranges from 0% - 32.5% per month. However, the authors are unable to predict marginal returns to capital on the basis of the observable information they collected. Thus the problem of identifying high-growth potential microentrepreneurs remains of first order importance for institutions interested in providing capital to microfirms.

In the absence of asymmetric information, finding and empowering potential high-growth entrepreneurs would be easy. The presence of screening costs, however, poses a serious barrier to providing finance (Stiglitz and Weiss, 1981). Though credit registries for small scale borrowers are emerging in many parts of the developing world, they are still very much in their infancy, and asymmetric information continues to plague the informal lending industry in these areas. As such, there is a serious need for tools that can overcome hidden information problems and accurately target capital to firms that can use it well. Low-cost, reliable tools to discern high-growth potential firms may allow financial institutions to extend larger or more flexible loan contracts to these entrepreneurs, fostering economic development while maintaining profitability (Field, Pande, Papp, and Rigol 2013). Additionally, as evidence accumulates about the substantial positive impacts of unconditional transfers, there is rising interest among governments and non-profits in disbursing grants to poor households (e.g. Haushofer and Shapiro 2015; Blattman, Fiala, and Martinez 2013; De Mel, McKenzie, and Woodruff 2008; Fafchamps, McKenzie, Quinn and Woodruff 2014; McKenzie and Woodruff 2008). Enabling targeting on the basis of micro-entrepreneurs' ability might increase the viability and development impact of these initiatives.

- Targeting entrepreneurs is hard for a few reasons: we don't have the necessary data collected reliably. Even if we did have the necessary data, predicting who will be a good entrepreneur with that data is not straightforward. Unlike targeting assistance to the poor, where the criterion is straightforward (you have to be poor), the idea of an entrepreneur is an amorphous concept.

What progress has been made thus far in targeting high-growth entrepreneurs? Existing research

focuses on developing metrics for assessment of entrepreneurs' individual characteristics. The earliest literature uses family characteristics and personal background; De Mel, McKenzie, and Woodruff (2008) use tests of cognitive ability and personality which evaluate focus, competitiveness, etc.; Khwaja and Klinger (2013) develop a psychometric test that analyses ethics and character, intelligence, attitudes and beliefs, and business skills. Yet, as De Mel, McKenzie, and Woodruff (2008) recognize, most measures of entrepreneurship have been developed using high-income country entrepreneurs who may face vastly different constraints from their developing country counterparts. Such tests might thus be a poor metric of what makes a successful entrepreneur in the developing world.

We report on experimental evidence that suggests community information can be used to target high growth entrepreneurs. We pursue this case along two dimensions. First, we document that community members have valuable information regarding the entrepreneurial characteristics of their peers. Second, we establish that, while the natural inclination of community members is to distort their reports in favor of their family and friends, an array of simple techniques from mechanism design theory are effective in realigning community incentives towards truthfulness, substantially increasing the value of reports.

Specifically, we asked community members to rank their peers on various metrics of business growth potential (such as profits and marginal return to capital), borrower reliability (probability of a late payment or default), and other key firm characteristics. We then randomly allocated grants to some of these firms to induce growth. At the time of eliciting rankings we varied whether respondents were told their reports would influence the likelihood that their peers would receive grants. This induced an organic incentive to shift reports in favor of family and friends. We then cross randomized various mechanisms to encourage truth-telling: paying respondents for the accuracy of reports using peer prediction mechanisms, eliciting reports in public or private, and using cross-reporting techniques to infer which peers each community member was most likely to favor.

We find that respondents have valuable information about one another and that peer reports contain residual information that is not predicted by observable household and business-level characteristics. Using baseline data, we find respondents are able to predict one anothers' household income, value of household assets, business profits, work hours, medical expenses, and score on a digit span memory test. On average, a 1 percentile increase in the average rank is associated with a 0.23 percentile increase in the outcome variable. Other than work hours, respondents continue to predict these outcomes even after controlling for a plethora of easily verifiable and also harder-to-observe demographic, household, and business characteristics. Thus ranks contain valuable residual information about community members

beyond that which a loan officer may be able to observe.

More surprisingly, community members were also able to predict the impact of cash grants. Using follow up data on business profits for those who randomly did and did not receive our grants, we find that community members are well able to predict which of their peers enjoy high marginal return to capital. Though on average the impact of our grant on business profits was modest, those who were ranked one standard deviation above the mean enjoyed an 8% monthly marginal return to capital, and those who were ranked two standard deviations above the mean enjoyed a 14% monthly marginal return to capital. The value of community rankings is not diminished by the inclusion of a variety of respondent and business characteristics.

However, we also find that community members distort their reports when they are told that their information will influence the distribution of grants. The correlation between community reports and true outcomes is on average 40% lower, implying a significant amount of manipulation in reports and severely limiting their usefulness. We see manipulations directly when examining how peers rank themselves, their family members, and people who are identified by group members as a respondent's close friends. Although respondents manipulate their reports, we find that our interventions can significantly improve their accuracy. Giving respondents monetary incentives when they report in private increases the accuracy of reports by 75%. Asking respondents to report in public increases the accuracy of reports by 89%. As with manipulations, we also document direct evidence that incentives and public reports make it less likely that respondents favor themselves, their family members, or their close friends in the group.

Our primary motivation is to identify entrepreneurs with the highest growth potential (as measured by marginal returns), since allocating grants to these individuals may be efficient. However institutions targeting the poor may have a variety of welfare priorities. We demonstrate that community information can be used to allocate resources using other types of selection criteria such as wealth, health burden, and average business profitability. By additionally asking community members how they would prefer to allocate the grants, we also estimate how communities themselves trade-off efficiency and redistribution.

This project builds on work done by Rigol and Roth (2016). In a lab in the field experiment, we tested the effectiveness of two different types of monetary incentive schemes to elicit truthful reporting: one that relies on ex-post verification of outcomes and one *peer prediction* incentive scheme that relies on a correlation in community reports to pay respondents based only on their own reports and the contemporaneous reports of their peers. The lab experiment took place in a controlled setting in which farmers were presented with a simple and transparent trade off - to lie or to tell the truth - when

asked to report on their farmer peers in the same village. Reassuringly, we find that the two rules are equally effective in eliciting truthful responses. However, relative to a payment rule that relies on ex-post verification of outcomes, peer prediction schemes are much easier to implement for a variety of reasons we discuss below. So, in the present experiment we paid our respondents using the Robust Bayesian Truth Serum of Witkowski and Parkes (2011), a peer prediction scheme identified to be particularly effective in our lab experiment. To our knowledge, this is the first large-scale field experiment to use a peer prediction mechanism to incentivize respondents in a developing country setting and we are among the first to use peer elicitation outside the lab.[1]

Our work adds to the literature that studies the use of community reports for targeting. Alatas et. al. (2012) explore the value of community information in targeting the poor in Indonesia. The authors ask villagers to select the village's poorest members who are then eligible for a limited number of government transfers. They find that communities have information about the consumption of their members over and above the information contained in the standard proxy means tests. Relative to theirs, our project furthers the understanding of community targeting along several directions. First, we quantify the value of community information over a much broader set of characteristics. Alatas et. al. (2012) measured community knowledge of wealth, which may be easily observable to community members and hence an easy test case. So our ability to target high growth potential entrepreneurs along arguably harder to observe metrics is an encouraging step. Second, we identify that incentives to report truthfully are of first order importance in the elicitation of community information. While Alatas et. al. do examine whether elite capture poses a problem for community reporting, we take a substantially closer look by manipulating respondents' incentives to misreport their information and identifying the likely beneficiaries of each respondent's manipulations. We conclude that there is a substantially larger amount of misreporting than Alatas et. al. are able to identify. Finally we identify a variety of mechanisms effective in encouraging truth-telling, substantially improving the informativeness of community reports.

Closely related is Bryan, Karlan, and Zinman (2014), who randomly give existing borrowers incentives for referring a new borrower and for the repayment outcomes of their referral. They find strong peer enforcement effects but no evidence of peer selection effects in excess of the screening that the bank already does. Their interpretation is that the original clients do not have useful information regarding the individuals they refer beyond the information available to the bank. In contrast to our study, theirs is conducted in South Africa, which has a well-functioning credit bureau. In settings without a well-

[1]The notable exception is John, Loewenstein, and Prelec (2012) who use BTS to measure the prevalence of questionable research practices among academic psychologists.

functioning credit bureau, one might expect community information to be of considerably more value.

The rest of the paper proceeds as follows. Section 2 describes our experimental setting and design, Section 3 describes the data, Section 4 discusses our key specifications, Section 5 discusses our results, and Section 6 concludes. All tables are relegated to the appendix.

2 Description of the Experiment

2.1 Sampling and Study Population

Our intervention took place in Amravati, a city of about 550,000 persons in the state of Maharashtra, India. Within Amravati, we selected nine neighborhoods that each had an abundance of microenterpreneurs.[2] In September 2015, we conducted a complete door-to-door census of these neighborhoods, which encompassed 5,573 households. Each person in the household who was engaged in self-employment activities responded to our census. Every respondent reported the total value of their enterprise's durable assets and inventory (excluding the value of land and buildings), their total number of permanent employees, and their business sector. We identified a sample of 1,576 households that had at least one enterprise with (1) USD 1,000 or less in total working and durable capital and (2) no paid, permanent employees. We excluded farmers and self-employed service persons, such as domestic helpers and teachers. Our selection follows the same criteria as recent "cash-drop" experiments (see e.g. De Mel, McKenzie, and Woodruff 2008).[3]

In October 2015, we recruited these households to participate in our study. At the time of recruitment, households were informed that we would be conducting a project to study entrepreneurship and business growth. Ultimately, 1,345 households agreed to participate in the study. We conducted baseline surveys of these households in December 2015 - April 2016. These 1,345 households were then organized into groups of five based on geographic proximity.[4] We created a total of 274 groups across all neighborhoods. After all respondents in a neighborhood had completed the baseline survey, the groups in that neighborhood were invited to participate in an activity at the community hall, where at the end of the activity, a lottery would be conducted and either 1 or 2 members of the group would receive a USD 100 grant. At the hall, respondents completed the ranking activity (described in section 2.2) and, upon its completion, a public

[2]The neighborhoods are: Belpura, Vilash Nagar, Mahajan Pura, Akoli, New Saturna, Old Saturna, Wadali, and Pathan Chawk.

