Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Peer Networks and Entrepreneurship: a Pan-African RCT

Fernando Vega-Redondo, Paolo Pin, Diego Ubfal, Cristiana Benedetti-Fasil, Charles Brummitt, Gaia Rubera, Dirk Hovy, Tommaso Fornaciari

Working Paper n. 648

This Version: May, 2019

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Peer networks and entrepreneurship:
a Pan-African RCT*

Fernando Vega-Redondo (Bocconi University & IGIER)
Paolo Pin (Bocconi University & IGIER)
Diego Ubfal (Bocconi University & IGIER)
Cristiana Benedetti-Fasil (European Comission)
Charles Brummitt (Harvard University)
Gaia Rubera (Bocconi University & BIDSA)
Dirk Hovy (Bocconi University & BIDSA)
Tommaso Fornaciari (Bocconi University)

May 2019

*Preliminary results on the experiment discussed in this paper have been presented at the PEDL 2017 Conference and at seminars held at the Paris School of Economics, Queen Mary University, University of Alicante, and the Barcelona GSE Summer Forum. We thank the useful feedback provided in all of them. We also want to acknowledge the strong support of Village Exchange Ghana – and in particular its director Christiane Milev, as well as Kofi Nyalimba and Richard Kpobi – for the small pilot experiment in Ho (Ghana) that started our work in the present project. In that first pilot, the research team included Fernando Vega-Redondo, Cristiana Benedetti-Fasil, Emmanuele Colonnelli, Paolo Pin, and Nicholas Yannelis. At a later stage, as the experiment grew substantially in size and scope, Chris Roth played an important role. And throughout the whole project, we have been fortunate to enjoy superb research assistance. The RA team working in the project during its different phases has included (at different times) the following members: Silvia Barbareschi, Chiara Dall'Aglio, Domenico Fabrizi, Gianluca Gaggioti, Francesca Garbin, Long Hong, Ruggero Jappelli, Giulia Mariani, Nelson Mesker, Laura Montenovo, The Lin Bao Nguyen, Giulia Olivero, Massimo Pulejo, Andrea Pugnana, Ferran Vega-Carol, and Krzysztof Zaremba. We are very grateful to all of them. Last, and certainly not least, we want to than the PEDL program, which provided initial funding for an exploratory project and then a major one for the full-scale experiment described and analyzed in this paper. Throughout, Chris Woodruff, the Scientific Coordinator of this program, has provided both encouragement and valuable advice. Naturally, they are all innocent of any mistakes or misunderstandings that might still remain.

Extended Abstract

Here we report the results of a large RCT conducted at the pan-African level that wants to shed light on the impact of peer effects on innovation and entrepreneurship. The experiment involved around 5000 entrepreneurs (some established, other just aspiring) from 49 African countries. All of those entrepreneurs completed an online business course, while only the treated ones had the additional possibility of interacting with peers, within groups of sixty, and in one of three different setups: (a) ***face-to-face***, (b) ***virtually "within"*** (where interaction was conducted through an Internet platform in groups of entrepreneurs of the same country), (c) ***virtually "across"*** (where the virtually connected groups displayed a balanced heterogeneity across countries). After two and a half months, all participants were asked to submit business proposals. The ones submitted were then evaluated in a two-stage procedure. First, they were graded by a panel of African professionals; subsequently, the pool of highest-graded proposals were again assessed and graded by senior investors, who selected some for possible funding. Two outcome variables follow from this evaluation exercise: the (optional) decision of whether to submit a proposal, and the grades (1 to 5) obtained by the proposals that were submitted.

Next, we outline our main results concerning the ***effect of the treatment*** on the two aforementioned outcomes – ***submission*** and ***quality*** (measured in the intensive margin) – as well as the combination of both of them that we call, for short, ***extensive quality***.

(1) ***Virtual-within interaction*** has a positive and significant treatment effect on the three dimensions: submission, intensive quality, and extensive quality. Instead, when interaction is face-to-face (thus also "within") only submission and the extensive quality margin are affected (positively so).

(2) ***Virtual-across interaction*** yields no significant effect on any of the former three dimensions.

(3) When effective on quality (cf. (1)), the ***treatment operates*** by shifting up, on average, the evaluation grade of business proposals from low levels (grades 1 or 2) to high ones (grades 4 or 5).

(4) The ***baseline quality*** of entrepreneurs has a positive effect on performance. However, the average such quality of the peers in one's own group has a ***negative composition effect*** on intensive quality. In fact, a similarly negative effect is also induced by peers' average ***experience*** level.

(5) As a ***robustness test***, the core treatment effects described in (1)-(2) are confirmed to remain essentially unchanged under a full range of control (baseline) variables, while the composition effects identified in (4) are found to survive a standard placebo test.

As a second step in the analysis, we construct a social network in each group by defining the weigh of a directed link between two entrepreneurs as the amount of information (overall size of messages) written by one of them for which there is evidence that the other has been exposed to, then writing a subsequent message. Then, on the basis of the network structure so defined, we estimate the induced peer effects and arrive at the following conclusions.

(6) In ***large countries*** (the only ones for which a sufficient number homogeneous groups can be formed), ***virtual-within interaction*** leads to positive and significant peer effects on submission and extensive quality, but not intensive quality. Instead, when entrepreneurs of large countries are exposed to virtual-across interaction, no significant peer effects arise in any of the three outcomes.

(7) In the set of ***small countries***, where only ***virtual-across interaction*** is possible, there are positive and significant peer effects on both extensive and intensive quality but not on submission.

(8) ***Composition effects*** on network peers are weak, largely captured by (outcome-based) peer effects.

(9) Results (6)-(7) are ***structurally robust*** to redefining the network links in the following two ways:
(a) they are limited to involve less than a maximum communication lag, suitably parametrized;
(b) they are two-sided, their weight tailored to the flow information channeled in both directions.

A combined consideration of (1)-(9) reveals an interesting contrast between treatment and peer effects. For example, in view of (1)-(3), we may conclude that whereas some group homogeneity – or face-to-face contact– bring about positive treatment effects, the group heterogeneity induced by virtual-across interaction fails to deliver significant such effects on all three dimensions. Instead, (6)-(7) indicate that network-based peer effects deliver an intriguingly different pattern. For, under virtual-within interaction, we find that entrepreneurs' peers exert a significantly positive influence on submission (and the extensive margin) but not so on quality *per se* (in the intensive margin, while a somewhat polar behavior arises in small countries who undergo virtual-across interaction. This suggests that whereas homogeneity leads to peer interaction that is rather independent of peer performance, heterogeneity has peer performance play an important role (both in positive or negative terms, depending on the quality of that performance). Overall, this induces an effect of the treatment that is significantly positive under homogeneity (virtual-within interaction for large countries) but not strong enough to be significant under full-fledged heterogeneity (virtual-across interaction for small countries).

The aforementioned contrast between the nature and implications of the treatment effects stated in (1)-(4) and the network peer effects in (6)-(8) is interesting and deserves further investigation. A possible explanation for it might hinge upon the positive role that homogeneity/familiarity may play as a source of encouragement (and hence participation), as opposed to the negative impact it could have in reducing the novelty of ideas and/or highlighting the fear of competition (thus dis-incentivizing information sharing and thus a genuine effect induced by peer performance). To gain a good understanding of these issues, however, one needs the help of theory as well as a detailed investigation of how communication actually unfolds in our context. Both lines of work are part of our ongoing research. Here, we provide a preliminary account of the latter, which is included in the final part of the paper.

Our approach to semantic analysis relies on the machine-learning tools developed by the modern field of Natural Language Processing (NLP). This methodology is applied to the vast flow of information exchanged by entrepreneurs (over 140,000 messages) in order to identify, first, what have been the modes/categories of peer communication more prevalent in our context, e.g. business focus, sentiment/encouragement, target audience, etc. Then we use this information to understand what are the different patterns of communication most prevalent in our context, as captured by a corresponding set of conditional and unconditional distributions that show and how communication is associated to: (a) endogenous variables such as behavior or performance; (b) exogenous variables, such as treatment type or individual baseline characteristics.

The main conclusions obtained so far can be summarized as follows. Messages are quite polarized in either the business or sentiment dimension, showing an inverse dependence in the (strong) FOSD sense between the respective distributions. Applying the same comparison criterion, we also find that highly performing agents use more business-focused messages, which are not only neutral in sentiment but also targeted to specific peers (rather than being general messages). Interestingly, however, the treatment arm (virtual-within or -across) has no significant effect on the type of communication, while baseline quality and a measure of "motivation" do have an effect analogous to that described before for performance.

Finally, we also rely on the message categorization induced by the NLP analysis to construct semantically weighted networks on two specific features/categories: *business relevance* and *sentiment*. Quite remarkably, the corresponding peer effects are found to be unaffected by either of these "semantic projections" of the social network. This suggests that, even though entrepreneurs' messages focus heavily on business issues, their communication displays a feature that is often observed in ordinary (non-virtual) interaction: there is a balance between business focus and a comparable amount of sentiment-laden talk.

Keywords: Social networks, peer effects, peer networks, entrepreneurship, semantic NLP analysis.

JEL classif. codes: C93, D04, D85, O12, O31, O35.

Contents

1 Introduction **7**

2 Experimental design and main objectives **13**
2.1 Recruitment . . . 14
2.2 Baseline survey . . . 14
2.3 Treatment . . . 16
2.4 Randomization . . . 17
2.5 Online course . . . 19
2.6 Timeline . . . 20
2.7 Outcomes . . . 21

3 The core treatment effects **22**
3.1 Submission . . . 23
3.2 Business quality . . . 25
3.2.1 The extensive margin . . . 25
3.2.2 The intensive margin . . . 27
3.3 Peer-composition effects . . . 29
3.4 The second-stage evaluation . . . 32
3.5 Robustness . . . 33

4 Network analysis **34**
4.1 Network construction . . . 34
4.2 Network peer effects . . . 38
4.2.1 The identification problem: reflection and homophily . . . 38
4.2.2 The core peer effects . . . 40
4.3 Robustness . . . 43
4.3.1 Network-composition effects . . . 43
4.3.2 Parametrizing the communication lag . . . 46
4.3.3 Two-sided communication flows . . . 46

5 The semantic analysis **51**
5.1 Natural Language Processing: the annotation setup . . . 51
5.2 Communication patterns . . . 53
5.2.1 Semantic distributions . . . 53
5.2.2 Conditional relationships . . . 57
5.2.3 Semantic indicators . . . 73
5.3 Semantically weighted networks . . . 75

6 Summary, conclusions, and ensuing research **77**

Appendix: supplementary material **83**
A.1 Balance checking . . . 83
A.2 The core treatment effects: full set of controls . . . 87
A.3 Peer composition: robustness check . . . 90
A.4 Peer effects: complete results . . . 92
A.5 Peer effects: semantically weighted networks under bidirectional influence . . . 96

List of Tables

1 Summary statistics, full sample . . . 15
2 Treatment effect on the submission of business proposals . . . 24
3 Treatment effect on business quality: extensive-margin analysis . . . 26
4 Treatment effect on business quality: intensive-margin analysis . . . 27
5 Marginal quality effects (ordered Probit), intensive-margin, M0 completers . . . 30
6 Effects of peer composition . . . 31
7 Second-stage outcomes . . . 32
8 Network statistics: description . . . 37
9 Network statistics: correlations . . . 38
10 Network peer effects, benchmark setup . . . 42
11 Network-composition effects: baseline business quality . . . 44
12 Network-composition effects: business experience . . . 45
13 Network peer effects, bidirectional influence . . . 50
14 Semantic node-based statistics . . . 55
15 Semantic message-based statistics . . . 57
16 Cross-category correlations: message-based approach . . . 58
17 Cross-category correlations: entrepreneur-based approach . . . 58
18 Semantics-performance relationship: p-values on distributional hypotheses . . . 65
19 (b) Semantic-characteristic relationship: p-values on distributional hypotheses . . . 69
20 Network peer effects, business network . . . 76
21 Network peer effects, sentiment network . . . 76
22 bilateral interaction flows: benchmark, business & sentiment networks. . . . 77
A1 Balance check, Uganda sample . . . 84
A2 Balance check, large-country sample . . . 85
A3 Balance check, small-country sample . . . 86
A4 Replicating Table 3 (extensive margin), with controls . . . 87
A5 Replicating Table 4 (intensive Margin), with controls . . . 88
A6 Replicating Table 5 (marginal-quality effects), with controls . . . 88
A7 Replicating Table 7 (second-stage outcomes), with controls . . . 89
A8 Effects for peer composition (I): placebo for control group . . . 90
A9 Effects of peer composition (II): surprise component . . . 91
A10 Network peer effects, benchmark setup, complete results on submission . . . 93
A11 Network peer effects, benchmark setup, complete results on extensive quality . . . 94
A12 Network peer effects, benchmark setup, complete results on intensive quality . . . 95
A13 Network peer effects, business network under bidirectional influence . . . 96
A14 Network peer effects, sentiment network under bidirectional influence . . . 96

List of Figures

1 **Age distribution of the overall entrepreneur population.** 16

2 **Schematic representation of experimental design** 19

3 **Structure of the online entrepreneurship course.** 20

4 **Peer network for a virtual-within interaction group based in Nigeria.** 36

5 **Peer network for a virtual-across interaction group where red nodes correspond to entrepreneurs from large countries whereas those colored green correspond to entrepreneurs based in small countries.** 36

6 **Peer effects on submission, across reaction lags** 47

7 **Peer effects on business quality in the extensive margin, across reaction lags** 48

8 **Peer effects on business quality in the intensive margin, across reaction lags** 49

9 **Node-based densities over semantic categories** 54

10 **Message-based densities over semantic categories.** 56

11 **Message-based densities of business content and audience cond. on sentiment** 59

12 **Message-based cdf's on business content and audience conditional on sentiment** . 61

13 **Message-based cdf's on sentiment and audience conditional on business content** . 62

14 **Message-based cdf's on business content and sentiment conditional on audience** . 63

15 **Message-based density of sentiment conditional on business content** 64

16 **Entrepreneur-based cdf's over business content conditional on performance** 66

17 **Entrepreneur-based cdf's over sentiment conditional on performance** 67

18 **Entrepreneur-based cdf's over individual audience conditional on performance** . . 68

19 **Entrepreneur-based cdf's over semantic categories conditional on treatment arm** 70

20 **Entrepreneur-based cdf's over semantic categories conditional on M0 complete** . 71

21 **Avge. business focus of entrepreneur cond. on age, gender, and baseline quality** . 72

1 Introduction

Peer networks play a prominent role in many different social and economic contexts. They are essential, for example, to understanding how firms, markets and, generally, large production and financial systems operate.[1] They are, of course, particularly important as well for innovation and entrepreneurship,[2] for which successful peer networks typically involve the following rich set of peer-based phenomena:

1. *cooperation* – peers cooperate and therefore must trust each other;
2. *learning* – new information is gathered from peers as it diffuses through the network;
3. *competition* – peers compete for funding, and often operate in related markets as well.

There is, in the network literature, a substantial amount of research that studies ***separately*** each of the above dimensions of the problem, mostly from a theoretical viewpoint.[3] But, to the best of our knowledge, there is no existing body of research that jointly studies the interplay between the three of them. We argue, however, that in order to understand (and hence model) how peer networks operate on innovation and entrepreneurship it is of fundamental importance to account for such an interplay. The issue, however, is that, being the problem at hand so rich and complex, it is difficult to identify what are the main forces at work and then be in a position to develop a satisfactory approach. How do agents (attempt to) reconcile their incentives to collaborate with the fear of favoring a potential competitor? How important is it for agents to enjoy broad and flexible channels of communication? How does the structure of the network (itself emerging rather than imposed) bear on both the diffusion of information and the support of trust? What interaction rules and communication protocols render interaction most effective?

To address at least some of the previous questions, we need a solid modeling approach, which in turn requires a comparably solid empirical evidence to build upon. Since, as mentioned, such an evidence is not yet available, to gather it systematically has been one of the primary motivations of our present research. The aim, that is, has been to study in detail how, in a large and rich environment, the key phenomena discussed above shape a process of peer-based innovation and entrepreneurship. More specifically, we wanted to observe and document the following evidence:

(a) how, over time, entrepreneurs search for useful peers;
(b) what sort of communication they undertake and how does their behavior evolve;
(c) what effect (a) and (b) have, in the end, on their innovation performance.

To this end, we have conducted a randomized control trial (RCT) involving a large population of African entrepreneurs (almost 5,000 of them) from all over the continent (49 countries represented). The experiment

[1] The recent Handbook on the topic of social networks edited by Bramoullé, Galeotti, and Rogers (2016) provides a good panoramic view of the state of the art. The reader may check, specifically, the contributions by Dessein and Pratt (2016) for a discussion on how network structure affects the operation of firms and other organizations; by Beaman (2016) on how social connections impinge on hiring and employment in labor markets; by Breza (2016), on how informal networks strive to compensate for poorly performing institutions in less-developed economies; by Acemoglu, Ozdaglar and A. Tahbaz-Salehi (2016), as well as Cabrales, Gale and Gottardi (2016), on how the network pattern of connections (input-output relationships in the first case, assets and liabilities in the second) determine the overall systemic risk faced by large economies.

[2] On this matter, see below for a discussion of relevant literature.

[3] On cooperation in social networks, we may refer to the papers by Karlan, Mobius, Rosenblat, and Szeidl (2009), Lippert and Spagnolo (2011), and Jackson, Rodriguez-Barraquer, and Tan (2012); on learning in networks, those of Bala and Goyal (1998), DeMarzo, Vayanos and Zwiebel (2003), and Golub and Jackson (2010); on competition/conflict in networks, Franke and Ozturk (2015), König, Rohner, Thoenig, and Zilibotti (2017), and Heijnen and Soetevent (2018).

itself lasted for about two and a half months. During this period, the whole population (treated and control alike) underwent an online business/entrepreneurship course, specifically tailored to the nature of the experiment. In parallel to it, the treated subpopulation also interacted in peer groups consisting of 60 entrepreneurs, according to three different sub-treatments. In one of the treatment arms (implemented in Uganda), groups interacted face-to-face. In the other two, they interacted virtually through an Internet-based chatting platform, which was designed to allow for flexible communication, providing the entrepreneurs' full discretion on how to modulate privacy and organize content.

The two virtual arms differed in the composition of the groups. In the treatment arm we labeled "virtual-within" every group was nationally homogeneous, i.e. all its members belonged to the same country. Instead, in the other arm, labeled "virtual-across," groups were formed randomly in a nationally heterogeneous (and balanced) manner. Whereas the virtual-within arm was circumscribed to entrepreneurs from the five countries in the sample with the largest representation (Nigeria, Ghana, Kenya, South Africa, and Tanzania), the virtual-across arm involved individuals from all African countries represented in our population.

All along the interaction (which lasted for slightly more than two months), we recorded the evolution of the process exhaustively, including the full networking activity and the whole set of messages exchanged in the virtual treatment arms. Then, at the end of the intervention, all entrepreneurs – treatment and control alike – were asked to submit a business proposal. This in turn allowed us to obtain a measure of innovation performance through the following two-stage procedure. First, all submitted proposals were evaluated and graded in a 1-5 scale by a panel of 15 African professionals belonging to investment firms, entrepreneurship hubs, and accelerators. Then, in a second stage, the business proposals that ranked highest in the first stage (together with a smaller additional set selected by a European academic panel) were again evaluated, ranked, and some funded by a panel of 24 senior investors (VCs, investment-fund managers, and angel investors), who acted as the financial partners of our project.

The main results can be succinctly advanced as follows.

First, consider those concerning the ***effect of the treatment*** on the two aforementioned outcomes – ***submission*** and ***quality*** (as measured by the aforementioned evaluation panels) – as well as on the combination of both of them into a single variable that we call, for short, ***extensive quality***.

(1) ***Virtual-within interaction*** has a positive and significant treatment effect on the three dimensions: submission, (intensive) quality, and extensive quality. Instead, when interaction is face-to-face (thus also "within") only submission and the extensive quality margin are affected (positively so).

(2) ***Virtual-across interaction*** yields no significant effect on any of the former three dimensions.

(3) When effective (cf. (1)), the ***treatment operates*** by shifting up, on average, the evaluation grade of business proposals from low levels (grades 1 or 2) to high ones (grades 4 or 5).

(4) The starting ***baseline quality*** of entrepreneurs has a uniformly positive effect on performance. However, the average baseline quality of group peers have a negative composition effect on intensive quality. In fact, a similarly negative composition effect is also induced by peers' average experience level.

(5) As a test of ***robustness***, the core treatment effects described in (1)-(2) are confirmed to remain essentially unchanged under a full range of control (baseline) variables, while the composition effects identified in (4) are found to survive a standard placebo test.

As a second step in the analysis, we have constructed a social network in each group in which a directed link is postulated to exist between two entrepreneurs if there is evidence that one of them has written a message that the other party has been exposed to, after which the latter has written back on the same "channel." In that case, the weight of the link is simply identified with the amount of information (overall size of the messages) triggering all such communication instances between the two entrepreneurs. Then, on the basis of the network structure thus constructed, we estimate the induced peer effects and arrive at the following conclusions.

(6) In ***large countries*** (i.e. the only ones for which a sufficient number of nationally homogeneous groups can be formed), ***virtual-within interaction*** leads to positive and significant peer effects on submission and extensive quality, but not on intensive quality. Instead, when entrepreneurs of large countries undergo virtual-across interaction, no significant peer effects arise.

(7) In ***small countries***, the alternative (complement) case, ***virtual-across interaction*** (the only context that happens to relevant for these countries) induces positive and significant peer effects on both extensive and intensive quality but not on submission.

(8) Neither ***peers' baseline quality***, nor their ***experience***, have any significant effect on performance.

(9) The results in (6)-(7) are robust to redefining the links of the network in two alternative ways:
(a) they cannot involve a communication lag higher than some parametrized upper bound;
(b) they are two sided, their weight reflecting the amount of information flowing in both directions.

A combined consideration of (1)-(9) reveals an interesting contrast between treatment and peer effects. Specifically, in view of (1)-(3) we may conclude that whereas some group homogeneity – or face-to-face contact– brings about positive treatment effects, the group heterogeneity induced by virtual-across interaction fails to deliver significant such effects on all three dimensions. Instead, (6)-(7) indicate that network-based peer effects deliver a quite different pattern. For, under virtual-within interaction, we find that entrepreneurs' peers exert a significantly positive influence on submission and the extensive quality margin but not so on quality *per se*, i.e. on the intensive margin. A somewhat polar behavior arises in small countries (whose treated entrepreneurs only "interact across" in heterogeneous groups). These also enjoy positive and significant peer effects on extensive quality, but this is now a reflection of the impact of peer interaction on intensive quality, *not* submission. Finally, there is the third group of entrepreneurs, consisting of those belonging to large countries who interact within heterogeneous groups. In a sense, for these entrepreneurs, the forces that, in the former two contexts, rendered peer effects non-significant on, respectively, submission and intensive quality appear to play a preeminent role. This, in the end, combines into peer effects on the remaining outcome dimension, extensive quality, being non-significant as well.

The contrast between the nature and implications of the treatment effects stated in (1)-(4) and the network peer effects in (6)-(8) is intriguing, so we shall put forward some conjectures about the possible forces at work. In a nutshell, our suggestion can be described as follows: while the (national) homogeneity/familiarity of the virtual-within treatment may be a good basis for encouragement (and hence participation), it may also be detrimental to the novelty of ideas that can be shared and/or increase the fear of competition (thus dis-incentivizing information sharing in the first place). In contrast, we shall propose, the opposite forces could be at work in the virtual-across treatment, with a particularly strong bearing on entrepreneurs from small countries, where the considerations involved seem strongest. Clearly, in order to gain a proper understanding of these issues one needs the help of a suitably formulated theory as well as a detailed investigation of the actual process of interaction/communication taking place among entrepreneurs. Both of these – the development of a suitable ***theoretical framework*** and a thorough ***semantic analysis*** of the information flows – are the primary objectives of two separate companion pieces, still in the making. In the final part of the paper, however, we provide a preliminary account of the latter route, which is based on the application of the machine-learning tools developed by the field of Natural Language Processing (NLP).

First, we start by describing the NLP methodology, as we have applied to the vast amount of information (over 140,000 sentences/messages) exchanged throughout the experiment. The main purpose of this first stage of the analysis is to identify what features of interest (so-called modes or sentiments) characterize peer communication in our context. Then we turn to the question of which of those features are statistically associated to

(a) various baseline characteristics and treatment arms (exogenous conditions);

(b) different type of behavior and performance (endogenous variables).

We find that messages are quite polarized in either the business or sentiment dimension – i.e. they are markedly identified to be of one type or the other, showing an inverse dependence in the (strong) FOSD sense between the respective business-based and sentiment-based distributions. Applying the same comparison criterion, we also find that high performing agents use more business-focused messages, which are not only more neutral in sentiment but also targeted to specific peers (rather than being general messages). Interestingly, it turns out that the treatment arm (virtual-within or -across) has no significant effect on the type of communication – that is, both of them induce an essentially equal distribution over semantic content. Instead, individual characteristics such as baseline quality and an indirect measure of "motivation" do have an effect analogous to that described before on performance: entrepreneurs displaying those characteristics are more business-focused and the their messages are usually targeted to specific peers.

Finally, we rely on the NLP-induced semantic identification to construct alternative representations of the peer network, each differing in what alternative features (in particular, business relevance or sentiment) are used to compute the weights of the network links. An interesting observation in this respect is that semantically weighted networks where either ***business relevance*** or ***sentiment*** are the alternative focal features yield very similar estimates of the peer effects. As we then find out, the simple reason underlying this observation is that, as indeed happens is so many realms of life, entrepreneurs typically mix "business talk" and "sentiment-laden talk" in a rather balanced manner. In future work, an important objective of a refined semantic analysis will be to refine our construction of the peer network, extracting more precisely the actual flows of bilateral information and feedback from the free-format chatting data.

We end this introductory section with a brief review of related literature. One of the leading ideas underlying the fast development of the field of social networks has been that inter-agent connections are an important conduit for sharing information, and hence they are important as well in promoting innovation (see e.g. Granovetter (1973), Burt (1992, 2004)). For the specific case of firm-based innovation in the marketplace, there is ample empirical evidence showing that much of the R&D conducted in the most dynamic sectors of the economy is carried out through inter-firm collaborations. As a very small account of this large literature, good concrete illustrations can be found in Zhao and Aram (1995), Powell, Koput, and Smith-Doerr (1996), Edquist, Eriksson and Sjögren, (2000), or Dahl and Pedersen (2005), while for a panoramic and aggregate view of the phenomenon, Hagedoorn (2002) is a good empirical source.

The previous papers mostly focus on inter-firm research collaboration that aims at the development of new products or better production processes. But, of course, interaction among firms can also serve them to improve management practices, as studied by Fafchamps and Quinn (2016) or Cai and Szeidl (2017). The first paper reports on a field experiment conducted in three African countries (Ethiopia, Tanzania and Zambia), while the second one discusses an experiment involving Chinese entrepreneurs. Both of them find that, indeed, peer-manager interaction improves management practices. And in the case studied by Cai and Szeidl, the authors also document a positive effect on profitability and growth.

Naturally, peer effects may also involve, rather than established firms, aspiring entrepreneurs who are exploring fresh business ideas in the hope of starting up a *new* firm. We are not aware of much literature concerned with this important context. Interesting exceptions are provided by the papers of Nanda and Sorensen (2010), Lerner and Malmendier (2013), and Hacamo and Kleiner (2018). The first paper studies, for an exhaustive data set on Danish workers, the influence exerted by peers on the probability that any given worker may create a firm thereafter – here, two individuals are declared to be peers if they have been been co-workers in a firm at some point during the preceding years. The latter two papers, on the other hand, ask an analogous question concerning the MBA students of two different American universities. There, peers are defined to be those students who, upon entrance, are (randomly) assigned to the same section or cohort. The results reported in these three papers are not perfectly aligned. While the first one reports a positive peer effect on average, the latter two find that *experienced* students (i.e. those with some previous management record) induce a negative effect on the probability of firm creation by their peers. Further analysis of the data suggests that such a negative impact can be attributed to the discouragement effect resulting from the following simple observation: a good fraction of experienced students have gone through a prior *unsuccessful* venture. In fact, our experiment identifies a somewhat related negative effect by experienced peers. However, as we shall explain, this composition effect is conceptually quite different from the one that is channeled through an endogenous peer network and an endogenous choice of actions, both of which are our main concern in this paper.