[3]If there were multiple business owners in the household, we selected households that had less than USD 2000 in combined business capital.

[4]We organized respondents into groups that would minimize the geographic distance between study households. The total number of respondents per neighborhood was not always a multiple of 5, so some groups had 4 or 6 clients.

lottery was held and the winners of the USD 100 grant were announced.

2.2 Experimental Design

The aim of this study is to investigate whether knowledge that community members hold about one another can be useful for making decisions regarding the allocation of business grants, we ask a number of related questions. First, do community members have accurate knowledge regarding their entrepreneur peer's business growth potential? To address this question, we instructed community members rank their peers across several personal and enterprise characteristics that are predictive of business growth (such as profits and marginal returns to capital) and borrower reliability (probability of a late payment or default). We compare community reports to business and household outcomes elicited in a baseline household survey to determine the predictiveness of these measures. To assess whether respondents can predict one another's marginal return to capital, we utilize follow up data on business profits for those who did and did not randomly receive our grants (discussed below). Importantly, the baseline household survey was conducted *before* respondents knew they would be ranking their peers. In addition, we assess whether respondents' reports provide useful information beyond what could be gathered through measurement of observable household and business characteristics or psychometric tests.

Community knowledge–even if accurate– is only useful for allocative decision-making if those collecting the community information can be reasonably confident that they will gather *truthful* reports. But eliciting reliable information from community members is not a straightforward task. If respondents know that information they provide about their peers will be used for targeting resources, they may distort their reports to favor persons with whom they have stronger social ties. Rigol and Roth (2016) find that, in a low stakes lab experiment, farmers in a village setting distort their reports in favor of family and close friends. In this paper, we measure the level of distortion that occurs when stakes are high: half of our study participants, whom we refer to as being in the *revealed* treatment, were told that their reports would influence the probability that their peers would receive cash grants of USD 100 (grants are randomly allocated to a third of our sample).[5] This is a shock approximately equivalent to one and a half months of business profits for the average business in our sample. The control group, whose rankings do not influence grant allocation, provide an estimate of the quality of information embedded in communities in the absence of incentives to distort responses. Comparing the informativeness of the two groups' reports establishes whether community members strategically misreport information to affect

[5]The treatment name, *revealed,* refers to community reports being revealed to a principle who would make use of them.

allocative decisions (i.e. the possibility of receiving a grant).

If distortions can be corrected, their possibility does not eliminate the utility of community reports. We experimentally test whether two simple techniques from mechanism design theory can effectively correct distortions: (1) asking respondents to report in public or in private and (2) offering monetary incentives for the accuracy of reports. Screening mechanisms that are common in practice and that rely on community information (such as joint liability and participatory rural appraisals) require public reporting. If respondents fear negative social repercussions of truth-telling, mechanisms that involve public reporting may discourage reliable reporting. On the other hand, if respondents fear that misreports will cause them to be perceived as liars by their peers, public reports may lead to more truthful elicitation. To test whether reporting in public would increase or decrease the accuracy of reports, our respondents were randomly assigned to make reports in public (ranks were visible to the entire group) and or private (group members did not see ranking reports). The difference in the accuracy of reports between these two groups reveals whether privacy is an important criterion for elicitation of peer information.

In order to test whether monetary incentives impact accuracy of reports, we randomly assigned respondents to either receive no monetary incentive or to receive an incentive delivered via the *Robust Bayesian Truth Serum* of Witkowski and Parkes (2011). We discuss the details of the payment rule and our rationale for choosing it in the following section.

In addition to these treatments, we included two non randomized features of the elicitation exercise. First, we asked all respondents to report about their peers using two different methods: Ranking the five members of the group relative to one another, and separately placing the five members of the group in quintiles relative to the entire neighborhood. The relative ranking may be a coarse measure of the group's knowledge but has the advantage that it induces a zero-sum game in the sense that elevating the status of any member necessarily lowers the status of another, which may deter misreporting. On the other hand, quintiles may contain more information but are also easily manipulable; everyone can be placed into the highest position. The last elicitation mechanism we evaluate is closely related to the cross reporting techniques which play a prominent role in mechanism design and implementation theory (see Maskin, 1979). We asked each respondent to identify who each other respondent was closest to (and hence most likely to lie about) within the group. We also asked each respondent to identify who in the group would give the most accurate ranking.

In summary, our respondents were cross-randomized (at the group level) to give their ranking reports under the following three treatment conditions, for a total of eight treatment cells: Revealed vs Not

Revealed (R0 vs R1), Public vs Private (P0 vs P1), and Incentives vs No Incentives (I0 vs I1). Since important predictors of our outcomes of interest such as occupation, religion, caste, and other socio-demographic characteristics are frequently correlated with the area in which respondents live, we stratified our treatments by geographic clusters. Since there are eight treatment groups, we created geographic clusters of eight proximate groups. We then randomized our 8 experimental interventions amongst the eight groups in each cluster.[6] We also randomly selected one-third of our sample to receive USD 100 grants. Selection of grant winners was done via public lottery and is explained in detail in Section 2.4

2.3 Peer Prediction and the Robust Bayesian Truth Serum

Delivering monetary incentives is a somewhat thorny endeavor. A natural inclination might be to pay respondents based on the closeness of their reports with ex-post, objectively measured outcomes. However this approach has several drawbacks. First, it may be costly or even impossible to measure ex-post outcomes; indeed, this is one of the primary motivations of relying on community reports. Second, in cases where it is possible to measure ex-post outcomes via surveys, respondents who know their reports will determine the incentive payments of their peers may lie - a particularly undesirable outcome in a research setting where survey reports are a primary measure of ground truth. Thus we were in need of an alternative method.

As stated above, we delivered monetary incentives via the *Robust Bayesian Truth Serum (RBTS)* of Witkowski and Parkes (2011). RBTS is part of a class of mechanisms known as *peer prediction mechanisms.*[7] The major innovation in these payment rules is that they eschew reliance on ex-post verifiable information to calculate payments. The tradeoff is that these are very complicated payment rules, which rely on comparison between reports from the respondent about the question of interest (which we refer to as *first order beliefs*) and community expectations about the distribution of responses to the question of interest (which we refer to as *second order beliefs*). Roughly speaking, the payment rule rewards respondents who choose responses that are under predicted by the community more than those who choose responses that are over predicted by the community. Details about the implementation of RBTS and the intuition behind its incentive compatibility are relegated to the appendix.

Whether these payment rules are as effective as simpler payment rules for respondents with low numeracy is an empirical question addressed in Rigol and Roth (2016). In a population similar to that of the current study, Rigol and Roth (2016) test the difference between paying incentives via a simple

[6] As some neighborhoods had a number of groups indivisible by 8, some clusters have fewer than eight groups. For these clusters, we randomly selected groups to receive one of the eight treatments.

[7] See Prelec, 2004 for a seminal contribution to this literature.

rule based on ex-post accuracy and via a peer prediction payment rule. The first rule was easy for surveyors to explain and for respondents to understand; respondents were paid based on the closeness of their report and the ex-post realized outcomes. On the other hand, as in our current study, surveyors did not attempt to explain the peer prediction method. They merely elicited first and second order beliefs and asserted to respondents that they would maximize their incentive payments by telling the truth. Rigol and Roth (2016) find that the additional accuracy induced by the simple ex-post incentive is statistically and economically indistinguishable from that induced by the peer prediction method for eliciting information about borrower reliability and entrepreneurial ability.

That upon first exposure, respondents treat the two payments schemes equivalently is reassuring; because in our experiment each respondent was only exposed to our payment scheme once, we feared no loss from relying on the peer prediction scheme. However, if RBTS does not function as theoretically predicted, over a longer period respondents may experiment and discover that truth telling is not their optimal strategy. As such, Rigol and Roth (2016) take the analysis a step further by estimating the higher order beliefs of respondents in the sample. Not only do respondents report as accurately when facing the complicated incentive scheme, but we provided evidence that truth-telling is the strategy that maximizes their subjective expected payments from RBTS. Thus, even over periods of repeated exposure, respondents should continue to tell the truth.[8] Details of this exercise are replicated in the appendix of this paper.

2.4 Ranking Questions and Implementation Method

We asked respondents to rank their peers on a series of dimensions. We collected information about the following criteria: highest level of education attained, marginal returns of the peers' business if she were to receive an Rs.6000 grant, household average monthly income over the past year, projected monthly profits of the peers' business if she were to be given an Rs.6000 grant, total value of household assets, number of hours that their peers work, total household medical expenses in the previous 6 months, loan repayment trouble over the past year, and digit span memory test. For marginal returns, income, profits, and assets, we asked respondents to rank their peers relative to one another as well as to place them in quintiles of the community distribution. For the remainder of the questions, respondents were asked to report only relative ranks. We also asked a subset of groups to report who they thought deserved to

[8]Note, this is the standard to which we hold commonly used mechanisms even in the developed world. For example, medical students are placed in their first residency via the *deferred acceptance algorithm*. They largely largely do not know how it functions but they are assured that truthfully stating their preferences is their best strategy. And because we have theoretical guarantees that stating their true preferences is indeed their best option, the profession feels confident in giving them this advice.

receive the grant. We did not provide any criteria for this ranking and asked respondents to choose what they thought were important criteria.

To minimize respondent fatigue, each respondent answered only a subset of these questions. All members of the same group were asked the same ranking questions. In Figure 1 below, we lay out the question randomization structure. Because incentivized groups also had to report second order beliefs in addition to ranks (first order beliefs), they were only be asked to answer a total of 7 questions, while non-incentivized groups answered 10. The order of the first 3 questions was always the same and groups were cross-randomized between P0/P1 and I0/I1. The first question was always about education and we primarily intended it to be a practice round. We chose to elicit education quintiles as it allowed us to explain the quintile rankings early. The next two questions were always about marginal return quintiles and relative rankings. In the relevant groups we elicited marginal return information in public and with incentives but we never used marginal return information to affect the distribution of grants because we did not want reports in this dimension to be adulterated by strategic behavior. For questions 4-7, we randomly picked two of three questions: income, assets, and profits. These were cross-randomized with all 3 of our treatments and we elicited both relative rankings and quintiles. Lastly, we randomized questions 8-10 with the public and private treatments only so as to minimize the amount of time respondents spent doing the rankings exercise. Notice that because income, assets, and profits were also in the rotation for Q8-Q10, we have more data on relative rather than quintile rankings for these questions.

After all baseline surveys were completed in a particular neighborhood, groups were invited to a large community hall to conduct the ranking exercise. One group was invited to conduct the exercise at a time. As soon as a respondent arrived in the hall, he or she was seated behind a privacy screen along with a surveyor. The screen was placed both to reassure the respondents in the privacy treatment that their responses would never be visible to others in the group, but also to avoid potential coordination. Respondents were given name cards with the names of all of the peers they would be ranking. To explain complicated concepts and to minimize variation across surveyors in implementation of the treatments, we created animated videos to guide respondents through the exercise. In the videos we explained the definition of a quintile and how to do a quintile ranking, and the definitions of marginal returns to capital, profits, income, and assets.

For groups in the *Public* treatment, although respondents gave their ranks behind their privacy screens, they were asked to move with their rankings to the center of the hall at the completion of each ranking. While the pretext of the move to the center was that the lead surveyor had to record everyone's

answers, the purpose was actually that peers could clearly observe each others' rankings.[9]In the privacy treatment, respondents never interacted with other people in the group until all rankings were completed.

For those who received the incentives treatment, the videos explained that incentives would be paid for truthfulness of the responses. Respondents were told that people who reported what they truly believed were more likely to receive higher incentive payments than those who did not report what they truly believed. Since RBTS incentive payments also required respondents to report their second order beliefs, the videos were used to explain what second order beliefs were. For each of her peers, each respondent was given 20 orange coins and was asked to place the coins in proportion to how she thought others would rank her peer. Payments were calculated and distributed in private by the surveyor at the end of each ranking question. Groups that did not receive incentive treatments were not asked to report second order beliefs and were not paid for their reports.

At arrival, respondents were told that at the end of the exercise, a lottery would be conducted to choose the grant winners. Each person was given 20 lottery tickets and was told that at the end, all people present in the room would put their lottery tickets inside a basket and the winner would be selected by picking out lottery tickets. For groups in the revealed treatment, after completion of the marginal returns relative rankings, the video explained that for the next 4 rankings they would be able to help determine the lottery winner. Respondents were told the person that was ranked the highest by the group for each round would receive extra lottery tickets. Since we wanted, as much as possible, to keep the probability of selecting the winner balanced across relative ranks, only 1 extra lottery ticket was awarded for winning a round. Respondents, however, did not know how many extra lottery tickets we were awarding each round until all of the ranking exercises were over. At that point the winners were given their extra tickets and the lottery was conducted in the presence of all respondents.

[9]Surveyors report that respondents always looked at their peers' rankings.

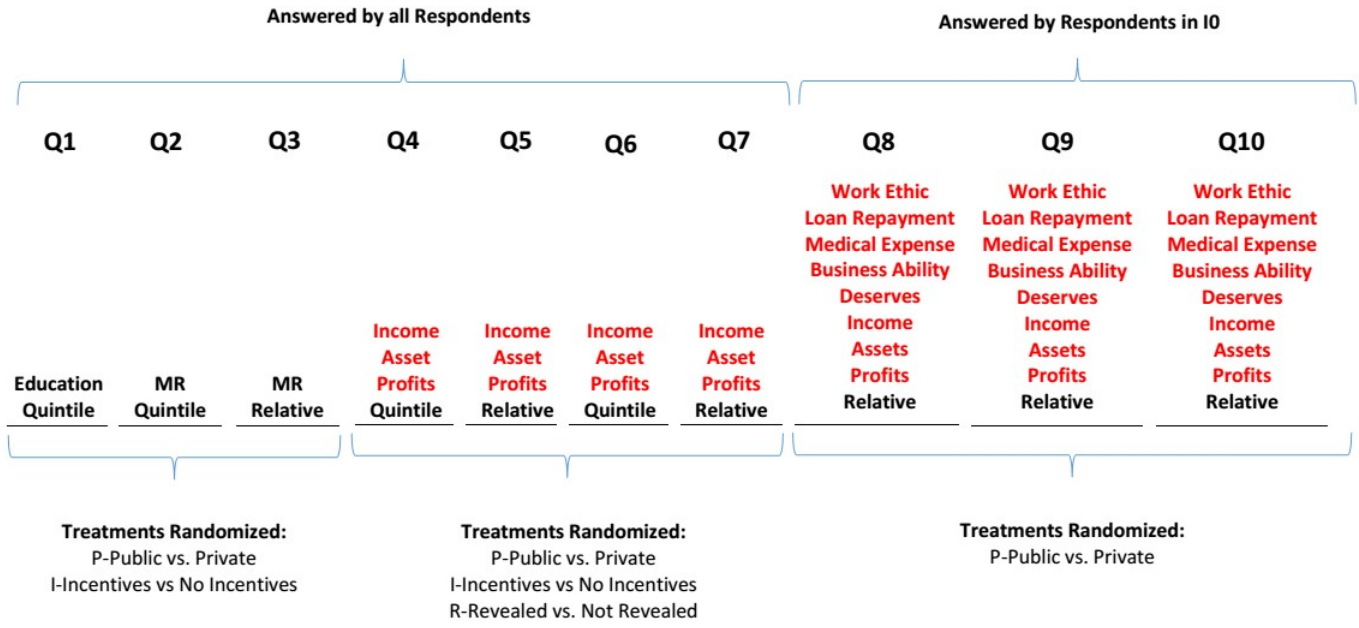


Figure 1: Question Randomization

3 Data and Sample Characteristics

3.1 The Data

The data for the experiment come from respondent surveys. Baseline surveys were conducted between December 2015 and April 2016 in the privacy of the respondent's home. Respondents were asked detailed information about the economic activities of the household. Specifically, households were asked demographic, occupation, work days and hours, and income information about all household members. They were asked to report about health expenditures and well-being as well as outstanding loans and loan repayment issues. Each business owner was asked to report about the revenues, costs, and profits, inventories, business assets, as well as other aspects of their own businesses. Importantly, all of this was done before respondents knew anything about the ranking exercise. Therefore we have detailed business module of **all** household businesses.

Business owners were asked to complete a psychometric questions section as well as the digit span memory test. A new method that has been utilized to identify credit-worthy entrepreneurs, developed by the Entrepreneurial Finance Lab, largely relies on psychometric questions. Although the precise test used by EFL in commercial institutions is proprietary, it relies on psychometric assessments similar to those

used in De Mel, McKenzie, Woodruff (2008), who find that psychometric tests predict whether people will become entrepreneurs versus wage workers in Sri Lanka. We therefore pose the same questions to our respondents, found in the appendix of this paper, as those utilized in their study. Respondents answered each question on a scale of one to five indicating whether they strongly disagreed, disagreed, do not agree or disagree, agreed, or strongly agreed with the statement. The questions are organized according to categories developed by industrial psychologists: polychronicity measures the willingness to juggle multiple tasks at the same time (Bluedorn et al. 1999) ; impulsiveness is a measure of the speed at which a person makes decisions and savings attitudes (Barratt Impulsiveness Scale); tenacity measures a person's ability to overcome difficult circumstances (Baum and Locke 2004); achievement is a measure of satisfaction in accomplishing a task well (McClelland 1985); and locus of control measures a person's willingness to put themselves in situations outside of their control (Rotter 1996).

Households were asked to report on a household assets list in which we asked whether the household owned a particular type of asset, how many pieces, and the current resale value of those assets. Our surveyors were trained to verify that the household actually owned the assets about which they reported.

Baseline data is utilized to validate whether respondents can predict information about their peers. We emphasize that this data was collected before any respondent in a particular neighborhood was aware that they would be asked to report about the household or business activities of their neighborhoods. The follow up data allows us to test whether respondents' ranks predict future outcomes such as the observed marginal returns to the grant.

Peer predictions data comes from the group ranking exercises.

3.2 Community Knowledge and Sample Characteristics

Given the selection criteria for the sample described in Section 2.1, it is no surprise that our sample is composed of relatively poor microentrepreneurs. Our sample households report earning approximately Rs.9000 per month in total earnings from all income-generating activities, or USD 5 per day, and have approximately $7000 in total household wealth. As is common in developing countries, poor households diversify across different types of income-earning activities: in 50% of our sample households there is at least one fixed salary or daily wage worker and approximately a fifth of our sample has more than one household business. Medical expenses contribute to a very large portion of household expenses: on average respondents report spending nearly 30% of their monthly earnings on health-related expenditures.