Another interesting context where peer effects are also expected to be at work is that of the so-called "ecosystem accelerators" – a relatively recent type of organizations, arising worldwide, whose objective is the support of early entrepreneurship. Despite their growing importance, however, there have been only a few papers studying them systematically. An interesting example is provided by the paper of Gonzalez-Uribe and Leatherbee (2018). They focus on a Chilean accelerator, Start-Up Chile, for which they have data of circa 1,000 different start-up ventures, spread over a period of five years. The paper shows that the basic services provided by the accelerator (including, of course, the peer networking facilitated by its co-working facilities) has a significant effect on the ensuing business performance of entrepreneurs *only if* combined with formal training. This conclusion resonates well with our work. For it is precisely an analogous complementarity, conjectured to be important also in our case, that has motivated the incorporation of a substantial training component in our experimental design.[4]

To motivate an additional key feature of our design – its outcome variables – we now refer to the literature concerned with understanding the relative merits of alternative proposal evaluation schemes. Recall that

[4]Even though, as discussed in the survey by McKenzie and Woodruff (2014), the effects of training on performance reported by the literature are at best weak, those authors indicate that there is "stronger evidence that training programs help prospective owners launch new businesses more quickly." This is consistent with the fact that, both in the context studied by Gonzalez-Uribe and Leatherbee (2018) as well as ours, since many of the entrepreneurs are involved in start-up ventures, business training should have a significant impact.

the key outcome variable that measures the quality of business proposals is obtained through a two-stage procedure. First, the submitted proposals are assessed and ranked by junior professionals of the African entrepreneurial ecosystem, then those identified to be among the best are evaluated (and possibly funded) by the senior investors. This procedure raises the question of how relevant such evaluation may be as an indicator of economic potential and firm survival. Two recent papers, Fafchamps and Woodruff (2017) and McKenzie and Sansone (2017), study this issue by contrasting in specific cases the performance of two alternative evaluation methods: surveys and expert panels.[5] Their essential conclusion is that neither of the methods does a good enough job at "picking winners."[6] Why do we then use panel rankings as an outcome? The main reason is that, in the minds of the entrepreneurs, their short-run objective is crystal clear: to get funding for their project. Thus, in this sense, it is reasonable to argue that their "innovation efforts" will be largely steered in this very direction, i.e. towards matching the criteria they anticipate will be used by the investors to select their "winners." Whether or not such criteria are good at predicting economic success is, of course, a very important question, as emphasized by the two aforementioned papers. We an abstract from them here, however, since the primary focus of our research is to understand how peers strive to improve their *perceived* business quality, as this is reflected by investor evaluation criteria and hence translates into funding success.

Finally, we close this literature review by revisiting the basic question of why is it that peer interaction affects innovation. One reason for it has already been discussed: peers can be a source of *motivation*, although (as we have discussed) in some cases they can also operate negatively by inducing discouragement. Another channel through which innovation can also be promoted is diversity. For, in some cases, social networks allow individuals to access and combine information, skills, and different approaches to problems in ways that expand substantially the range of possibilities that would be available otherwise. But does this mean, at least as far as innovation is concerned, that maximum group diversity is optimal? Again, the empirical literature documents that there can be various effects at work, offsetting each other, and sometimes leading to a non-monotonic relationship. That is, in general, while some diversity is welcome, too much of it can be detrimental. As some illustrative (and, in line with the present theme, also "diverse") set of examples, we can mention the following papers: Lungeanu and Contractor (2015), who study a large collection of research teams world-wide in the medical field of oncofertility; Letaifa and Rabeau (2013), who focus on the large ecosystem of Canadian ICT firms in the Montreal area; Dayan, Ozer, and Almazrouei (2017), who consider the product development teams of manufacturing firms operating in the Ankara area; or Tavassoli and Carbonara (2014), who study the comparative innovation performance of the 81 different areas categorized in Sweden as distinct "functional regions." In all these empirical analyses, a suitable version of the aforementioned trade-off between the positive and negative effects of diversity is encountered, highlighting the complex way in which this feature enters into the social (peer-based) process of innovation. In this paper, we study very partially this issue by considering how the national diversity of groups affects, on average, its entrepreneurial performance. Of course, much more needs to be done in this important respect, which by its very nature is very multisided and therefore complex.

[5]Fafchamps and Woodruff (2017) focus on the business plans submitted to a business competition they ran in Ghana. McKenzie and Sansone (2017), use data from a very large business competition (YouWin!) run by the Nigerian government in collaboration with the World Bank. Besides the two methods mentioned in the text, McKenzie and Sansone also consider machine-learning approaches.

[6]Whereas Fafchamps and Woodruff (2017) find some complementary predictive power to both methods, it is a quite weak one. Instead, McKenzie and Sansone find that some survey variables are weakly correlated with firm performance, panel decisions display essentially no correlation.

The rest of the paper is organized as follows.

In Section 2 we describe the different components of the experimental design and anticipate our main objectives.

Section 3 focuses on the the estimation of the treatment effect on the three performance dimensions considered here: submission, extensively-measured quality, and intensively-measured one. In the latter case, it undertakes a detailed discussion of the identification problem induced by possible selection bias on the set of submitting entrepreneurs. It also studies peer-composition effects and revisits the estimation exercise from the viewpoint of the project evaluation conducted by VCs and investors in the second stage of the experiment. This section ends with the analysis of some robustness issues.

In Section 4 we turn to studying the phenomenon of peer interaction from a network perspective. We start this section by formulating an operational specification of the social (peer) network. Then, before proceeding to the estimation of peer network effects, we describe and address the important econometric identification problem posed by the so-called reflection problem first raised by Manski (1993). Concerning the estimation itself, we again decompose it between the core peer effects and the composition effects. Finally, we address two robustness issues that concern the operationalization of the network: first, we contemplate bounding the maximum communication lag; second, we turn to measuring inter-agent influence through the two-way flow of information exchanged by every pair of entrepreneurs in both directions.

Section 5 closes the analysis of the paper by conducting a preliminary semantic analysis of the communication (messages/posts) channeled through the Internet platform implementing virtual peer interaction in our experiment. First, we describe the basic annotation setup that codifies patterns and meaning. Then, we relate such patterns/meaning to exogenous variables (baseline characteristics), interaction conditions (treatment arms), and endogenous variables (outcomes and networking). Finally, we rely on the semantic categorization of messages resulting from the aforementioned annotation procedure to construct and analyze semantically-projected networks, then estimating the sign and magnitude of the peer effects induced.

Finally, Section 6 ends the main body of the paper with a summary, conclusions, and description of ongoing research. Additional material (in particular, secondary tables) is included in the Appendix.

2 Experimental design and main objectives

For the sake of clarity, our description of the experiment is organized into various subsections, each of them separately focusing on one of its several constituent parts. We start by explaining how recruitment of entrepreneurs was conducted, then describe the baseline survey, the treatment, the randomization, the online course, the time line of the experiment, and the range of outcomes obtained. As important partners in the project, we involved a team of 15 young professionals working in the the entrepreneurship ecosystem of various African countries, and a panel of 24 senior investors (mostly VCs and angel investors). The role of both groups will be described below in some detail. The overall intervention was named *Adansonia*, a reference to a widely quoted African proverb that suggests the importance of "peer interaction" in the generation of knowledge. [7]

[7] A Ghanaian version of the proverb reads: "Wisdom is like a baobab tree; no one individual can embrace it." Adansonia is the *genus* of the trees generally knows as baobabs.

2.1 Recruitment

The entrepreneurs were recruited from all over Africa through a variety of different methods. The main ones were the following.

- **Social media**: We implemented a massive campaign on social media – mainly, Google and Facebook – that was targeted to reach either aspiring or operating entrepreneurs throughout the continent. These campaigns were based on a selected set of profiles, keywords, and sites, where our intervention was advertised.
- **Viral campaign**: Registered participants were encouraged and incentivized to bring further participants into the experiment. The incentives were induced through a lottery to participate in a two-week program of different events (training, networking, and pitching) that were organized at Bocconi University at the end of the experiment. Further details on this program are provided below.
- **Information events on the ground**: A set of face-to-face events were organized in universities and entrepreneurial hubs where local partners explained the intervention and encouraged established entrepreneurs and aspiring ones to join.

In the end, our recruitment efforts led to almost five thousand participants (4958) from 49 different African countries. They were typically young and well educated individuals spreading quite widely over different sectors and demographic characteristics. Information on these characteristics was collected by an extensive baseline survey, as discussed in the next subsection.

2.2 Baseline survey

Prior to the intervention, and a prerequisite to be part of it, all recruited individuals were asked to complete a survey consisting of 92 questions which took around 30 minutes to complete.

Table 1 below provides a selected summary of the baseline information gathered from this survey. We may divide it as follows.

- **Demographics**: Geographically, the population was quite disperse throughout the continent, although there were some significant biases as well. For example, Northern Africa was hardly represented in our full sample (most of the population originated from Sub-Saharan countries) and there was a larger representation of West Africa as compared to East Africa, even though each part of Africa is quite similarly populated. Slightly over 30% of the population were female and 63% already had a business. They were also generally well educated (more than 80 % had completed university education), about 1/3 were married, and they were quite young as well (the average age was around 30 years old). For a graphic representation of the age distribution in the population, see Figure 1.

Table 1: **Summary statistics, full sample**

	(1)	(2)	(3)
	Mean	S.D.	Obs.
Panel A. Demographics			
Northern Africa	0.01	0.11	4958
Western Africa	0.53	0.50	4958
Eastern Africa	0.35	0.48	4958
Southern Africa	0.07	0.25	4958
Middle Africa	0.04	0.19	4958
Female	0.31	0.46	4958
Has Business	0.63	0.48	4958
Age	30.96	7.79	4958
University Complete	0.82	0.39	4958
Married	0.38	0.49	4939
Panel B. Business Idea			
Has Idea about Existing Business	0.38	0.48	4958
Has Idea about New Business	0.55	0.50	4958
Has no Idea yet	0.08	0.27	4958
Sector: Agriculture	0.25	0.43	4958
Sector: Services	0.17	0.37	4958
Sector: Technology	0.13	0.34	4958
Sector: Manufacture	0.09	0.29	4958
Sector: Social Entrepreneurship	0.12	0.33	4958
Sector: Retail	0.05	0.21	4958
Has written business plan	0.59	0.49	4958
Participated in business competition	0.37	0.48	4958
Has employees	0.36	0.48	4947
Panel C. Financial Access			
Saves at a Bank	0.90	0.30	4958
Got Bank Loan for business	0.09	0.29	4958
Prefers equity debt to loans	0.45	0.50	4958
Prefers loans to equity debt	0.16	0.37	4958
Prefers either equity or loans	0.34	0.48	4958
Panel D. Labor Market Outcomes			
Reservation Wage (in USD)	1604	2102	4914
Years of work experience	5.11	3.27	4958
Has a job	0.55	0.50	4958
Panel E. Networks			
Number of People discuss business	4.63	3.37	4958
Prefers to discuss with different sector	0.11	0.31	4958
Prefers to discuss with different gender	0.13	0.33	4958
Prefers to discuss with different country	0.18	0.38	4958
Panel F. Personality Traits			
Risk Aversion (choice among 6 lotteries)	3.44	2.05	4948
Trust Measure (0 to 10)	4.81	2.77	4957
Position in your country: current (0 to 10)	4.82	1.63	4650
Position in your country: expected (0 to 10)	7.84	1.55	3642
Position in your country: desired (0 to 10)	9.06	1.54	4577

Notes: The table uses values of the variables collected in the online application form completed during May 2017.

- **Business profile**: Most of the subjects entered the experiment with some business idea – for more than half (55%) this idea was about a new business, while for 38% it was for an existing business. Concerning the sectoral distribution, agriculture, service, and technology were the dominant ones, in that order. On financial matters, most of them (90%) had savings in a bank account, but few (9%) ever got a loan from a bank. Concerning their desired source of funding, a sizable fraction (45%) preferred equity-based funding to loans, although there were also many (34%) that would be happy with either of them. They had an average of slightly more than 5 years of experience and 55% were employed at the time.
- **Networks**: The average number of people with whom our subjects typically discussed business topics was 4.6. To identify their taste for diversity, they were asked the profile of the person they would prefer to discuss business with. Only 11% expressed their preference for someone from another sector, of different gender (13%), or different country (18%).
- **Personal traits**: The personality of the participants was recorded along several dimensions. They displayed a moderate level of risk aversion in that they chose a lottery of intermediate risk out of six hypothetically offered. On the other hand, they exhibited an intermediate trust level, with an average of 4.81 in a 0-10 scale that measured their trusting attitudes through their behavior in a trust game. Concerning expectations and desires, while on average they placed themselves in a middle position 4.82 in a 0-10 scale, they expected to escalate up to an average of 7.84 and desired to reach an average of 9.06.

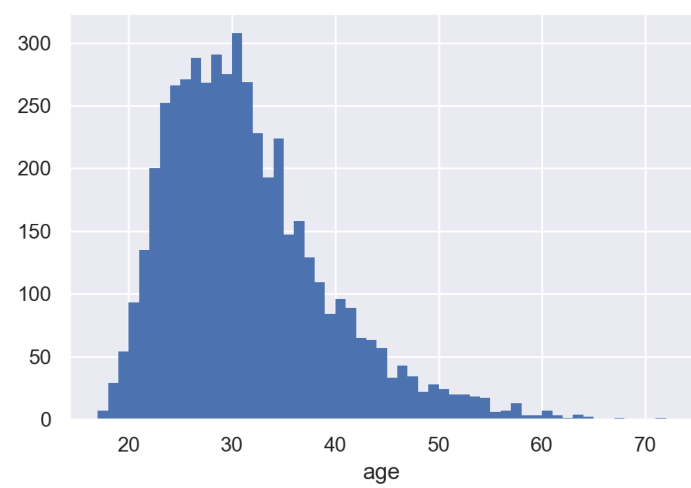

Figure 1: **Age distribution of the overall entrepreneur population.**

2.3 Treatment

As advanced, in our intervention the treatment involved being exposed to peer interaction. Such an interaction was conducted in randomly formed groups of 60 individuals (with some small deviations due to indivisibilities). We contemplated two different mechanisms and three different treatment arms.

- **Face-to-face** (f2f): In one of the treatment arms, the interaction was conducted *face-to-face*, the group meeting regularly for around three hours every two weeks in a certain location (the large premises of a partner institution). They met bilaterally or multilaterally, also subdividing themselves in subgroups of varying sizes to discuss their business ideas. This was mostly done in a self-organized manner. Around 6% of the treated individuals (all from Uganda, as explained below) were part of this arm.