Approximately 60% of the business owners in our sample are male. On average, the firm owners are

40 years old and have about 7 years of formal education. 30% of them work in manufacturing, another 30% in retail, and another 30% work in services with the remainder being spread across construction and livestock rearing. They work 45 hours during an average week and earn about Rs.4500 in profits per month from the businesses. When asked to project the increase in monthly profits that they expected if they were to receive the USD 100 grant, they report expecting to increase their monthly profits by an average of Rs.2000 (USD 35).

Business owners in our sample are well informed about the peers in their group. We find that in less than 1% of the cases a respondent reported not recognizing a peer in her group. On average, they visited one another 22 times over the previous 30 days. 23% report discussing private family or business issues with the peer. To establish a baseline level of knowledge that peers in the group had about one another, we asked them to report some observable and easily verifiable information about one another. We asked respondents whether their peers owned a motorcycle (50% of our sample has at least one motorcycle in the home). 83% of respondents correctly identified when a peer did not own a motorcycle and 70% correctly identified when a peer owned a motorcycle. This number rises to 80% if we include other motorized vehicles in the criteria. We also asked whether there were children living in their peers' homes. As with motorcycle ownership, approximately 80% were able to correctly identify when there were and there were no children in their peers' homes.

4 Empirical Strategy

We present the regression specifications to answer the questions posed in Section 2.2. To address how well peers can predict one another's baseline outcomes we use the following model.

$$Outcome_{ikjc} = \alpha_0 + \alpha_1 Rank_{ikjc} + \gamma_c + \epsilon_{ikjc} \tag{1}$$

where i indexes the person being ranked (rankee), j indexes his group, and k indexes the person doing the ranking (respondent). $Outcome_{ikjc}$ is the percentile at which person i's outcome value lies in the distribution of that outcome for the sample. $Rank_{ikjc}$ is the percentile at which the rank that person k provides for person i in group j lies in the distribution of ranks for the sample. γ_c is a fixed effect for each cluster of groups that form the basis for our stratification. Standard errors are clustered at the group level. In all Tables presented in the paper we additionally include a surveyor and month of baseline survey fixed effect.

We measure outcomes and ranks in percentiles to maintain comparability across outcomes and reporting methods (ranks vs quintiles), and to ease interpretation. A 1 percentile increase in the rank is, therefore, associated with an α_1 percentile increase in the outcome variable distribution.

To estimate the residual information that ranks provide over and above observables, we will estimate regression 1 with the addition of $\mathbf{X}_{ijkc}$, a vector of person i's baseline characteristics that may impact the outcome value. If ranks have residual predictive power, we expect α_1 to remain positive and significant. We will also estimate a related version of this model in which, rather than using each $i-k$ pair as the unit of observation, we will average across all of the reports that peers in group j gave about person i. $\bar{Rank}_{ijc} = \sum_{k=1}^{n} \frac{1}{n} * Rank_{ikjc}$ where n is the total number of group members in group j

$$Outcome_{ijc} = \alpha_0 + \alpha_1 \bar{Rank}_{ijc} + \gamma_c + \epsilon_{ikjc}$$

To address whether respondents can predict one another's marginal return to capital we estimate the following model.

$$Profit_{ikjct} = \alpha_0 + \alpha_1 Winner_{i,t} + \alpha_2 Rank_{ikjc} + \alpha_3 Winner_{i,t} \times Rank_{ikjc} + \alpha_4 \mathbf{X}_{ijkc} + \alpha_5 Winner_{i,t} \times \mathbf{X}_{ijkc} + \gamma_c + \gamma_t + \epsilon_{ikjct}$$

where $Profit_{ikjct}$ is the profit for person i being ranked by person k in group j, cluster c and time t, $Winner_{i,t}$ is an indicator for whether person i has (randomly) received a grant at time t, γ_t is a wave fixed effect and all other variables are as defined above. α_2 captures the average increase in marginal return to capital associated with being one rank higher in the community prediction. In some specifications we control for borrower and business characteristics $\mathbf{X}_{ijkc}$ as well as its interaction with $Winner_{i,t}$ to measure the additional information conveyed in community reports.

To answer question (2), whether respondents manipulate their rankings when they have a strategic incentive to do so, we test how the accuracy of their responses changes when they are assigned to the revealed treatment (they are told that their ranks would affect the probability that their peers win the lottery) . Specifically, we estimate

$$Outcome_{ikjc} = \beta_0 + \beta_1 Rank_{ikjc} + \beta_2 Rank_{ikjc} * Revealed_{jc} + \beta_3 Revealed_{jc} + \gamma_c + \epsilon_{ikjc} \qquad (2)$$

where $Revealed_{jc}$ is a dummy for whether group j was assigned to the revealed treatment. β_1 indicates the accuracy of reports in the group that was told that was told nothing about reports being used to inform the distribution of grants and so respondents had no explicit incentive to misreport. β_2 is then the

differential effect of being assigned to the revealed group. If respondents manipulate rankings, we expect β_2 to be negative - the correlation between reports and true outcomes should be lower for this group.

A more direct way to test whether respondents are lying is to check whether they manipulate the reports that they give about peers in the group that they share a close relationship with (measured by self and cross reports). Specifically we estimate,

$$Rank_{ikjc} = \pi_0 + \pi_1 Member_{ijc} + \pi_2 Member_{ijc} * Revealed_{jc} + \pi_3 Revealed_{jc} + \gamma_c + \epsilon_{ikjc}$$

$Member_{ijc}$ is a dummy that will signify one of three types of relationship between the rankee and the respondent: they are family members, they are close peers (as reported by other group members), or the it is the rank that the respondent gives herself. If $\pi_2 > 0$, then respondents up rank these individuals irrespective of their true characteristics.

Lastly, to answer question (3) - whether monetary incentives or reporting in public/private improve the accuracy of reports - we estimate the following model

$$Outcome_{ikjc} = \theta_0 + \theta_1 Rank_{ikjc} + \theta_2 Rank_{ikjc} * Public_{jc} + \theta_3 Rank_{ikjc} * Incentives_{jc} + \tag{3}$$

$$\theta_4 Rank_{ikjc} * Public_j * Incentives_{jc} + \theta_5 Public_{jc} * Incentives_{jc} + \theta_6 Public_{jc} + \theta_7 Incentives_{jc} + \gamma_c + \epsilon_{ikjc}$$

$Public_{jc}$ is a dummy indicating whether the group is assigned to report rankings in public (vs private). $Incentives_{jc}$ is a dummy for whether the group is assigned to receive personal incentives for its reports. θ_1 indicates the accuracy of reports when respondents report in private and receive no incentives. θ_2 is the differential effect of reporting in public without incentives. θ_3 is the differential effect of reporting in private and receiving incentives, so if incentives improve rankings, we expect this to be positive. Lastly θ_4 is the differential effect of reporting in public and receiving incentives relative to either treatment on its own.

4.1 Randomization Check

In Appendix Table 1, we present the randomization check of baseline characteristics by treatment. To check for balance we estimate the model

$$Characteristic_{ijc} = \tau_0 + \tau_1 Treatment_{jc} + \gamma_c + \epsilon_{ijc}$$

where $Treatment_{jc}$ is a dummy for whether the group was assigned to the $Revealed_{jc}$ treatment (columns 1 and 2), the $Incentives_{jc}$ treatment (columns 3 and 4), and the $Public_{jc}$ treatment (columns 5 and 6). The odd columns show the average of each characteristic for the control group in each block. So column 1 shows the means of characteristics for groups that were assigned to *Not Revealed*. The even columns show τ_1 for each treatment (the difference between treatment and control characteristics) . The characteristics in rows 1-6 are at the household level. The remainder pertain to the business owner that was selected to participate in the study. As would be expected by chance, there is some imbalance across a few characteristics between the various treatments. Groups assigned to *Revealed* report lower yearly income and slightly fewer years of education, and the incentivized group has somewhat lower baseline assets than the unincentivized group and lower projected monthly profits if they were to receive the grant. There is no significant difference in the means of characteristics in the *Public* treatment. We also present the results of a joint test of statistical significance, and cannot reject that all groups are drawn from the same population.

5 Results

5.1 Can Communities Predict Marginal Returns To Capital?

We ask whether respondents can predict one another's true marginal return to the grant. To answer this question, we generate a simple graph that compares the profits of the grant winners and losers by the marginal return rankings given to each individual entrepreneur at baseline. Figure XX shows a local polynomial regression of the log profits at followup for the treatment (grant winners) and control (grant losers) groups on the percentile of the average rank assigned by their peers.[10] The difference between the two lines is the marginal returns to the grant.

Both the treatment and control lines are upward sloping, implying that the marginal returns rank is

[10]To calculate the average marginal returns rank we take the simple mean of every individual report given about person i by all of her peers in the group.

positively correlated with the level of profits in the followup period.[11] Remarkably, it is also correlated with *marginal* profits. Profits in the treatment and control groups below the $35th$ percentile of the marginal returns distribution are indistinguishable. Above the $35th$ percentile, however, the distance between the treatment and control profits increases with the marginal returns rank, indicating that community rankings are highly predictive of marginal return to capital. In this simple formulation, while the returns at the $1st$ percentile are 0, the grant winner at the $99th$ percentile of peer ranks earns an extra 45% in profits per month. This increase in profits translates to an 18% monthly return to the grant for the entrepreneur in the $99th$ percentile of the ranks distribution.

Had we asked our grant winners to repay their in grants with a 26% APR - the Reserve Bank of India mandated cap on MFI loan interest rates - the principal and the interest payment would amount to Rs.570 per month. While there are many reasons why a loan may have induced different selection and investment patterns[12], as a simply back-of-the-envelope calculation, all entrepreneurs above the $58th$ percentile of the ranks distribution would have been able to earn a net positive return on their investments.

We confirm the results of the local polynomial regression using a difference-in-differences model, which, unlike the local polynomial regression, allows us to additionally exploit variation in baseline profits and income .[13] As outlined in Section 4, our primary specification is

$$Y_{ijt} = \alpha_0 + \alpha_1 Winner_{it} + \alpha_2 Winner_{it} \times \bar{Rank}_{ij} + \alpha_3 Winner_{it} \times \mathbf{X}_{ij} + \gamma_i + \sum_{t=1}^{3} \delta_t + \theta_m + \tau_s + \epsilon_{ijt} \quad (4)$$

where Y_{ijt} measures either business profits or household income of person i in survey round t. As we use 3 followup rounds of data, t ranges from 0 (baseline) to 3.$Winner_{it}$ is an indicator for whether person i won a grant at or before survey round t. Note that $Winner_{it}$ is 0 at period t for all persons i. $\bar{Rank}_{ij}$ is the average marginal returns rank assigned to person i by the members of group j. We also include the following fixed effects: person (γ_i), survey round (δ_t), survey month (θ_m), and surveyor (τ_s). Standard errors are clustered at the group level. α_3 measures the average additional marginal return to capital an entrepreneur enjoys for every additional rank he is assigned by his group. Since this differences-in-differences specification assumes a linear interpretation of the treatment effect, we also present a non-linear model in which the ranks distribution is divided into terciles and interacted (as above) with $Winner_{it}$.

In Tables XX-XX, we present the effect of the grant on household income and total household profits.

[11] The marginal returns rank is also positively correlated with baseline profits.
[12] See Fiala (2014) for an experiment that randomly allocates loans and grants to entrepreneurs.
[13] This is the regression model used in De Mel, McKenzie, and Woodruff (2008).

Following De Mel, McKenzie, and Woodruff (2008), we remove the outliers of the household income and total profits distributions (levels) by trimming the top 0.5% of both the absolute and percentage changes in profits measured from one period to the next. We also present the $log(Y_{ijt}+1)$ of the untrimmed income and profits distributions.[14]

Across the various specifications in the odd columns of Table XX, α_3 is large and positive, with varying levels of statistical significance. An extra average rank increases profits and income by between Rs. 283.2 and Rs. 848.1, or 4.7% and 14% returns to the grant, per month. The mean and standard deviation of the average ranking of marginal return to capital are 3.46 and 0.66, respectively. Therefore, while the estimates in Column 1 imply that the average marginal return to capital in the whole population was about 7.1% per month, the average entrepreneur ranked one standard deviation above the mean enjoyed an 16.4% monthly marginal return to capital, and an entrepreneur ranked two standard deviations above the mean enjoyed a 25.7% marginal return to capital.

These calculations, however, assume a linear impact of the grant by average rank level. To check whether this is a reasonable assumption, we divide the ranks distribution into terciles and present these results in the even columns of Table XX. Across the board, we cannot reject that the bottom tercile has no differential returns to the grant, with the coefficient on $Winner_{it}$ implying a negative return. Consistent with Figure XX, the coefficients on log income and log profits for the middle tercile are positive, but not significant, while the level effects are almost precisely 0.[15] We find that the strongest effects are concentrated in the top tercile of the distribution and that we can statistically reject that the effect in the middle and top terciles are the same. Using the coefficients on the $Winner_{it} \times \bar{TopTercile}_{ij}$ in columns 2 and 6, the returns to the grant for the top tercile of the ranks distribution ranges between 16.6% and 26.7%.

Regression specification 4 identifies the treatment effect off the within-household differences in profits and income in the pre and post grant disbursal periods. As a robustness check, we also present the results using an ANCOVA specification in Appendix Table XX. In ANCOVA, the treatment effects are identified by comparing the cross-sectional differences between treatment and control in the post grant-disbursal periods, while controlling for the baseline value of the outcome characteristic. [16] The estimated model is:

[14]The results remain nearly identical whether we log-transform the trimmed or untrimmed income and profits distributions.

[15]Mechanically, since the middle tercile is fixed, the difference between the level and log results occurs because there are some extreme right-tail observations in the distribution of income and profits for the middle tercile ranks. The weight of these outliers in the regression is diminished when the distributions are log-transformed.

[16]For a detailed discussion of ANCOVA, please refer to McKenzie (2012).

$$Y_{ijt} = \beta_0 + \beta_1 Winner_i + \beta_2 Winner_i \times \bar{Rank}_{ij} + \beta_3 \bar{Y}_{ijPRE} + \delta_t + \theta_m + \tau_s + \epsilon_{ijt} \quad (5)$$

where Y_{ijt} are post-treatment outcomes (so t ranges from 1 to 3 rather than 0 to 3 as in regression model 4) and $\bar{Y}_{ijPRE}$ is the pre-treatment (time 0) value of the outcomes.

Results in the ANCOVA specification are qualitatively very similar. First, the bottom tercile has a negative (although often noisily estimated) return to the grant. In the following section we explore the grant's effects on business inputs, but one potential reason for this negative return is that winners in the bottom tercile work fewer hours per week and fewer days per month due to the grant compared to business owners in the bottom tercile of the control group. The treatment effects on the top tercile are always larger than those in the middle tercile, but in the ANCOVA specification the middle tercile is more clearly benefiting from the grant. In the profits and log profits specifications, we can reject that the treatment effects are the same for the top two terciles but not in income and log income. Overall, returns to the top tercile lie in the range of 18.2%-20.5% per month.

5.1.1 How Do Top-Ranked Entrepreneurs Behave in Response to the Grant?

How do the top-ranked entrepreneurs attain such high rates of return to the grant? We begin by analyzing how they invest the money that they receive. We asked all grant winners to report how they invested the grant in followup rounds of data collection. Grant expenditures were divided into business expenses (inventory, durable assets, labor, and other) and non-business expenses (loan repayment, relending, household repairs, and other household expenses). We also asked respondents if they added any money to the grant to make business purchases. In Table XX we examine the relationship between how respondents report spending the grant and the average community ranking of marginal return to capital. To do so we regress total grant expenditures in each category on whether an entrepreneur is in the top or middle terciles of the marginal returns ranks distribution. Since the omitted group are entrepreneurs in the bottom tercile, the coefficients on top and middle terciles tell us the differences in grant expenditures between the bottom and upper terciles.

While we are underpowered to detect statistically significant differences between the top and middle tercile, business owners in the top tercile invest an extra Rs. 903.1, or 25% more of their grants, in their businesses than those in the bottom tercile. Most of this difference appears as higher expenditures in inventories. Both the top and middle terciles spend significantly less money on "Other Household Expenses" - generally medical expenses, education, food consumption, etc. - and are less likely to have

saved their grant money.

Note that respondents who did not win the grant (our control group) were excluded from this exercise. Furthermore, money is fungible so reports of grant expenditures do not necessarily hold constant pre-grant consumption in these different categories. To investigate whether this pattern of grant investments translates to real marginal increases in business inputs, we compare the inventories, business assets, and labor outcomes of grant winners and losers.

In Table XX, we again utilize regression specification 4 to analyze differences in business inputs between grant winners and losers by terciles of the marginal returns distribution. Consistent with the pattern of investments in the previous table, the grant induces top and middle ranked entrepreneurs to accumulate higher capital stocks: top tercile grant winners report an extra $Rs.\,4541.0$ worth of inventory and an extra $Rs.\,8281.1$ of durable assets. The treatment increases the capital stock (inventory plus durable assets) by approximately 200% of the grant amount, which is large, but within the confidence bound of increases in capital stock in De Mel, McKenzie and Woodruff (2008).

The marginal returns rankings are also correlated with marginal increases in productivity: in columns (1) and (2) we show that grant winners in the top tercile spend an extra 10 hours per week and an extra 4 days per month working when compared to their untreated counterparts. *XX IS THIS RIGHT??XX*. The treatment also has an impact on the amount of household and non-household labor. At baseline, 21% of households in our sample utilize household labor for an average of 30 hours per week (conditional on having a worker) and almost none of these workers are officially paid a wage. Despite our selection criteria (see section 2.1 for a detailed explanation), 9% of households report utilizing non-household labor in their businesess. Conditional on having a paid non-household laborer, the weekly wage bill for these businesses is on average $Rs.\,3221$. The grant induces top-ranked entrepreneurs to be 8.1% more likely to have a household laborer and 6.4% more likely to have a non-household laborer.

Overall, we see that top ranked households invest higher proportions of their grants into the business, turn those investments into higher business stock, and devote more time to working in their businesses.

5.1.2 Credit constrained?

We find that for a significant portion of our sample, returns exceed the interest rates charged by banks and microfinance institutions. De Mel, McKenzie and Woodruff (2008), argue that heterogeneity in the marginal returns to firms in their sample is driven by credit constraints rather than missing insurance markets. Indeed, there is a deep experimental and non-experimental literature arguing that

credit constraints are an important impediment to microenterprise growth and potentially responsible for the overabundance of small firms in low-income countries relative to the distribution of firm sizes in high-income countries (Hsieh and Olken, 2014).

Over 92% of enterpreneurs in our sample report having a desire to borrow, but only 6% have ever borrowed from a bank and less than 20% have ever borrowed from a microfinance institution. Like De Mel, McKenzie and Woodruff (2008), we also find evidence that returns are larger among entrepreneurs who appear more credit constrained: more able entrepreneurs (more highly educated and can remember more digits in the Digit Span Memory test) and entrepreneurs who should have a higher shadow value of capital (as the authors posit, entrepreneurs who have less household wealth and fewer wage workers).