- **Virtual**: Most of the peer interaction (the remaining 94% of the treated population) was performed virtually through an open-source Internet-based platform that we customized, adapting it to our needs. Individuals were again partitioned in groups (also called "instances"), each of them also consisting of 60 individuals. Every individual could create as many "channels" (or rooms) as he/she wanted, and could also choose their focus and modulate their privacy at wish. As it turns, around 30% of the channels created were bilateral. In all instances there was a general channel where anyone could participate and communication was public.
  As explained later in more detail, there were two different treatment arms among the entrepreneurs interacting virtually:
  - **Virtual-within** (vw): In this arm, groups were formed with all individuals in any given group being of the same nationality.
  - **Virtual-across** (va): In this case, all groups were formed of mixed nationalities in proportion to their respective frequency in the overall population.

- **Tokens**: To stimulate interaction, we provided an incentive mechanism that was intended to induce entrepreneurs to search for the best "matches" among their peers. It worked as follows. Every two weeks (or, in the case of the f2f arm, at the start of every bi-weekly meeting) individuals were provided with 10 tokens. They were advised to hand these tokens to those peers whom they found most helpful, both as a reward and as a means of keeping them "faithful" and involved, thus continuing to give fruitful feedback. The value of tokens derived from the fact that, at the end of the intervention, the tokens given by the 50 highest-ranked entrepreneurs (as arising from the first-stage evaluation described in Subsections 2.6 and 2.7 below) acted as lottery tickets for one of the 30 (attractive) rewards that we labeled a "Bocconi Prize."[8] In fact, the tokens were meant to operate in both directions – that is, not only as a way to attract good peers on the part of the token-giving entrepreneurs, but also as a stimulus for the peers to find "promising" entrepreneurs for whom they could prove to be most beneficial (improving their chances of being highly ranked). Indeed, as we shall see, the messages exchanged by entrepreneurs in the virtually implemented arms show that tokens played often a significant role in their discussion.

2.4 Randomization

In order to conduct a proper randomization that accommodates the restrictions imposed by the three treatment arms described above, it was convenient to divide the overall population into the following three (disjoint) samples.

[8]They involved a fully-covered two-week trip to Milan that, among other things, included the following: an intensive entrepreneurship course offered by the Bocconi Business School, a pitching event where they would submit their project to European VCs for funding, the participation in a three-day high-level conference on Africa, and some tourism (e.g. a trip to Venice).

- **The Uganda Sample** (UgS). This sample included all the individuals (568 of them) originating from Uganda who lived in the vicinity of Kampala. They were randomly partitioned into three essentially equal-sized subsamples: **control, f2f-interaction, and vw-interaction**. The first subsample was offered no peer interaction, while for the latter two the interaction provided was of the f2f or vw type, respectively.

- **The Large-Country Sample** (LCS): This sample included all individuals ($3,333$ of them) who originated from one of the five countries that we labeled as "large," i.e. Ghana, Kenya, Nigeria, South Africa, and Tanzania. These were the countries for which the number of corresponding nationals in the total population was high enough to allow for the formation of enough groups of that nationality as well as also enough groups with a suitably mixed nationality profile. This was, therefore, the only sample where vw-interaction could be suitably tested against va-interaction. The test was implemented by creating three equal-sized subsamples: **control, vw-interaction, and va-interaction**. While individuals in the first one were offered no peer interaction, the other two had access to interaction of the corresponding type.

- **The Small-Country Sample** (SCS): This sample consisted of all individuals who originated from the 44 small countries obtained by excluding the five large ones. The total number of them was 1057. Here, unlike in the previous case, all interaction among the treated entrepreneurs was of the va-type.

An illustration of the experimental design can be found in Figure 2.

For each these three samples, the randomization across the constituent subsamples was conducted by stratifying according to the following baseline characteristics: gender; country (for the LCS) or region (SCS); having a prior business; and submitting the first "milestone" of the course on time (which was prior to the randomization – see below for a description of the online course). This stratification was implemented in order to ensure control-treatment balance on those three characteristics, which was judged particularly important. However, it is shown in Tables A1-A3 in the Appendix that our randomization also produced well-balanced outcomes in all three samples: UgS, LCS, and SCS.

Our segmented three-sample design was motivated by the following two-fold objective:

(a) to identify whether **some** type of treatment (i.e. peer interaction) may be effective *vis--vis* the control;

(b) to single out (if (a) is answered affirmatively) **what** kind of treatment is most effective.

To see how our design addresses (a)-(b), first note that the **Uganda Sample** (UgS) permits a comparison of **face-to-face** and **virtual interaction**. This is specially relevant for the key issue of **scalability**. For, if the population is large, only virtual interaction will typically be feasible. From this point of view, therefore, if virtual interaction is found to be significantly less effective than the interaction conducted face-to-face, the difference can be interpreted as the "price" of scalability.

On the other hand, the **Large-Country Sample** can be used to assess the treatment effect along another dimension: **interagent diversity**, specifically focused on how peer interaction in **nationally heterogeneous** groups performs compared to that conducted in **nationally homogeneous** ones. Here, the motivation is that, while some heterogeneity seems an essential requirement for peer interaction to be genuinely fruitful for innovation, it is *a priori* unclear what features are most relevant: nationality, experience,

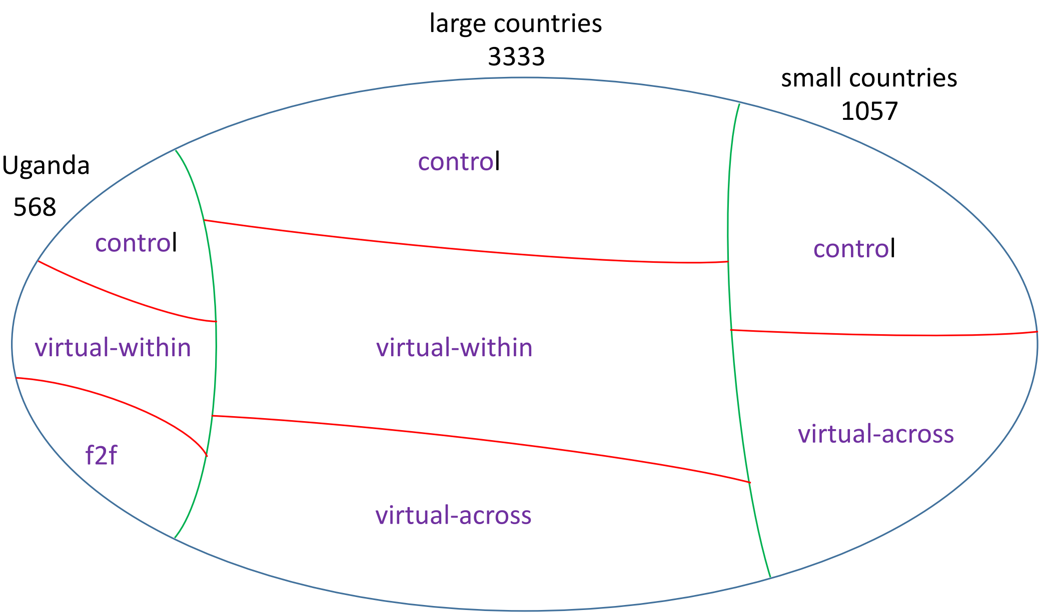

Figure 2: **Experimental design: schematic representation of the randomization conducted for each of the three samples: UgS, LCS, and SCS.**

technological or market knowledge, age, etc. Our exploration of the nationality dimension is just a first step in understanding this important problem.

Finally, the **Small-Country Sample** – including the countries where not enough nationally homogeneous groups can be formed – acts as a sort of "residual." The treated individuals in this case undergo virtual-across interaction among themselves and those individuals in the large countries assigned to the virtual-across treatment arm.

2.5 Online course

In parallel with the process of peer interaction, an online course on entrepreneurship was provided to all participants, treated and control alike. The completion of this course was required in order to be able to submit the business proposals at the end of the intervention. The course was structured in six modules and four milestones. The first of these milestones, labeled Milestone 0, was completed early on, before the randomization was performed and the peer interaction started. It included a short description of the initial business plan of the entrepreneur, which was externally evaluated and then identified with the *baseline quality* of the entrepreneur.[9] On the other hand, the last milestone, Milestone 3, involved the submission of the final business proposals, which were then evaluated (and selected for possible funding) by the two-stage procedure described in Subsection 2.6. For a schematic description of the course, please refer to Table 3.

[9]The evaluation was conducted by one of the 15 African professionals who undertook the first-stage evaluation of the full-fledged business proposals submitted at the end of the experiment.

ADANSONIA TRAINING PROGRAM - COURSE STRUCTURE						
	WHAT	TOOL	WHO	VIDEOS	SLIDES	ADDITIONAL MATERIAL/LINKS
Module 0	Introductory session	recorded video	F. Vega-Redondo	1	no	
	You don't have to be an entrepreneur	recorded video	Ojijo Pascal (GoBigHub)	1	no	
	video youtube on campuslife	https://www.youtube.com/watch?v=xWbCi5t7ZKA		1		
	syllabus	downloadable pdf				
	website tutorial	downloadable pdf				a tutorial guide to be sent out on Uploadable as well on the platform
	mobile number request	1 question quiz				
MILESTONE 0 - fill in the business proposal form with your initial idea- not evaluated						
TRAINING STARTS HERE						
Module 1 Ideating a new value proposition: involving consumers in innovation	The nature of start-ups in Africa	recorded video	Ojijo Pascal (GoBigHub)	1	no	
	Identifying white spaces for innovation	recorded video	Paola Cillo	1	yes	
	Involving Customers in Innovation	recorded video	Paola Cillo	1	yes	
	Customers as a Source of Ideation	recorded video	Paola Cillo	1	yes	
	multiple-choice quiz - evaluated					
Module 2 Testing an Idea before launch	Testing a product before launch	recorded video	Paola Cillo	1	yes	
	Testing a market before launch	recorded video	Paola Cillo	1	yes	
	multiple-choice quiz - evaluated					
	Market Research (in data short markets)	recorded video	Leticia Browne (ICG)	1	no	
MILESTONE 1 - individual assignment due / ananlysis of real life examples / general fb with solutions						
Module 3 Launching a new idea: the marketing dimension	The Business Model Canvas	recorded video	Gaia Rubera	1	yes	
	Developing your value proposition	recorded video	Gaia Rubera	1	yes	
	Customer segmentation	recorded video	Gaia Rubera	1	yes	
	multiple-choice quiz - evaluated					
	Activity based marketing and sales system	recorded video	Ojijo Pascal (GoBigHub)	1	no	
Module 4 Launching a new idea: the cost dimension	Identifying your key partners	recorded video	Giuseppe Stabilini	1	yes	
	Identifying your key activities and resources	recorded video	Marco Morelli	1	yes	
	Pricing strategy	recorded video	Gaia Rubera	1	yes	
	Pricing strategy	recorded video	Gaia Rubera	1	yes	
	multiple-choice quiz - evaluated					
	Systems are more important than products	recorded video	Ojijo Pascal (GoBigHub)	1	no	
MILESTONE 2 - individual assignment due / BMC exercise //general FB						
Module 5 Finalizing your business plan	Financial forecast in your BP	recorded video	Leonardo Etro	1	yes	
	How to Value a Start Up	recorded video	Leonardo Etro	1	yes	
	multiple-choice quiz - evaluated					
Module 6 Getting ready to present your startup	Build your startup team	recorded video	Massimo Magni	1	yes	
	What are the key elements of an investment pitch	recorded video	Leticia Browne (ICG)	1	no	
	Guidelines for submission for funding/ mini tutorial "How to pitch your idea"	recorded video	Massimo Pulejo	1		
	multiple-choice quiz - evaluated					
MILESTONE 3 - fill in the official business proposal form						

Figure 3: **Structure of the online entrepreneurship course.**

The online course and the peer interaction were conceived as complementary and implemented in parallel with a less-than-perfect overlap (see below for the precise timing). The lectures were given by Bocconi professors and African VCs. They were video-recorded and supported by ongoing slides. At the end of each module, the participants had to take an online multiple-choice quiz, which was corrected and evaluated automatically. The content of the course was a relatively standard entrepreneurship course, but with an approach and content (e.g. the case studies) adapted to the African environment. The main objective was simple: by the end of the course, each participant should be able to draft an economically coherent business proposal and have some basic notions on how to pitch it.

2.6 Timeline

The field experiment was conducted during the spring-summer of 2017. The following steps took place in sequence:

1. The **online course** lasted for ten weeks, from *May 22nd to July 31st*. Module 0 (where the participants summarized their initial business idea) had to be completed no later than May 29th.

2. The **peer interaction** (in all treatment arms) ran from *June 6th to August 14th.* Thus it had the same duration as the course, ten weeks, but started two weeks later. This meant, in particular, that the treatment had no effect on Milestone 0.

3. The **submission of business proposals** could be performed at any point within the period running from *August 8th to August 15th.* All participants were allowed to resubmit a revised version up to the deadline.

4. **The first-stage evaluation** of business proposals took place from *September 15 to October 15.* It was conducted by a panel of 15 African professionals with ample experience in entrepreneurship programs (mentoring, training, and investment) throughout the continent. They graded all the submitted proposals in a 1-5 scale. These grades defined one of the primary outcome variables of the experiment (see Subsection 2.7).

5. The **second-stage evaluation** took place from mid November 2017 to February 2018 and a group of 36 partner investors affiliated to the Adansonia program (mostly VCs and angel investors) were involved. The evaluation, applied to a subset of selected "good" proposals, was carried out as follows:

   (a) First, the proposals that were assigned one of the top two grades (4 or 5) in the first stage by the panel of African professionals were included into the selected set. This amounted to 444 proposals.

   (b) Second, in order to increase the size and diversity of the pool of proposals considered, an evaluation analogous to that of the first stage was conducted by a panel of 15 advanced Ph. D. students from the Bocconi Ph.D. program of Business Administration. Any ***new*** proposal assigned a top-two grade by these evaluators was also included in the selected set. This added 164 proposals.

   (c) Third, the 608 proposals obtained through (a) and (b) were distributed to investors, according to their expertise and declared interest. Eventually, only 24 out of the 34 original investors completed their task on time to be considered,[10] which reduced to 454 the set of proposals actually evaluated. All these proposals were not only graded (again in a 1-5 scale) but also considered for possible funding. A total of 93 proposals actually moved into a concrete follow-up by the investor. Up to 30 of them had been funded or were still followed up as of November 2018.