5.1.3 Who Are the Top Ranked Entrepreneurs?

To understand whether top-ranked entrepreneurs look different than those ranked in the middle and bottom terciles we compare the baseline charactersitics of households in all three in Table XX. In Column 1, we present the mean of each characteristics for the bottom tercile group. Columns 2 and 3 are the coefficients from a regression of the characterstic on whether the person is ranked in the middle (2) or top (3) terciles. So each coefficient is interpretable as the impact of being in the upper terciles relative to being in the bottom tercile.

Top-ranked entrepreneurs are 8 percentage points more likely to be male, are younger, and are less likely to be married. While they do not appear to be more educated, they do remember an average of 0.57 digits more in the digitspan memory test than their counterparts in the bottom tercile.

While the total number of businesses in the households of top-ranked entrepreneurs is the same as those in the bottom, the businesses of top ranked entrepreneurs have 52% larger in terms of assets and earn 40% higher profits per month. The business owners in top-ranked households work an extra 6.8 hours per week and 1.8 days per month. Top-ranked entrepreneurs are more slightly more likely to be engaged in a food preparation business than the bottom-ranked entrepreneurs, but otherwise they differ little by business type. We asked business owners how much a salaried job would have to pay per month in order for them to exit self-employment. Top ranked entrepreneurs report that they would require 22% higher monthly wages to shut down their businesses.

For the most part, entrepreneurs in the middle tercile have baseline characteristic means between that of bottom and top ranked entrepreneurs. Two notable exceptions are that they have

5.1.4 Predictions with Observables

While it is clear that community information is highly predictive of realized marginal returns, a first-order question is whether the information contained in the reports can be predicted with observable characteristics about business owners and her household. Table XX examines how much the value of community reports diminishes when controlling for baseline business and household characteristics, by including in the regression model in the interaction of these characteristics with $Winner_{it}$. The odd columns in Table XX control for a set of characteristics that would be easily observable and verifiable by a loan officer and which tend to be predictive of the marginal return: the entrepreneur's gender, marital status, age, education, and business type. We find that community information is almost orthogonal to these characteristics; the estimates in Table XX are strikingly similar to those obtained without controls.

In the even columns of Table XX we also control for harder to observe characteristics: value of business assets at baseline and average yearly profits. If anything, controlling for baseline profits seems to enhance the value of community information, as the estimates in Table 13 are substantially larger than the preceding ones. The reason for this is that the marginal returns rank is positively correlated with baseline profits and baseline profits are negatively correlated with the marginal return to capital (implying that there are diminishing returns to capital).[17]

Another set of characteristics that have been identified as predictive of credit worthiness and entrepreneurial aptitude are psychological questions that identify characteristics such as tenacity, polychronicity, and optimism (see Klinger, Khwaja, and Carpio 2013). In Appendix Table XX, we therefore test how well psychometric questions perform at predicting marginal returns and how they compare to community rankings. The regressors are labeled according to the psychological trait for which they are meant to proxy (the specific wording of the statement is found in the Appendix). The traits that are strongly predictive of marginal returns fall into two categories: optimism and achievement. Optimism negatively predicts marginal returns: business owners who are more likely to agree with the statements "In times of uncertainty I expect the best" and "I'm always optimistic about the future" and those who are more like to disagree with "If something can go wrong with me, it will" have lower self-reported marginal returns. People who agree with the statement "Part of my enjoyment in doing things is improving my past performance" tend to have higher marginal returns. Importantly, the value of the predictions remain almost identical to the original results presented in Table XX.

[17] To see this, we regress income and profits on the interaction of baseline profits with winner and winner with the fixed effects specified above. We find that the coefficient on the interaction between winner and baseline profits is negative and significant at the 1%. Regression available from the authors.

Overall, community information is valuable in identifying entrepreneurs with high marginal return to capital even if the implementing organization has access to a large variety of observable demographic and business information.

5.2 Do Peers Distort Their Responses When There Are Real Stakes?

The analysis in Tables XX-XX has shown that communities are well informed about the marginal returns of its members. Why is that information not being leveraged by lending institutions? One reason may be that acquiring this information *truthfully* is not straightforward. When faced with allocating a highly desirable resource (such as a loan), community members may alter their reports to benefit or hurt particular individuals in a social network. Aside from quantifying how much information neighbors have about one another, the second major goal of this project is to quantify whether and how much peers misreport in high stakes settings. Half of our sample was informed that they could affect the probability that their peers (or themselves) would win the $100 grant. The other half was not given any information about how their rankings would be used and were clearly informed that, at the completion of the ranking exercise, a lottery (with all group members have equal probabilities of winning) would be help to allocate the grant. To test whether respondents behave strategically, we will compare the accuracy of reports in groups in which ranks affected the grant allocation (*Stakes*) and groups in which ranks had no effect on grant distribution.

5.2.1 How Much Do They Know?

As explained in Section 2.4, we implemented this treatment only after the marginal returns rankings were completed. We will therefore estimate misreporting on the remaining set of questions. As with the marginal returns predictions, we begin by showing that respondents have information about their peers in the groups. Unlike with marginal returns in which we can only calculate an average return as the difference between the outcomes of the grant winners and losers, for the remaining questions asked during the ranking exercise, we have an individual outcome variable for every respondent that is ranked. We quantify how much respondents know about community members by estimainge the correlation between the ranks given to a respondent by her peers and a respondents' true outcome value for that ranking (regression model 1).

In Table XX, we pool responses across all treatments (*Stakes*, *Incentives*, and *Public*). The outcome variables, denoted in the column headings, were elicited from clients (rankees) during the baseline survey,

which was conducted before respondents had any knowledge of the rankings portion of the study and before any ranking exercises had been conducted in the community. They correspond exactly to the outcome that peers were asked to rank on during the ranking exercises. In columns 1 and 2, the outcome is average monthly household income over the past year. In columns 3 and 4, the outcome value is the clients' predicted monthly profits if they were to receive a $100 grant.[18] In columns 5 and 6, the outcome is the total value of household assets, which were visually verified by our survey team. In column 7, we report a households total medical expenses in the past six months. In column 8, the outcome variable is the average number of hours the client works per week and in column 9 we report the total number of digits the client remembered during a digit span memory test.[19] *Rank* represents the rank given to the rankee by each ranker in the category represented in the column. For example, *Rank* in column 1 is the income quintile rank, in column 2 it is the income relative rank, etc.

Following Alatas et al. 2012 *XX VERIFY XX*, in Panels A and B, both the outcome variables and *Rank* have been converted to percentiles of the distribution of observed outcome values. So a 1 percentile increase in the rank is associated with an α_1 percentile increase in the outcome variable distribution. In Panel C, the outcome variable and regressors are presented in levels.

Panel A reports the results at the ranker-rankee pair level of observation. For all outcome measures, the reports are highly predictive. Peers are informed about relatively observable aspects of their neighbors' lives, such as household assets, but also much less easily observable characteristics including working memory, work ethic, and business profits. The level of accuracy of the prediction varies depending on the question. Respondents are most accurately able to predict peers' household assets: a 1 percentile increase in the relative rank of assets is associated with a 0.154 percentile increase in the distribution of household assets in our sample population. To contextualize how "predictive" these estimates are, we regress the income percentile on the percentile of the education of the household head and percentile of assets: a 1 percentile increase in the education distribution is associated with a 0.128 percentile increase in the distribution of household income; the $\alpha_1$on households assets is 0.308. Notably, like with the marginal returns regressions, community information is again predictive of outcomes over and above what is predicted by observables.[20]

We asked respondents to rank their peers relative to others in the group and also relative to the

[18] We use this outcome as it corresponds to the ranking question posed to respondents. In Appendix Table 1 we show that the results remain nearly identical if we use average yearly profits as the outcome variable.

[19] Digit span is a commonly used test for working memory. Respondents are shown flashcards with an increasing number of digits and asked to recall the numbers from memory. The surveyor records the total number of digits that the respondent correctly repeated back.

[20] This regression can be found in Appendix Table XX.

community by reporting the quintile of the outcome distribution that they believe the peer to be in. In theory, quintile rankings could be more useful as they could contain more information about the true position of the peer vis-a-vis other similar microentrepreneurs. In both the even and odd columns 1-6 of Table XX the outcome variable is the same, but what changes is the method of reporting. In the odd columns, the regressor is the percentile in the relative rank distribution and in the even columns the regressor is the percentile of the quintile rank distribution. By comparing the odd and even columns, we find that relative and quintile rankings are equally informative.

Respondents may have idiosyncratic preferences for misreporting about certain peers in their group and may otherwise make idiosyncratic errors. One way to reduce the influence of the errors is by averaging across all reports given about a particular group member. We do so and present these results in Panel B, where the unit of observation is the rankee. We observe that the average reports are significantly more predictive of all outcome variables. While a 1 percentile increase in the profits rank leads to a 0.145 percentile increase in the profits distribution in Panel A, it is associated with a 0.246 percentile increase in Panel B. Averaging reports is therefore a costless way to nearly double the predictiveness of community reports.

5.2.2 Do Peers Distort their Responses in a Real Stakes Setting?

Having established that respondents have valuable information about baseline data, we can evaluate whether they distort responses. In Table XX, we pool the percentilized outcomes and ranks across all 9 columns presented in Table XX. Column 1 shows the regressions at the ranker-rankee level of observation (analogue of Panel A in Table XX) and Column 2 is the regression with the average rank (analogue of Panel B in Table XX). The *Rank* variable captures the accuracy of the report in the control group (*No Stakes*). The $Rank * Stakes$ coefficient tells us whether the rankings are differentially informative when respondents are told their ranks will be used to help determine grant allocation.[21]

First, as observed in the previous table, the average predictiveness of ranks in the control group increases significantly when reports are averaged: in Column 1, a 1 percentile increase in the rank distribution is associated with a 0.162 shift in the outcome distribution in the individual regressions and a 0.252 shift in the average regression (Column 4). Secondly, the coefficient on $Rank * Stakes$ is large, negative, and significant. We should note that this was not ex-ante obvious: the *Stakes* treatment may have had a positive effect since revealing ranks may have caused respondents to focus and take the exercise

[21] To reduce clutter in the regression table, we have ommited the *Stakes* coefficient from the regression report as it does not contain information relevant for the interpretation of results. *XX XX*

more seriously. The regression implies that responses are significantly less accurate when respondents have an incentive to behave strategically: in the pooled individual regression in Column 1, the responses become 34.6% less accurate in the *Stakes* group. While averaging reports increases accuracy in the control group, it does not reduce the effect of misreports. While the loss in accuracy in Column 4 is smaller as a percent of the information obtained in the treatment group (24.0%), it has nearly the same size negative impact (-0.060 in Column 4 versus -0.056 in Column 1).