2.7 Outcomes

Our experiment delivers a wide variety of outcomes, which in turn calls for a rich multisided approach in analyzing the evidence. The main outcome dimensions we will consider are the following.

1. **Submission**
   A first and important outcome variable is the number of participants (treated and control) who actually submit a business proposal at the end of the intervention (see item 3 above). As we shall see, there is a significant fraction of individuals who do not submit a proposal, and the alternative treatment arms have a differential effect on the extent of submission. To understand this pattern will be our first step in the analysis.

[10]This happened for a variety of reasons. Some investors quit business since the time when they were recruited, while others were too busy or simply lost interest in partnering in the intervention.

2. **Business quality**
   The business projects submitted by the participants are evaluated in two steps: first (see item 4 above), by a panel of African professionals; second (see item 5), a subset of "good" proposals is evaluated by a group of partner investors. In both cases, the evaluators assigned a grade from 1 to 5, in increasing order of quality. The first-step grade will be the outcome variable measuring business value.

3. **Network**
   As explained in Subsection 2.3, 96% of the treated population interacted virtually through the Internet-based platform designed for this purpose. For these treated participants, we have a comprehensive longitudinal (panel) record of their networking efforts throughout the intervention. Operationally, in constructing the social network from of these data, we shall consider different procedures, i.e. different ways of declaring that a link exists between any two entrepreneurs, as well as different forms of quantifying the density of every such link. For each of these alternative formulations, we shall want to shed light on how the pattern of (weighted) connections unfolds, depending on different circumstances (e.g. particular treatment) and individual characteristics (e.g. gender, age, or experience). And then, finally, our objective will shift in turn towards understanding how those connections that form over time bring about peer effects, i.e. behavioral or performance relationships across interacting entrepreneurs.

4. **Messages**
   Concerning the interaction taking place within the treated population, not only do we have information on inter-agent links and their density (the volume of posts they write) – we also have full access to the posts themselves. To analyze them, we shall focus on the constituent sentences as the simplest units of intrinsic meaning and will study them semantically. This we will do by heavily relying on the techniques developed by the field of Natural Language Processing (NLP), which will allow us to identify a consistent pattern of meaning across the whole communication undertaken (a total of 140,000 sentences written). More specifically, we shall primarily focus our efforts on three different categories: business relevance, sentiment (positive or negative), and audience scope (bilateral or general).

The four items listed above (submission, business quality, networking, and messages) are the main *outcome dimensions* to be studied here. Each of them, of course, can be studied separately, and this is what we shall partly do. But the full richness of the setup will become apparent only when we proceed to a multidimensional analysis of the evidence. Then, the effect on business quality arising from peer interaction will not just be studied as an abstract embodiment of the "treatment effect." Rather, we shall try to move beyond such a black-box analysis and attempt to shed light on how such peer effects are channeled – in a supporting or hindering manner – through the evolving pattern of connections and communication. All this, in sum, should provide not only a multidimensional perspective on the problem but also information on the *concrete mechanism* through which the treatment, if successful, ends up working its way through.

3 The core treatment effects

As explained, the treatment – in its different variants – may have an effect on a number of different experimental outcomes. In this section, most of the focus is on what may be regarded as the two most basic effects of the treatment, namely, the impact it has on the rate at which the subjects submit their business proposals (Subsection 3.1) and its influence on the quality of the proposals submitted (Subsections 3.2.1 and 3.2.2). In the last part of this section we will also study the following issues: (a) how such a treatment effect depends on the profile compositions of peer groups (Subsection 3.3); (b) what is its impact on the second-stage evaluation (Subsection 3.4); (c) whether it passes various robustness checks(Subsection 3.5). Then, in Sections 4 and 5 we will respectively turn to studying how the treatment impinges on inter-agent networking and then on their communication.

3.1 Submission

Naturally, prior to studying how the treatment affects the quality of the business proposals submitted, a preliminary important concern is how it influences the extent to which entrepreneurs submit their business proposals in the first place. The results in this respect are summarized for the different treatment arms in Table 2, where we code no-submission/submission as a binary variable $\sigma \in \{0, 1\}$, and partition the analysis into our three leading samples: UgS, LCS, SCS. This table has two columns. The first one considers the full samples in each of the three aforementioned cases, while the second column restricts attention, in each case, to the subsamples of those who submitted Milestone 0 (M0) on time in the business course – cf. Subsection 2.5. As already explained, since M0 had to be submitted prior to randomization, the subsamples of "M0 completers" can be regarded as exogenously constructed. We shall interpret those entrepreneurs (55% of the overall population) as being those who were especially motivated (or otherwise more available to focus on the intervention). An indication that this interpretation is a reasonable one follows from the observation that the submission rate under this condition is *uniformly* (and substantially) higher than in the unconditioned case (in particular, it is essentially around 20 points higher among the control subpopulation for each of the three original samples). The fact that increased take-up/submission enhances statistical power justifies studying how the treatment influences such a subsample of "motivated individuals."

The first observation to gather from Table 2 is that, overall, the submission rate in the treated subpopulations is higher than the control. Specifically, for the general samples we obtain a combined average of 41.7% across the different treatment arms and 38.1% in the control, while the respective rates for the M0 completers are 58.3% and 63%. This aggregate information, however, abstracts from what is a marked contrast in the magnitude and sign of the treatment effect on submission that is observed in Table 2 among the three samples.

Let us start by the first column of the table that separately includes all the entrepreneurs belonging to each of the three samples: the Uganda Sample (UgS), the Large-Country Sample (LCS), and the Small-Country Sample (SCS). For the UgS we find that whereas the treatment effect on submission is positive and very significant for the face-to-face (f2f-)treatment, the positive effect arising under virtual-within (vw-)interaction is not statistically significant. In view of the fact that the vw-treatment is positive and significant for the LCS, it is reasonable to understand the non-significance of the vw-treatment for the UgS as a consequence of the small sample size. This small-sample handicap, however, does not crucially impinge on the f2f-treatment, whose effect is strong enough to arise as significant despite the small sample.

Next, to understand the comparative effect of v-within and v-across scenarios we must focus on the LCS, which is the only context where the two treatment subarms are implemented. There we we find that, in contrast with the aforementioned positive and significant treatment of vw-interaction, the virtual-across (va-)treatment yields no significant result. Indeed, a detrimental effect of va-interaction on submission manifests

Table 2: **Treatment effect on the submission of business proposals**

	(1) Submitted, full samples	(2) Submitted, M0 completers
Panel A. Uganda Sample (UgS)		
Face-to-face treatment	0.126***	0.151**
	(0.038)	(0.065)
Virtual-within treatment	0.022	0.041
	(0.034)	(0.064)
Control Mean	0.34	0.55
p-value face-to-face = virtual within	0.00	0.05
Number of Entrepreneurs	568	291
Panel B. Large-Country Sample (LCS)		
Virtual-within treatment	0.036**	0.048**
	(0.015)	(0.023)
Virtual-across treatment	0.014	0.018
	(0.018)	(0.026)
Control Mean	0.38	0.58
p-value virtual-within = virtual-across	0.18	0.15
Number of Entrepreneurs	3,333	1,848
Panel C. Small-Country Sample (SCS)		
Virtual-across treatment	-0.057**	-0.036
	(0.024)	(0.037)
Control Mean	0.42	0.61
Number of Entrepreneurs	1,057	596

Notes: Estimated coefficients for OLS regressions of the submission outcome on indicators for treatment (face-to-face, virtual-within, and virtual-across interaction) in the three different samples considered: UgS, LCS, and SCS. Column (1) applies to the whole population in each of the samples, while Column (2) restricts to those who submitted M0 on time in the business course. We include strata fixed effects and cluster the errors at the group level for treated individuals (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively.

even more starkly on the SCS, where the effect on submission is both negative and significant (relative to the control). All this suggests that there are consequential differences between the two virtual treatments (within and across) which need to be not only clarified but also revisited for other outcomes (business quality, networking, and messaging). Indeed, this will be one of the primary objectives that will be pursued when studying, in particular, networking decisions and communication in Sections 4 and 5, respectively.

Finally, as a counterpart of the former conclusions, in the second column of Table 2 we report the results derived from a parallel analysis undertaken for the subset of "motivated" individuals who, in each of the three samples, submitted M0 on time. The treatment effect on submission displays, across the three setups, a similar behavior as before except for one difference: the significant negative effect formerly observed for the SCS now becomes statistically insignificant. Thus, to sum up, we may conclude that, while a restriction to M0 completers increases the rate of submission and the treatment effects, it has no substantial impact on the *qualitative pattern* induced by the different treatment arms on submission. Heuristically, that pattern can be informally described as follows. There is a progressively decreasing effect running from the context where interaction is most "***direct***" and the peer pool is more ***homogeneous*** and ***motivated*** (f2f in the UgS for M0 completers) to that where the interaction is virtual and the pool is more ***heterogeneous*** and on average ***less motivated*** (SCS, including both those who completed M0 on time and those who did not).

3.2 Business quality

Now we extend our previous analysis to the effect that the treatment has on business quality. First, in Subsection 3.2.1, we do this by exploring what we shall call the *extensive margin*. This amounts to defining an outcome that combines information on both submission *and* the quality of the submitted business proposals, the latter being the result of the first-stage evaluation described in Subsection 2.6. This combined outcome is measured by an index $\hat{\varpi} \in \{0, 1, ..., 5\}$ that integrates the information of no-submission ($\sigma = 0$) with the grade $\varpi \in \{1, 2, ..., 5\}$ reflecting the quality of a business proposal when submitted. In contrast, the analysis of the intensive margin, which will be studied in Subsection 3.2.2, focuses only on the submitted proposals and has as outcome the evaluation ϖ assigned to them by the evaluation procedure.

3.2.1 The extensive margin

As it turns out, most of the features highlighted before in Subsection 3.1 when analyzing submission reappear here in a similar vein. The results are summarized in Table 3.

Thus, indeed, as obtained for submission, Table 3 shows a positive and highly significant treatment effect for f2f-interaction in the UgS and for vw-interaction in the LCS, while all other effects are statistically insignificant at even the 10% level, both for the unconditioned sample and for that conditioned on the on-time completion of M0. (The latter contrasts with the mildly significant negative effect encountered for submission in the SCS, but only for the unconditional subsample.[11])

The very similar conclusions observed in Tables 2 and 3 raise the question of whether, in essence, the effects reported for the extensive margin are mostly a reflection of the impact of the treatment on submission. Or, somewhat differently formulated, the issue we raise is whether it might be the case that the treatment has no significant effect on the quality of *submitted* proposals, and the effects reported above are simply capturing the treatment impact on submission. To estimate this quality-enhancing effect – which reflects what we shall call the *intensive margin* – is the object of the next subsection.

[11] Another difference is that the positive effect for the vw-treatment in LCS is statistically more significant in the present case.

Table 3: **Treatment effect on business quality: extensive-margin analysis**

	(1)	(2)
	Quality, full samples	Quality, M0 completers
Panel A. Uganda Sample (UgS)		
Face-to-face treatment	0.411***	0.420**
	(0.106)	(0.204)
Virtual-within treatment	0.165	0.333
	(0.117)	(0.226)
Control Mean	0.88	1.47
p-value face-to-face = virtual within	0.01	0.69
Number of Entrepreneurs	568	291
Panel B. Large-Country Sample (LCS)		
Virtual-within treatment	0.128***	0.232***
	(0.047)	(0.077)
Virtual-across treatment	0.047	0.065
	(0.056)	(0.088)
Control Mean	1.03	1.58
p-value virtual-within = virtual-across	0.08	0.03
Number of Entrepreneurs	3,333	1,848
Panel C. Small-Country Sample (SCS)		
Virtual-across treatment	-0.135*	-0.045
	(0.073)	(0.116)
Control Mean	1.10	1.62
Number of Entrepreneurs	1,056	595

Notes: Estimated coefficients for OLS regressions of the outcome on indicators for treatment (face-to-face, virtual-within, and virtual-across interaction) in the three different samples considered: UgS, LCS, and SCS. The dependent variable takes discrete values from 1 to 5 representing the quality score (in ascending order) given by 15 African evaluators. We impute a quality value of 0 to proposals that were not submitted. Column (1) applies to the whole population in each of the samples, while Column (2) restricts to those who submitted M0 on time in the business course. We include strata and evaluator fixed effects. Standard errors are clustered at the group level for treated individuals (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively.

3.2.2 The intensive margin

The estimated effects of the different treatment arms on the business quality of those proposals actually submitted (and hence the only ones evaluated) are reported in Table 4. We observe that for the subsample of "motivated" entrepreneurs (those who submitted M0 on time), vw-interaction produced a statistical significant positive effect, both in the UgS and the LCS. Instead, for the other treatment arms (f2f- and va-interaction), no significant effect is identified. A similar pattern arises for what we label the *quality effect*, defined as the difference in average quality increase between the submitters that belong to the different treatment arms and the control subpopulation.

Table 4: **Treatment effect on business quality: intensive-margin analysis**

	(1)	(2)
	Intensive margin, full samples	Inten. marg., M0 completers
Panel A. Uganda Sample (UgS)		
Face-to-face treatment	0.180	0.059
	(0.267)	(0.230)
Virtual-within treatment	0.335^{*}	0.427^{**}
	(0.179)	(0.204)
Quality control mean if submitted	2.61	2.66
Quality effect face-to-face	0.18	0.03
Quality effect virtual within	0.30	0.38
Number of entrepreneurs who submitted	221	179
Number of entrepreneurs	568	291
Panel B. Large-Country Sample (LCS)		
Virtual-within treatment	0.070	0.138^{**}
	(0.070)	(0.068)
Virtual-across treatment	-0.021	-0.007
	(0.079)	(0.083)
Quality control mean if submitted	2.70	2.71
Quality effect virtual within	0.08	0.16
Quality effect virtual across	0.02	0.03
Number of entrepreneurs who submitted	1,322	1,118
Number of entrepreneurs	3,333	1,848
Panel C. Small-Country Sample (SCS)		
Virtual-across treatment	0.033	0.074
	(0.108)	(0.115)
Quality control Mean if submitted	2.65	2.65
Quality effect virtual across	0.05	0.09
Number of entrepreneurs who submitted	409	351
Number of entrepreneurs	1,056	595

Notes: The table presents, for each panel, the intensive margin results from OLS regressions of the quality obtained in the evaluation on indicators for treatment (face-to-face, virtual-within, and virtual-across interaction) restricting to the sample that submitted a proposal. The dependent variable takes discrete values from 1 to 5 representing the quality score (in ascending order) given by 15 African evaluators. Column (1) applies to the whole population in each of the samples, while Column (2) restricts to those who submitted M0 on time in the business course. We include strata and evaluator fixed effects. Standard errors are clustered at the group level for treated individuals (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively. Following Attanasio et al. (2011), we also report for each sample the corresponding *quality effect*, which is defined as the (absolute) increase in average quality induced by the treatment, which can be readily computed from the analysis on submission and the extension-margin displayed in Tables 2 and 3.

Clearly, the validity of the previous conclusions is subject to the risk of a selection (or composition) bias since the set of submitters is endogenous. And indeed, in our context, this possibility cannot be *a priori* discarded since a sizable subset of participants did not submit (recall Subsection 2.7) and, furthermore, treatment and control displayed quite different extents of attrition. In view of this, we need to contemplate additional considerations that may reasonably rule out (or at least bound) the potentially distorting impact of the aforementioned election bias.

To this end, we pursue the approach proposed in Attanasio *et al* (2011), also pursued by Bandiera *et al* (2017). The key assumption underlying it is that the population does not have any so-called "defiers," i.e. individuals who do not submit a proposal when treated, while they do so when not treated. This appears to be a plausible assumption in that it reflects the idea that, when an entrepreneur is given the ***opportunity*** to interact with peers (which, of course, can be simply ignored), this should not discourage her from submitting a proposal when she should have done it otherwise (i.e. as part of the control).

To formalize matters denote by $S(\theta) \in \{0,1\}$ the decision of whether to submit ($S(\theta) = 1$) or not ($S(\theta) = 0$) where $\theta \in \{0,1\}$ is an indicator for treatment ($\theta = 1$) or not ($\theta = 0$). With this notation in place, we can identify each of the four possible types of individuals with a corresponding pair $(S(0), S(1)) \in \{0,1\}$. The aforementioned *defiers* are characterized by the pair $(1,0)$, while those described by the pair $(0,1)$ can be called *compliers*. In contrast, the individuals of type $(1,1)$ will be called *always submitters*, and those of type $(0,0)$ *never submitters*. For each possible $\tau \in \{0,1\}^2$, we shall denote by $\mathbb{P}(\tau)$ the fraction/probability of individuals of type τ in the population. Then, Attanasio *et al.* (2011) show that, if the following assumption holds:

No Defiers (ND): $\mathbb{P}(1,0) = 0$,

i.e., the population has no defiers, then the expected intensive-treatment effect on business quality $\varpi \in \{1,2,...,5\}$ among submitted proposals (i.e. conditional on $\sigma = 1$) can be decomposed as follows:

$$\begin{aligned}\mathbb{E}[\varpi \,|\, \sigma = 1,\, \theta = 1] - \mathbb{E}[\varpi \,|\, \sigma = 1,\, \theta = 0] &= \{\mathbb{E}[\varpi \,|\, \tau = (0,1),\, \theta = 1] - \mathbb{E}[\varpi \,|\, \tau = (0,1),\, \theta = 0]\}\ \frac{\mathbb{P}(0,1)}{\mathbb{P}(0,1) + \mathbb{P}(1,1)} \\ &+ \{\mathbb{E}[\varpi \,|\, \tau = (1,1),\, \theta = 1] - \mathbb{E}[\varpi \,|\, \tau = (1,1),\, \theta = 0]\}\ \frac{\mathbb{P}(1,1)}{\mathbb{P}(0,1) + \mathbb{P}(1,1)} \\ &+ \{\mathbb{E}[\varpi \,|\, \tau = (0,1),\, \theta = 0] - \mathbb{E}[\varpi \,|\, \tau = (1,1),\, \theta = 0]\}\ \frac{\mathbb{P}(0,1)}{\mathbb{P}(0,1) + \mathbb{P}(1,1)}.\end{aligned} \tag{1}$$

The LHS of (1) is the effect we can measure from the data whereas the effect that we are truly interested in here is given by the sum of the first two terms of the RHS of that expression. The difference between these two magnitudes is given by the third term of (1), which involves precisely the composition bias – $\mathbb{E}[\varpi \,|\, \tau = (0,1),\, \theta = 0] - \mathbb{E}[\varpi \,|\, \tau = (1,1),\, \theta = 0]$ – that concerns us here. Specifically, an incorrect identification of the sign of the computed effect would be associated to a positive such term that is large enough to "mask" a non-positive effect resulting from the two first terms.

Attanasio *et al* (2011) explores two different ways of addressing the issue. One of them involves bounding the difference between $\mathbb{E}[\varpi \,|\, \tau = (0,1),\, \theta = 0]$ and $\mathbb{E}[\varpi \,|\, \tau = (1,1),\, \theta = 0]$ by relying on the observed distribution of business quality observed for the non-treated submitters. In their context, however, this approach proves unsuccessful because the wedge between the estimated upper and lower bounds is too wide, and hence a negative effect in the intensive margin cannot be discarded.[12] As it turns out, an analogous procedure applied to our context delivers as well inconclusive results.[13]

The alternative approach proposed by Attanasio *et al* (2011) to deal with the composition bias relies on the following assumption:

Better Always Submitters (BAS): $\mathbb{E}[\varpi \mid \tau = (1,1),\, \theta = 0] - \mathbb{E}[\varpi \mid \tau = (0,1),\, \theta = 0] \geq 0$

which states that that the expected business quality of those who submit in every case (both when they belong to treatment and control) is no lower than that of the entrepreneurs who switch into submitting only when they are treated. Arguably, this is quite a plausible assumption. It reflects the intuitive idea that those who, even in the absence of peer interaction, would always submit a proposal are entrepreneurs who are relatively more self-confident or/and were more motivated by the possibility of funding. (Another complementary reason might be that they were those who benefited the most from the online course.) Then one may be sure – under maintained assumption (ND) that underlies (1) – that the estimated effect given by the RHS (1) *can only underestimate* the true impact of the treatment on intensive margin, i.e. the influence of the treatment on the quality of a treated entrepreneur, whether she finally submits her proposal or not. Therefore, if the estimated treatment effect on quality – *conditional*, of course, on submission – is positive and significant, this is in effect an unambiguous indication of a positive *unconditional* effect as well.

To complete our analysis of the intensive margin in the treatment effect, we conclude this subsection with an estimation of the marginal effects induced on business quality. Conducting such an exercise through an ordered Probit, the results are displayed in Table 5. As expected, we find that the treatment has the effect of shifting weight from the lowest scores (1 and 2) to the highest ones (3 to 5) in the virtual-within arms (with a non-significant effect for score 3 in the UgS) while it yields no significant weight reassignment in all other cases.

3.3 Peer-composition effects

It is interesting to study how the effect of the treatment depends on the characteristic profile of the members of the group within which the individual interacts. Even though, obviously, such an effect must be mediated by interaction, it is conceptually different from the network-based interaction that will be studied in Section 4 in at least two key respects. First, it is group-based (and hence *exogenously* determined by randomization) rather than dependent on the (endogenous) relations that an individual chooses to establish within her group. Second, it is related to the baseline (exogenous) characteristics of the peers in the assigned group rather that on the their (endogenous) behavior.

In Table 6 we focus on two prominent baseline characteristics. One of them is the average baseline quality of peers, as obtained from the external evaluation of the initial business outline (see Subsection 2.5). The second characteristic is the average experience of peers, as measured by the fraction of those who were already operating a business at baseline. For the sake of focus and contrast, the results are reported only for our largest sample, LCS, which is also the only context where we find positive and significant treatment effects for all the three outcomes: submission, the intensive quality margin, and the extensive one. Panel

Table 5: **Marginal quality effects (ordered Probit), intensive-margin, M0 completers**

	(1) Score = 1	(2) Score = 2	(3) Score = 3	(4) Score = 4	(5) Score = 5
Panel A. Uganda Sample (UgS)					
Face-to-face treatment	-0.012	-0.010	0.007	0.012	0.004
	(0.050)	(0.042)	(0.028)	(0.048)	(0.016)
Virtual-within treatment	-0.078**	-0.094**	0.026	0.101**	0.045*
	(0.036)	(0.042)	(0.022)	(0.046)	(0.024)
Number of entrepreneurs = 179					
Panel B. Large-Country Sample (LCS)					
Virtual-within treatment	-0.032**	-0.019**	0.007*	0.028**	0.015**
	(0.016)	(0.009)	(0.004)	(0.014)	(0.007)
Virtual-across treatment	0.002	0.001	-0.001	-0.002	-0.001
	(0.020)	(0.010)	(0.006)	(0.017)	(0.008)
Number of entrepreneurs = 1,118					
Panel C. Small-Country Sample (SCS)					
Virtual-across treatment	-0.016	-0.009	0.005	0.016	0.005
	(0.027)	(0.016)	(0.008)	(0.026)	(0.009)
Number of entrepreneurs = 351					

Notes: The table presents the *marginal* effects on the predicted probabilities for the different business-quality scores estimated by an ordered Probit regressions on indicators for treatment (face-to-face, virtual-within, and virtual-across interaction) in the three different samples considered: UgS, LCS, and SCS. We restrict the sample to entrepreneurs who submitted a business proposal. We include strata and evaluator fixed effects. Standard errors are clustered at the group level for treated individuals (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively.

A presents the results for the full sample, while Panel B restricts to the subsample that submitted M0 on time, before the randomization.

Interestingly, we find that both the average baseline quality and the share of experienced peers exert a negative effect on the business quality of treated entrepreneurs, which is significant in every case at least in the intensive margin. It is is worth highlighting that this occurs together with an (opposite) positive treatment effect (also significant in the intensive margin) and for both panels considered (the unrestricted subsample and that restricted to M0 completers), and even if one controls for own baseline quality. Such negative composition effects are reminiscent of those identified by Lerner and Malmendier (2013) and Hacamo and Kleiner (2018), discussed in the Introduction. Recall that these authors suggest that the reason why the influence of peers with prior business experience tends to be detrimental is that those peers often tend to provide discouraging feedback. Here, not only do we find a negative effect associated to peer experience but also encounter a similar one for peer baseline quality. This highlights that peer influence is, in general, a double-edged sword, whose positive or negative impact may depend on both context and the particular type of networking and communication established by the individuals. Networking and communication will be the object of Sections Section 4 and 5, respectively. As we shall see, however, the preliminary analysis conducted there still falls short of shedding much light on the problem, which will require in the future a more thorough study specifically focused on this dimension of the empirical evidence.

[12]See Bandiera *et al* (2017) for a recent application of the same approach to a different context.

[13]In particular, for the two samples where our estimation yields a significant treatment coefficient (cf. Table 4), the bounds implied do not remain in the positive region, not even if one restricts to the M0 subsample. For UgS the bounds computed to the aforementioned subsample define the interval $(-1.76, 1.82)$, while for the LCS the corresponding interval is $(-0.30, 0.35)$.

Table 6: **Effects of peer composition**

	(1)	(2)	(3)	(4)	(5)	(6)
	Submission	Quality ext.	Quality int.	Submission	Quality ext.	Quality int.
Panel A. Full Sample in large countries						
Virtual-within treatment	0.024	0.089^{**}	0.116	0.024	0.092^{**}	0.124^{*}
	(0.015)	(0.043)	(0.071)	(0.015)	(0.044)	(0.070)
Own baseline quality	0.116^{***}	0.450^{***}	0.246^{***}	0.116^{***}	0.452^{***}	0.247^{***}
	(0.015)	(0.050)	(0.046)	(0.015)	(0.050)	(0.045)
Average baseline quality of peers	0.017	-0.124	-0.424^{**}			
	(0.046)	(0.144)	(0.168)			
Share of peers with business				0.078	-0.321	-1.539^{***}
				(0.107)	(0.344)	(0.502)
Number of entrepreneurs	2,222	2,222	899	2,222	2,222	899
Panel B. Sample who completed M0 on time in large countries						
Virtual-within treatment	0.034	0.178^{**}	0.147^{**}	0.035^{*}	0.186^{***}	0.167^{***}
	(0.020)	(0.073)	(0.066)	(0.020)	(0.068)	(0.059)
Own baseline quality	0.116^{***}	0.445^{***}	0.249^{***}	0.116^{***}	0.447^{***}	0.250^{***}
	(0.014)	(0.050)	(0.046)	(0.015)	(0.050)	(0.045)
Average baseline quality of peers	-0.028	-0.244	-0.350^{*}			
	(0.051)	(0.157)	(0.191)			
Share of peers with business				-0.082	-1.230^{**}	-1.600^{***}
				(0.135)	(0.469)	(0.363)
Number of entrepreneurs	1,231	1,231	758	1,231	1,231	758

Notes: The table uses data for entrepreneurs in the virtual interaction arms in the Large-Country Sample (LCS), either virtual-within or virtual-across interaction. In all regressions we include strata fixed effects and control for baseline level of quality obtained from evaluations of a summary of proposals submitted before treatment assignment; we replace missing values in baseline quality with zeros and include a dummy for missing observations. Standard errors are clustered at the group level (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively. The outcome in columns (1) and (4) is submission of the proposal, for columns (2) and (5) quality evaluation replacing missing values with zeros, for columns (3) and (6) quality evaluation dropping missing values. Average Peer Quality is the average of the baseline quality score of other group members. Share Peer with Business is the share of other group members that had a business at baseline. Panel A uses the full sample and uses 0 as the quality score of those who did not submit the the Milestone-0 (baseline) proposal before the randomization. Panel B restrict the sample to those who submitted Milestone 0 before the randomization and for which we have a baseline quality score.

3.4 The second-stage evaluation

As explained in Subsection 2.6, the set of submitted proposals that ranked highest in a first stage of the evaluation procedure were then evaluated in a second stage (and considered for possible investment) by a panel of investors. In principle, there is no reason to expect that the average quality of the proposals selected into the second stage should be any different depending on whether they originated from the control or the treated subpopulation. It may be conjectured, therefore, that *no* significant "treatment effect" should be found in the evaluation conducted in the second stage. Some results bearing on this question are reported in Table 7.