Lastly, we hypothesized that while quintile ranks could theoretically contain more valuable information about rankings because entrepreneurs are compared to the commmunity more broadly than only the group, it could also mean that it is easier to lie about peers. We separate the pooled questions into quintile ranks questions (the aggregate of Columns 1, 3, and 5 in Table XX) and relative ranks questions (Columns 2, 4, 6, 7, 8, and 9 in Table XX) and test for distortions. The coefficients on *Rank* in the individual (Columns 2 and 3) and the average regressions (Columns 5 and 6) are very similar, implying that in the absence of high-stakes, the value of information from relative and quintile ranks is very similar. While the coefficient on $Rank * Stakes$ in the quintile regressions is larger in magnitude both in the individual and average models, we cannot reject that respondents misreport by the same amount in either type reporting method.

5.3 Can Mechanism Design Tools Correct Incentives to Misreport?

5.3.1 Monetary Incentives and Public Reporting

In the previous section we provided evidence that respondents distort their reports when they have a strategic incentive to do so. These distortions have a substantial impact on the accuracy of reports, particularly when respondents have an incentive to misreport such as when their information is used to allocate a desired good. Can we use mechanism design tools to generate incentives for truthful reporting? We test two randomly assigned tools: incentive payments for the accuracy of reports and reporting in public versus private.

In Table XX, we provide evidence of the accuracy of reports in the *Public*, *Incentives*, or *PublicandIncentives* treatments. The coefficient on *Average Rank* indicates the accuracy of reports in groups in which respondents do not receive incentives and report in private. Reports are averaged. Columns 2 and 3 show the results in the *No Stakes* and *Stakes* treatments, respectively. In Column 1, we pool across these two treatments. In the *No Stakes* treatment, public reporting and incentives both improve the accuracy of reports by $66\% - 70\%$. What this implies is that in reporting in the no-stakes, private, and no incentives

setting is not the best benchmark for how accurate reports could be. *XX SHOULD WE SHOW THE INDIVIDUAL? XX* The coefficent on the treatment in which respondents receive montetary incentives and report in public is negative and significant. We can, however, reject at the 10% level that the accuracy of information in this group is the same as in the private reporting and no monetary incentives group. We interpret the negative coefficient, therefore, as being indicative that the best we can do in improving accuracy is by giving either monetary incentives or asking people to report in public. Assigning both cannot improve upon each one individuallly.

Our main goal in these treatments, however, was to test whether mechanism design tools can correct incentives to misreport created by elicitation in high-stakes settings. How does these tools perform in the *Stakes* treatment? In Column 3, we see that the treatment effects have large standard errors, but we can surmize that *Public* has no effect and *Incentives* could be relevant to increasing accuracy.

To understand this result we explore how the two treatments affected the reports of different people in the group. In all tables presented up to this point, both individual and average reports included how the respondent ranking herself. In Table XX, we look separately at how the treatments affected reports about self (Columns 1 and 2) and about everyone else in the group (Columns 3 and 4) in the *No Stakes* (odd columns) and *Stakes* treatments (even columns). First, we see in Columns 1 and 2 that providing stakes decreases the accuracy of reports about self by 28% in the private and no incentives treatments. Secondly, neither incentives nor public reporting have any impact on the accuracy of reports about self. Notably, though, while our treatments were not successful in incentivizing respondents to be truthful about themselves in a high-stakes setting, even when they do lie, the accuracy of their reports about themselves (0.146) is approximately the same as the accuracy of reports about others in the group in the private and no incentives treatment (0.120 in Column 4).

Turning to Columns 3 and 4, we see that incentives significantly increase the informativeness of reports about others in both the *No Stakes* and *Stakes* treatments: in both cases, monetary incentives improve accuracy by over 100%. The *Public* treatment on the other hand continues to yield no differential effect in the *Stakes* setting. While monetary incentives will unambiguously incentivize truth-telling, public reporting has an ambiguous prediction on accuracy: in the *No Stakes* treatment, there may be pressure to report close friends or family in the group as better than they actually are for the ranker to allow them to save face. At the same time, there is pressure by non-family members and other peers to be truthful. When we introduce stakes, both of these forces are intensified: family members and close friends want the ranker to sway the lottery in their favor, but it may also be especially important to the community

that the ranker be truthful when there are high stakes.

To shed some light into how respondents trade off these alliegence, we turn to Table XX. In this table, we examine the rankings that respondents give to family and non-family members in the group under different treatment conditions. The outcome variable is what rank respondent i assigned to peer j in the group. We exclude ranks given to ones own self. $Family$ indicates whether respondent i is related to peer j. We see that compared to non-family peers (the omitted group), family members receive 0.4 higher ranks on average when reports are given in private and in the no-stakes setting. The coefficent on $Family * Public$ indicates the additional effect of ranking a family member in public without real stakes. The $Public$ treatment fully corrects the extra ranks that are given to family members in private. This is consistent with the result in Column 3 of Table XX which shows that public reporting in a no-stakes setting improves the quality of reports. We cannot reject that introducing stakes in a private setting has no additional effect over how family members are ranked.[22]Lastly, when respondents are asked to rank in public and with stakes, the correcting effect of the $Public$ treatment is mostly undone: respondents rever to ranking their family members more highly than non-family members. This again is consistent with the result in Column 4 of Table XX which shows that public reporting in a stakes setting has no differential impact over private reporting in a stakes setting.

6 Conclusion

We find that community members have valuable residual information about their peers that can be useful in targeting. Not only can community members identify one another's business characteristics, but they can also predict which of their peers have high return to capital. By distributing cash grants and measuring profits over time we find that we can reliably identify entrepreneuers with profitable investment opportunities. In particular, while the average impact of our grant on business profits was modest, community information can reliably identify entrepreneurs with upwards of 14% *monthly* marginal return to capital.

While community information appears to be quite valuable for targeting, we also find that its accuracy is sensitity to the conditions under which it is elicited. In particular, we identify a natural tendency for respondents to favor their friends and family members. Moreover this tendency is amplified when respondents are told that their reports influence the distribution of grants. However we also find that a variety of techniques motivated by mechanism design theory are effective in realigning incentives for

[22]In fact, private stakes drive respondents to rank themselves more highly.

truthfulness. In particular, small monetary payments for accuracy, eliciting reports in public, rather than in private, and cross reporting techniques to identify which community members are likely to favor one another all substantially improved the accuracy of reports. We therefore hope that the techniques identified in this paper may prove useful in improving the targeting of financial services.

References

- Alatas, V., Banerjee, A., Hanna, R., Olken, B. A., & Tobias, J. (2012). Targeting the Poor: Evidence from a Field Experiment in Indonesia. The American Economic Review, 102(4), 1206-1240.
- Alatas, V., Banerjee, A., Chandrasekhar, A. G., Hanna, R., & Olken, B. A. (2012). Network structure and the aggregation of information: Theory and evidence from Indonesia (No. w18351). National Bureau of Economic Research.
- Attanasio, O., Augsburg, B., De Haas, R., Fitzsimons, E., & Harmgart, H. (2011). Group lending or individual lending? Evidence from a randomised field experiment in Mongolia.
- Beaman, Lori, and Jeremy Magruder. "Who gets the job referral? Evidence from a social networks experiment." The American Economic Review (2012): 3574-3593.
- Bhattacharya, V. and Roth, B. (2015). Rank-Based Scoring Rules for Binary Events. Working Paper
- Brier, Glenn W. "Verification of forecasts expressed in terms of probability."Monthly weather review 78.1 (1950): 1-3.
- Bryan, G. T., Karlan, D., & Zinman, J. (2012). Referrals: Peer screening and enforcement in a consumer credit field experiment (No. w17883). National Bureau of Economic Research.
- Carpena, F., Cole, S., Shapiro, J., & Zia, B. (2013). Liability structure in small-scale finance: evidence from a natural experiment. The World Bank Economic Review, 27(3), 437-469.
- Delavande, Adeline, Xavier Giné, and David McKenzie. "Measuring subjective expectations in Developing Countries: A Critical Review and new evidence." Journal of Development Economics 94.2 (2011): 151-163. APA

- Ghatak, Maitreesh. "Group lending, local information and peer selection."Journal of development Economics 60.1 (1999): 27-50.
- Ghatak, Maitreesh. "Screening by the company you keep: Joint liability lending and the peer selection effect." The Economic Journal 110.465 (2000): 601-631.
- Giné, X., & Karlan, D. S. (2014). Group versus individual liability: Short and long term evidence from Philippine microcredit lending groups. Journal of Development Economics, 107, 65-83.
- Hossain, T., and R. Okui (2013): The Binarized Scoring Rule," Review of Economic Studies, 80(3), 984-1001
- Karni, E. (2009): A Mechanism for Eliciting Probabilities," Econometrica, 77(2), 603-606.
- Levitt, Steven D., and John A. List. "What do laboratory experiments measuring social preferences reveal about the real world?." The journal of economic perspectives (2007): 153-174.
- Offerman, T., J. Sonnemans, G. Van De Kuilen, and P. P. Wakker (2009): A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes," Review of Economic Studies, 76(4), 1461-1489.
- Prelec, Dražen. "A Bayesian truth serum for subjective data." science 306.5695 (2004): 462-466.
- Witkowski, Jens, and David C. Parkes. "A Robust Bayesian Truth Serum for Small Populations." AAAI. 2012.

APPENDIX

Details for Robust Bayesian Truth Serum

This discussion is based on Rigol and Roth (2016)

Theory and Intuition

In this appendix section we discuss the details of the Robust Bayesian Truth Serum, an intuition for the underlying incentive properties, and our implementation of the payment rule in the field. The following discussion of the model is based on Witkowski and Parkes (2012).

Suppose there is a binary state of the world $t \in (h, l)$ (high, low) representing the entrepreneurial quality of a community member. Agents get a binary signal which is informative of the state of the world.

That is each agent receives a signal $s \in \{h, l\}$ which may represent what they observe about their peer (e.g. they appear responsible, smart etc). Suppose further that all agents share a common prior about the state of the world such that they all agree on the prior probability of a high state, and they all agree on the distribution of signals conditional on the state. Let $p_h = P(s_j = h|s_i = h)$ be the probability an agent assigns to one of his peers receiving a high signal conditional on himself receiving a high signal, and analogously let $p_l = P(s_j = h|s_i = l)$. We say the common prior is *admissable* if $p_h > p_l$, which in English implies that the probability that one's peer receives a high signal is higher if the agent himself receives a high signal. Many natural distributions satisfy this weak requirement.

In order to define the RBTS we must first define the quadratic scoring rule. Let

$$R_q(y, \omega) = \begin{cases} 2y - y^2 & \text{if } \omega = 1 \\ 1 - y^2 & \text{if } \omega = 0 \end{cases}$$

Imagine an agent trying to predict whether some true state ω is 1 or 0. The quadratic scoring rule has the property that his expected score is maximized by reporting his true belief about the probability the state ω is 1 (see e.g. Selten, 1998).

The RBTS is implemented as follows. Every agent states their first order belief (their signal), in a report $x_i \in \{0, 1\}$ (imagine $x_i = 1$ corresponding to $s_i = h$). Further they report their second order belief $y_i \in [0, 1]$ (this is the fraction of the population they believe will report a high signal, $x_k = 1$). For each agent i, assign them a peer agent j, and a reference agent k, and calculate

$$y_i' = \begin{cases} y_j + \delta & \text{if } x_i = 1 \\ y_j - \delta & \text{if } x_i = 0 \end{cases}$$

for arbitrary δ. The RBTS payment for agent i is

$$u_i = R_q\left(y_i', x_k\right) + R_q\left(y_i, x_k\right)$$

The main theorem of Witkowski and Parkes is that under the assumption of an admissable prior and risk neutral agents, there is a Bayes' Nash Equilibrium in which all agents report their first and second order beliefs truthfully.

The intuition behind the payment rule is fairly straightforward. The payment rule has two components. The second component incentivizes the agent to be truthful about his second order beliefs. That

is, the agent is paid via the quadratic scoring rule to predict what some reference agent k will announce as his signal. And by the discussion above, agent i maximizes his expected payment from this component of the scoring rule by truthfully announcing his belief y_i about the likelihood agent k will announce a high signal. In simpler terms, the payment rule rewards agent i for choosing a second order belief as close as possible to the truth (the realized distribution of first order beliefs).

The first component of the payment rule incentivizes the agent to be truthful about his first order beliefs. The term y_i' takes an arbitrary person j's second order belief y_j and either raises or lowers it depending on i's report x_i. RBTS pays agent i $R_q\left(y_i', x_k\right)$, and so i wants y_i' to be as near as possible to the true distribution of responses in the population. The admissibility assumption guarantees that if person j were to know that person i's signal were high, then person j would increase his assessment as to the number of people in the group who received high signals. Likewise, if j were to learn that i's signal were low, j would lower his assessment about the number of people in the group who received high signals. In effect the mechanism raises or lowers j's assessment based on i's report, and then pays i based on the closeness of this modified report to the truth. Thus i can do no better than to tell the truth.

Practical Implementation

We used this payment rule in the field to incentivize rank order responses about members of each group. The model and payment rule, however, were designed for binary responses. Thus while responses contain a rank ordering of 5 people, we treat each ranking as a composite response to 25 yes/no questions of the form "Is person i the highest ranking individual in the group?", "Is he the second highest?" and so on. We elicited second order beliefs of the form "How many people will say person i is the highest ranking individual in the group?" "How many will say he is the second highest?" and so on. From there we directly applied the payment rule, calibrated so that the expected difference between payments arising from truthful and deceptive answers was large. Note that the accuracy of responses across various questions in a single ranking were correlated, but under the assumption of risk neutrality (which is maintained throughout the peer prediction literature and may be empirically reasonable with respect to moderate sums of money), these correlations are irrelevant.

Incentive Compatibility Exercise for RBTS

Throughout the experiment we told respondents that they would maximize their personal payoffs if they reported truthfully. While RBTS is truthful under certain reasonable assumptions about how beliefs are

formed, its incentive compatibility under the empirical distribution of beliefs in practice remains an open question. We therefore evaluate whether respondents are maximizing their subjective expected utility by telling the truth.

Due to the coarseness of our elicited measures of belief, we cannot verify directly whether or not the respondents' priors are admissable. However we can determine the distribution of payoffs respondents can expect to receive under alternative responses to see whether they succeeded in maximizing their subjective expected utility. Respondents' payments depend on the distribution of first order beliefs (i.e. the empirical distribution of responses about the question of interest) and on the distribution of second order beliefs. Therefore, to determine whether truth telling is incentive compatible, we must understand what the respondent believes are the distributions of first and second order beliefs in the population. We obtain the former for free; respondents' beliefs about the distribution of first order beliefs are their second order beliefs, and we elicited these in our survey. We did not, however, elicit their beliefs about the distribution of second order beliefs: their third order beliefs. We must therefore construct those. The intuition behind the construction is presented in the following three steps:

1. We assume that respondents hold a common prior. If so, we can back out their third order beliefs from (a) the distribution of second order beliefs conditional on each first order belief and (b) their belief about how probable each first order belief is. The latter corresponds to her second order beliefs. [23]

2. We approximate the distribution of second order beliefs in the population conditional on any given first order belief with the (sparse) empirical distributions we observe.

3. Given second and third order beliefs, we can calculate a respondent's subjective expected utility from reporting the truth (her stated first order belief) and from any other report.[24] Specifically, we assume that the report the respondent has given is her true belief and calculate her payment. Holding constant her own second order belief and the first and second order beliefs of her peers, we then calculate her payments for every other possible report she could have given.

The results from this exercise are presented in Figure 2 below. Column 1 of the figure depicts the percentage of instances in which telling the truth gives the largest payment, column 2 depicts the percentage of instances that telling the truth results in the second largest payment, etc. Taking the graph at face

[23] If agents have common priors then conditional on the signal they receive, they would update to have the same posterior belief. We stress here that we elicited ranks and not signals. Therefore two agents who report the same rank do not necessarily have the same posterior as the rank is a coarse measure of a signal.

[24] Notice that we only utilize incentivized data since it is only for these data that we collected second order beliefs.

value, telling the truth maximizes the respondent's subjective expected utility in about 35% of instances and it minimizes her subjective expected utility in only about 10% of instances. An ideal graph would place all of its weight in the first column.



Figure 2: Incentive Compatibility of RBTS Using Collected Data

The observed departure from this ideal may be due to the failure of our assumptions required by RBTS holding in practice, or by our noisy approximation of third order beliefs. To evaluate this, we perform a simulation in which we generate data that perfectly abides by all of the assumptions required for the incentive compatibility of RBTS. We generate groups of artificial agents, each of whom holds a common prior and receives a signal about the skill level of their peers. Agents update their priors based on these signals and these form the basis of their second and third order beliefs, each of which we can compute.

Because the data is generated to be perfectly consistent with the assumptions of RBTS, the agent always maximizes his expected utility by telling the truth. Next we compress our simulation data to correspond exactly to the data we collected from our respondents: just first and second order beliefs about group rank. This allows us to have two data sets (collected and simulated) that contain identical level of detail. We then generate the same graph as we did for our collected data and present it in Figure 3.



Figure 3: Incentive Compatibility of RBTS Using Collected and Simulated Data

The graph produced with the collected and with the simulated data are strikingly similar. We therefore conclude that our noisy approximation of third order beliefs could be to blame for the observed weights in columns 2 through 5, and argue that our test yields the strongest evidence in favor of the incentive compatibility that could be achieved via this method. Therefore, telling respondents that they will maximize their expected payments by reporting truthfully may indeed be good advice.

Entrepreneurial Psychology

Impulsiveness:

- I plan tasks carefully.
- I make up my mind quickly
- I save regularly.

Optimism:

- In uncertain times I usually expect the best.
- If something can go wrong for me, it will.
- I'm always optimistic about my future.

- Generally speaking, most people in this community are honest and can be trusted

Locus of Control

- A person can get rich by taking risks.
- I only try things that I am sure of.

Tenacity

- I can think of many times when I persisted with work when others quit
- I continue to work on hard projects even when others oppose me.

Polychronicity:

- I like to juggle several activities at the same time
- I would rather complete an entire project every day than complete parts of several projects.
- I believe it is best to complete one task before beginning another.

Achievement

- Part of my enjoyment in doing things is improving my past performance
- If given the chance, I would make a good leader of people.

Organized person:

- My family and friends would say I am a very organized person