Table 7: **Second-stage outcomes**

	(1)	(2)	(3)	(4)
	Reach stage 2	Evaluated stage 2	Investor interest	Quality stage 2
Panel A. Uganda Sample (UgS)				
Face-to-face treatment	0.079*	-0.179	-0.069	-0.698*
	(0.041)	(0.168)	(0.191)	(0.334)
Virtual-within treatment	0.046	-0.120	0.128	-0.423
	(0.032)	(0.181)	(0.285)	(0.502)
Control mean	0.06	0.75	0.22	0.53
p-value face-to-face = virtual within	0.48	0.53	0.32	0.45
Number of entrepreneurs	568	60	37	37
Panel B. Large-Country Sample (LCS)				
Virtual-within treatment	0.026**	0.005	-0.090	-0.081
	(0.012)	(0.048)	(0.068)	(0.136)
Virtual-across treatment	0.003	0.049	-0.074	-0.051
	(0.012)	(0.048)	(0.062)	(0.141)
Control mean	0.12	0.77	0.28	-0.01
p-value virtual-within = virtual-across	0.02	0.31	0.79	0.83
Number of entrepreneurs	3,333	418	332	332
Panel C. Small-Country Sample (SCS)				
Virtual-across treatment	-0.020	0.054	-0.089	-0.096
	(0.018)	(0.090)	(0.099)	(0.253)
Control mean	0.13	0.64	0.20	0.04
Number of entrepreneurs	1,057	130	85	85

Notes: The table presents the results for OLS regressions of the outcome on indicators for treatment (face-to-face, virtual-within, and virtual-across interaction) in the three different samples considered: UgS, LCS, and SCS. In all regressions we include strata fixed effects. Standard errors are clustered at the group level for treated individuals (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively. The outcomes of the first three columns are all binary indicators. In column (1), the indicator specifies whether each entrepreneur is selected to reach the second evaluation stage. (As explained in Subsection 2.6, this selection was made on the basis of the evaluation conducted in the first stage.) In column (2), it indicates whether each proposal was actually evaluated by venture capitalists in the second stage – i.e. was not affected by exogenous disruptions. And in Column (3), the indicator records whether some investor reported interest in funding the project of the entrepreneur. Here, the sample is restricted to the 75% of the entrepreneurs selected to the second stage who were evaluated by venture capitalists. Finally, in column (4), also restricting the sample as in the previous case, the outcome considered is the quality score given to projects by the venture capitalists, these scores being standardized at the venture-capitalist level.

The first column of Table 7 shows, as expected, a positive treatment effect of reaching the second stage for f2f- and vw-interaction. Their corresponding (positive) coefficient is statistically significant only for f2f-interaction (in the UgS) and vw-interaction in the LCS. This is in line with the results obtained for the extensive-margin analysis displayed in Table 3, somewhat attenuated by the fact that, for the selection into the second stage, not only the evaluation by the African panel but also that by the European one was used as a subsidiary input.

The other three columns of 7 involve only the proposals that actually reached the second stage. First, Column 2 shows that the treatment had no significant effect on whether these proposals were actually evaluated by investors (i.e on whether or not the assigned investors did complete their evaluation task).[14] This indicates that, as one would expect, such investor-induced distortion did not have any significant differential effect on the issue at hand. Finally, the last two columns concern the different types of feedback provided by the investors. While Column 3 refers to the investors' interest in the project – as manifested in follow-up contacts with the entrepreneur for possible investment – Column 4 relates to the evaluation issued by the investors (in a 1-5 scale) on the quality of the proposals. We find that, for both of these outcomes, the treatment coefficients are statistically insignificant, except for a negative one (weakly significant) for the investor evaluation in the f2f-treatment in the UgS.

The latter results are in line with our prior discussion, but their interpretation can be only suggestive. For, in effect, the econometric exercise involves subsamples (from the control and different treatments) that are not exogenous and hence are subject to the possibility of some selection bias. This bias would be in effect an "evaluation bias," inducing some systematic difference on how the evaluation was conducted across control and treatment(s) for reasons that are *spuriously* associated to whether the proposals originated in one or the other. Instead, a seemingly more plausible explanation of the absence of treatment effect in the second stage is that, in fact, there was a reasonably good alignment in the evaluation criteria used in the two stages.

3.5 Robustness

We have conducted different robustness checks on the treatment analysis that are informally described in the present section, while the detailed results are reported in the Appendix.

On the one hand, we have conducted replicas that re-estimate the treatment effect, both in the extensive and the intensive margin, after adding controls for all the variables listed in the balance tables (see Tables A1-A3). The result of the exercise is displayed in the following tables in the Appendix: Table A4 (extensive-margin analysis), Table A5 (intensive-margin analysis), Table A6 (marginal quality effects), and Table A7 (second-stage outcomes). All the previous conclusions remain essentially unchanged.

On the other hand, we have revisited the effects of peer composition included in Table 6 and checked their robustness from two complementary viewpoints. First, in Table A8, we carry out a placebo test in which, relying on the same procedure used in the randomization of our experiment, groups are randomly formed among the entrepreneurs from the LCS that were assigned to the control. Then, we estimate whether significant effects of peer composition exist for such a randomization and find that, in contrast with the results found for the treated population, no composition effects are present. The second exercise is akin to that conduced by Cai and Szeidl (2017). We compute a "surprise effect" associated to the deviation of the average peer profile of every group from the one that one would expect in a randomly formed group. The difference between the two magnitudes reflects a surprise that, in the absence of bias, should be of the same type (sign and significance) as the formerly estimated effect. Indeed, Table A9 shows that the sign and pattern of significance obtained from such a construction is largely aligned with that displayed in Table 6.

[14] As explained in Subsection 2.6, some VCs did not finally fulfill their original commitment to evaluate the business proposals assigned to them for a variety of different reasons.

4 Network analysis

In the present section we turn to studying how peer influence is channeled among the treated entrepreneurs through the social connections they establish within their respective interaction group. Our primary objective is to estimate the sign and magnitude of those network effects, understanding as well how they are affected by the different treatment arms. To do so, of course, we first need to construct the social network out of the exhaustive communication data available. This will be the task undertaken in Subsection 4.1, where we also provide a preliminary descriptive information on the most prominent network characteristics resulting from the construction.

4.1 Network construction

Constructing the peer entrepreneur network requires specifying a way in which links are to be "operationalized" in terms of our data – in other words, we need to determine how to identify and measure the informational content flowing between any two entrepreneurs. Our approach will consider different complementary routes to it. The one we shall use as our benchmark formulation identifies a link from i to j as a message by i that is followed by another message by j such that, under certain conditions, can be interpreted as representing some feedback to i's message. More precisely, we posit:

(LF) (a) A *directed link* $i \to j$ is established for every message that i has posted in a particular channel of the chatting platform such that j has subsequently written another message in the same channel.[15] In general, we impose the requirement that j's subsequent message is written"not too long after" the message by i, as captured by a parameter τ to be described precisely below.

(b) When a link $i \to j$ exists, its *weight* is identified with the number of sentences included in the message sent by i that marks the start of the communication exchange. Then, adding the weight over all such links $i \to j$, we obtain what is defined as the *aggregate interaction flow* directed from i to j.

As suggested above, the simple idea captured by (LF) is that if j has been active on a certain channel shortly after i wrote a message on it, j has been exposed to (and probably read) the content of i's message. Therefore, there is some significant probability that whatever i wrote in that channel motivates j's subsequent message. In this sense, j's message can be viewed as a reaction to (or feedback on) the earlier i's message.

Clearly, the way in which we identify and measure communication between two entrepreneurs, i and j, can be conceived only as an approximate (and thus imperfect) assessment of their actual exchange of information and feedback. A more precise measurement of this flow would necessarily involve, *inter alia*, a detailed analysis of the content of the messages themselves. This, however, would require a semantic analysis of the messages that goes beyond what we have been able to carry out so far. It seems nevertheless quite clear that the NLP methods that we have used as yet for a more limited purpose (see Section 5) can be extended to more ambitious objectives, as the one described. This is indeed our plan for future research.

[15] Refer to Subsection 2.3 for a description of the chatting platform used in the experiment.

Still, one extension of our approach that we do explore in this paper pertains to the way in which the weight of a link is determined. For, in general, one may argue that a limitation of (LF) is that, whereas it quantifies the intensity of communication in one direction (i.e. on the formulation of the problem or the request of feedback), it is silent on the richness of the response. A simple way (and hence, still, also imperfect) to account for these reciprocal considerations is analyzed in Subsection 4.3.3, where the extent of communication across two agents is measured in both directions. We advance that, despite this extension having a significant impact on the induced network weights, our main identification of the peer effects is robust to this two-way specification.

Next we provide a concrete formulation of the phrase "not too long after" that has been informally used in (LF). As indicated, this notion will be parametrized in terms of some τ, understood as the maximum "delay" that may separate the two messages involved in the declared link. In fact, such a delay is not measured time-wise but in terms of the number of *intermediate* sentences that have been posted in the channel, thus detaching one from the other. Such post-indexed "time" appears to be a better choice than calendar time for our purposes in that the number of intermediate postings reflects how far down the timeline of j this entrepreneur finds i'message (hence affecting its call on j's attention). The maximum lag τ is a modeling parameter that, as it decreases, has the obvious effect of (weakly) reducing the density of the induced network. As we shall show in Subsection 4.3.2, our results are robust to (i.e. essentially unaffected by) this parameter, provided it is not too small. For the sake of focus, we defer addressing such robustness concerns to Subsection 4.3, thus carrying out our present analysis within what we shall label the *benchmark setup*. This is a setup where the network is constructed under (LF) with no bound contemplated on the communication delay, i.e., formally, for $\tau = \infty$.

Next, we introduce some useful notation. Let the matrix $M = (m_{ij})_{i,j=1}^{n}$ capture the pattern of directed communication across every two agents , each entry m_{ij} standing for what we have called the aggregate interaction flow directed from j to i (which, as indicated in part (b) of (LF), is given the total weight of the links $j \to i$). When no link $j \to i$ exists, we simply make $m_{ij} = 0$. Thus, for each $i = 1, 2, ..., n$, the ith row of matrix M embodies an absolute measure of the influence exerted on the specific entrepreneur i by all her peers $j \neq i$. From it, one can readily derive the matrix $G = (g_{ij})_{i,j=1}^{n}$ of *relative influence* by normalizing each row of M so that, for every pair of entrepreneurs $i, j \in N$,

$$g_{ij} = \frac{m_{ij}}{\sum_{j=1}^{n} m_{ij}} \tag{2}$$

so that $\sum_{j \neq i} g_{ij} = 1$ for all $i = 1, 2, ..., n$. Such a matrix G defines the adjacency matrix of the weighted influence/peer network on which we base our ensuing analysis.[16] A graphic illustration of the peer networks arising in our context is provided in Figures 4 and 5 below for two peer groups based in Nigeria, one interacting in a virtual-within scenario and another in a virtual-across one.

We close this introductory subsection by presenting, for the constructed (benchmark) network, a summary account of some canonical measures as well as their cross correlation. To focus on the truly engaged part of the population, throughout the whole of the present section we confine our analysis to the entrepreneurs who

[16] In general, one may be interested in accounting not only for relative weights but also for absolute intensities. Since the two perspectives – relative and absolute – may indeed be relevant, a richer specification would have to integrate both of them in some suitably combined manner. We choose, however, to abstract from these considerations in the present paper.

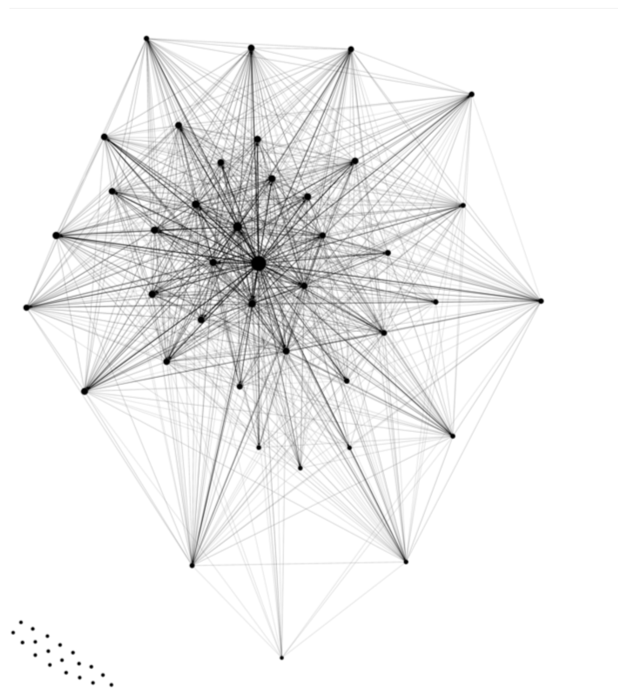

Figure 4: **Peer network for a virtual-within interaction group based in Nigeria.**

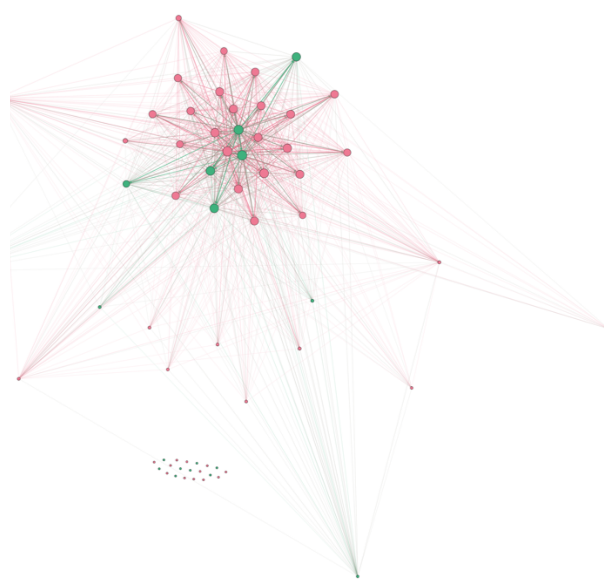

Figure 5: **Peer network for a virtual-across interaction group where red nodes correspond to entrepreneurs from large countries whereas those colored green correspond to entrepreneurs based in small countries.**

had some activity in the peer platform, among those whom we have called M0-completers, i.e. those who completed Milestone 0 on time. Platform-active entrepreneurs represented 66% of M0-completers, amounting to a total of 1082 entrepreneurs. For short, those entrepreneurs will simply be called hereafter the *active* ones.

The first table below (Table 8) displays some basic statistics on the following node-based information:[17]

(a) two reciprocal notions of *connectivity*, each reflecting one of the opposite directions in which peer influence flows, i.e. into a node (in-degree), or out from it (out-degree);[18]

(b) *clustering*, which captures the extent to which connections are transitive, i.e. whether an entrepreneur who influences another one, also influences those entrepreneurs the latter influences.

(c) three measures of *centrality* (pagerank, betweenness, and closeness), each assessing in a different way how crucial/prominent is the role played by an entrepreneur in globally connecting pairs of other nodes.

Table 8: **Network statistics: description**

	in-degree	out-degree	clustering	pagerank	betweenness	closeness
mean	443.5	443.5	0.039	0.036	0.052	0.627
stdev	731.6	1013.1	0.030	0.040	0.073	0.162
median	137.0	89.0	0.031	0.018	0.025	0.639
min	0.0	0.0	0.003	0.000	0.000	0.000
max	5445.0	10371.0	0.301	0.512	0.193	1.000

Notes: The table presents some basic statistics of selected network measures for entrepreneurs that delivered M0 on-time and also were active on the interaction platform. See the main text for a concise description of each network measure.

The statistics displayed in Table 8 show moderately left-skewed[19] distributions – except for closeness, which is slightly right-skewed. They also display moderate (normalized) variances, thus suggesting that mean values may be regarded as quite representative statistics for a typical entrepreneur. Concerning connectivity, in particular, an average of 443.5 sentences per entrepreneur "effectively" received/sent from/to peers represents a quite intense flow of communication (cf. Footnote 18).

Finally, Table 9 reports the bilateral correlations among the variables considered in Table 8.

The correlations described in Table 9 are quite intuitive. First, we find that agents who are quite engaged in writing messages/sentences[20] are also so in effectively receiving those written by others.[21] Entrepreneur connectivity (as measured by the in- and out-degree) is also quite highly correlated with the first two centrality measures considered (pagerank and closeness) but largely uncorrelated with the third one (betweenness).

[17]See Vega-Redondo (2007), Jackson (2008), and Bloch *et al.* (2017) for a formal description of these measures and their interpretation. Recall that, in the present context, the networks under consideration are weighted ones, with the weight g_{ij} being identified with the total number of sentences included in the posts sent by entrepreneur i that, according to (LF), suitably qualify as communication flowing from i to j.

[18]Note that, by an accounting identity, total/average in-degree must be exactly equal to total/average out-degree.

[19]Here, we identify skewdness in the simple-minded manner that declares a distribution to be left/right skewed if its median is below/above its mean.

[20]Since sentences or full messages represent alternative ways of capturing the same notion of communication intensity, henceforth we shall generally use the term "message" to refer to one or the other, depending on the context. This should generate no confusion.

[21]This, of course, is partly a consequence of how we identify links in constructing the network. It is not difficult, however, to construct simple examples showing that, in general, the in-degree and out-degree of individuals can be very different.

Table 9: **Network statistics: correlations**

	in-degree	out-degree	clust'ing	pagerank	closeness	betw'ness
in-degree	1					
out-degree	0.792	1				
clustering	0.377	0.368	1			
pagerank	0.590	0.517	0.490	1		
closeness	0.475	0.427	0.265	0.453	1	
betweenness	0.069	0.055	-0.093	0.163	0.293	1

Notes: The table presents the correlations among some selected network measures among active entrepreneurs (see the main text for a concise description).

Such a contrast is a reflection of the following twin observation. On the one hand, the first two centrality measures rely on ***all*** path connections between two entrepreneurs and hence the in-degree of a node (as well as its out-degree, due to the aforementioned correlation) tend to enhance its centrality. Instead, betweenness considers only ***shortest*** paths between nodes and therefore, in general, it may well happen that an entrepreneur with a low connectivity plays a major role in globally connecting others from the betweenness viewpoint.

4.2 Network peer effects

Now we turn to what is the main aim of this section: the estimation of the network-based peer effects under the different treatment arms considered in our experiment. We start in Subsection 4.2.1 by discussing how we deal with some important identification issues, while the estimation itself of the core peer effects is conducted in Subsection 4.2.2. The robustness of our formulation and econometric approach is studied in Subsection 4.3.

4.2.1 The identification problem: reflection and homophily

As explained, an important problem that needs to be tackled in estimating network peer effects in our context is one of identification. There are two main concerns in this respect. One is related to the well-known reflection problem raised by Manski (1993), which is a consequence of the fact that the choices and outcomes of peers are jointly determined endogenously. This renders it difficult in some cases to identify the structural parameters from the reduced form specification of the model. Another concern pertains to the confounding effect of homophily. In this case the problem derives from the fact that the network itself is also endogenous in our case (i.e. it is the object of choice by individuals). Therefore, any observed correlation of behavior among peers might be not a result of network-mediated influence but of the tendency to interact with others who are alike in some relevant respect – for example, in our case, similar in age, operating sector, or education level. For, as the ancient proverb goes, "birds of feather flock together."

The approach we pursue here to tackle the reflection problem follows Bramoullé *et al.* (2009) and De Giorgi *et al.* (2010). In essence, it relies on the fact that, if the network structure is irregular enough, it is possible to "instrument" peers' behavior in terms of the exogenous characteristics impinging on the behavior of second neighbors – i.e. entrepreneurs who are (strictly) two steps away. Such instrumenting can then be used to identify the underlying peer effects. Of course, the key assumption here is that every influence across entrepreneurs has to be mediated by the network, i.e. through interaction among those who are connected. Then, the second neighbors of an entrepreneur can only affect her through intermediate peers, and this in turn implies that the usual exclusion restriction is satisfied, which is key for successful identification.

Now we describe matters formally, as applied to our context. Denote by $\boldsymbol{y} = (y_i)_{i \in N}$ the outcome profile over the set N, the population sample of interest. The three outcomes that will be *separately* considered here are the same ones used in the treatment analysis: submission, extensive business quality, and the intensive one (recall Section 3 and see Subsection 4.2.2 below for a more detailed reminder). Let $G = (g_{ij})_{i,j=1}^{n} \in \mathbb{R}_{+}{}^{n \times n}$ be the network adjacency matrix, where n is the cardinality of N. A typical entry g_{ij} stands for the weight of the link directed from i to j (which, as will be recalled, is given by the number of messages written by i that can be "effectively" read by j).

As regressors in the estimation, besides the (endogenous) outcome variables of peers as listed in $\boldsymbol{y}$, we also include the set of m (exogenous) baseline variables given by the individual characteristics obtained for each participant from the initial survey. They are arranged in the matrix $X = (x_{i\ell})_{i,\ell} \in \mathbb{R}_{+}{}^{n \times m}$ where a typical entry $x_{i\ell}$ stands for the ℓth characteristic of individual i.

Having introduced the required notation, the two-stage structure of the IV estimation considered here can be formally described in a compact manner as follows:

$$\boldsymbol{y} = \alpha + G\boldsymbol{y}\,\beta + X\boldsymbol{\gamma} + GX\boldsymbol{\xi} + \varepsilon \tag{3}$$

$$G\boldsymbol{y} = \varrho + G^2 X\boldsymbol{\lambda} + X\boldsymbol{\delta} + +GX\boldsymbol{\varsigma} + \nu, \tag{4}$$

where ε and ν stand for stochastic noise satisfying the standard conditions. The estimation procedure relies on using the variables in G^2X (the total influence weights associated to the peers lying two links away) as respective instruments for $G\boldsymbol{y}$ (the influence weights bearing on each entrepreneur, associated to her direct peers).

Besides the vertical intercepts α and ϱ, the coefficients to be estimated in (3)-(4) are the following:

- $\beta \in \mathbb{R}$ is the peer effect reflecting how $G\boldsymbol{y}$ – the (G-weighted) average of (direct-)peers' outcomes estimated from (4) – affects own outcomes;
- $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, ..., \gamma_m)'$, $\boldsymbol{\delta} = (\delta_1, \delta_2, ..., \delta_m)' \in \mathbb{R}^m$ are the column vectors of regression coefficients associated to the m own baseline covariates being considered in the two stages of the estimation;
- $\boldsymbol{\xi} = (\xi_1, \xi_2, ..., \xi_m)'$, $\boldsymbol{\varsigma} = (\varsigma_1, \varsigma_2, ..., \varsigma_m)' \in \mathbb{R}^m$ are the column vectors of regression coefficients associated to the m baseline covariates of direct peers in the two stages of the estimation;
- $\lambda = (\lambda_1, \lambda_2, ..., \lambda_m)' \in \mathbb{R}^m$ is the column vector of regression coefficients associated to the m own baseline co-variates of second-order neighbors, as weighted by the entries in the matrix G^2;

The framework given by (3)-(4) defines the basic model within which we shall undertake the estimation of peer effects. As explained in Bramoullé *et al.* (2009), if G^2, G, and the identity matrix I are linearly independent (a condition satisfied in our setup), the above framework suitably identifies the peer-effect coefficient β of interest. From an econometric viewpoint, however, the IV estimation requires two further assumptions:

(**ER**) The instruments satisfy the familiar exclusion restriction, which in our case amounts to the requirement that all peer influence be channeled through the network.

(**NE**) The peer network is either exogenous or, if endogenous, the biases induced by the link-formation process can be suitably controlled for.

The essential idea underlying (ER) is that, if peer influence can only flow between entrepreneurs who interact in the peer network (i.e. communicate through the chatting platform), the way in which any given entrepreneur can be affected by second-order peers is by these indirect peers influencing first her direct ones. This then allows one to rely on second-order peer's characteristics (which are exogenous) as instruments generating exogenous variability on the behavior (and thus influence) of first-order peers.

Turning now to the second condition above, (NE), we can of course *not* claim that in our case the network is exogenous, for its endogeneity is indeed one of the distinctive features of the experimental setup. Therefore, what we put forward instead is the claim that, in view of the quite wide range of individual characteristics gathered at baseline from the survey and the course (recall Table 1), most of the correlation in peer performance induced by network-formation forces (e.g. homophily) can be suitably explained in terms of those baseline (exogenous) co-variates. Controlling for them, therefore, in the econometric specification, the estimated peer effects should be largely free of the bias that might otherwise be induced by the endogeneity of the network.

4.2.2 The core peer effects

The analysis of peer effects reported here focuses on estimating the impact of peer communication on the three same dimensions/outcomes that have been used throughout to measure entrepreneurial performance (see Section 3):

(a) submission,

(b) the extensive quality margin (accounting for the combined effect on submission and business quality),

(c) the intensive quality margin (measuring the effect on business quality among submitted proposals).

For all these three outcomes, we again partition the analysis into the following three interaction contexts, each one reflecting a different way in which virtual interaction takes place in our experiment:[22]

(i) virtual-within for large countries, with interaction taking place in nationally homogeneous groups,

(ii) virtual-across for large countries, with interaction taking place in nationally heterogeneous groups,

(iii) virtual-across for small countries, with interaction taking place in nationally heterogeneous groups.

Overall, therefore, the peer network effects will be estimated in *nine* different cases, each of them matching a corresponding counterpart in the treatment analysis conducted in Subsections 3.1, 3.2.1, and 3.2.2. As explained, initially, the peer network used for the analysis will be the one we have labeled the benchmark one (cf. Subsection 4.1), which imposes no upper bound on the communication lag (i.e. $\tau = \infty$) and

[22]For the sake of focus, we exclude the Uganda sample under virtual-within interaction since the effects are in line with those observed for large countries under this type of interaction

determines the link weights as specified in (LF). Also recall that our analysis focuses on those entrepreneurs who completed Milestone 0 on time and were active on the platform (66% of the M0 completers).[23]

To gain a sharper assessment of peer effects in our context, we find it useful to start with a specialized version of the general econometric specification (3)-(4) that restricts to peer effects to be only outcome-based, thus abstracting from the network counterpart of what we have called "composition effects" in the treatment analysis (cf. Section 3.3). The motivation is that, in the presence of outcome-based peer influence, composition-channeled influence would be largely redundant (since it is highly correlated) and hence obscure the estimation analysis.[24] Indeed, support for this view is obtained in Subsection 4.3.1 where, within the formal specification given by (3)-(4), we separately estimate their network composition effect for the two characteristics that, intuitively, seem most susceptible to display significant such effects: baseline quality and experience. We find that neither of them displays a significant effect while outcome-based influence carries the full, and highly significant, weight.

Thus, in view of the previous considerations, we start studying the specialized version of (3)-(4) in which it is posited that $\boldsymbol{\xi} = \boldsymbol{\varsigma} = \mathbf{0}$. The estimated peer effects are displayed in Table 10 below, while for a full description of the results, with estimates corresponding to all co-variates involved, we refer the reader to Tables A10-A12 in the Appendix.

In trying to understand the pattern of effects displayed in Table 10, it may be useful to put forward, in a heuristic and largely conjectural manner, the polar forces that might be at work in the different scenarios:

(a) Peer interaction in a homogeneous environment, because of the familiarity it provides (hence leading, for example, to common social norms and references as well as general trust) has a supporting/encouraging effect on participation/involvement.

(b) Peer interaction in a heterogeneous environment, because of the diversity it entails (hence exposing individuals to more novelty in knowledge and perceptions) promotes innovation/creativity.

(c) In contrast to (a), and for converse reasons, interaction in a homogeneous environment should not favor innovation.

(d) In contrast to (b), and for converse reasons, interaction in a heterogeneous environment should be hardly supportive.

In combination, (a)-(d) may be used to rationalize the results displayed in Table 10. First, one may view the subsample vw-LC as displaying the highest homogeneity in interaction. Thus, in line with (a) and (c), we find that peers effects are positive and highly significant for submission but insignificant for intensive quality (among submitted proposals). Then, the positive and significant effect on the extensive margin can be interpreted as a manifestation of the effect on submission.

[23] Given that active participation in the platform is an endogenous decision of entrepreneurs, this poses, in principle, issues of possible selection bias analogous to those studied in Subsection 3.2.2. The key observation to make, however, is that in the present case the selection mechanism is quite different from the one operating in our analysis of the treatment effect. For here the selection into "activity" is implemented by entrepreneurs with no prior differential exposure to any exogenous feature of the experiment. In view of this, therefore, one should not expect any systematic bias induced. Indeed, a support for this claim follows from the fact that the fraction of active individuals in both subarms is approximately equal to 68%.

[24] In fact, a reason why this should be a particularly true in our case is that individuals are identified in an anonymous manner (i.e. through an artificial username) and none of the characteristics underlying the aforementioned network composition effects are observed at the start of the interaction.

Table 10: **Network peer effects, benchmark setup**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Peer effect	.533***	.014	.098
	(.160)	(.132)	(.289)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Peer effect	.327**	.060	.537***
	(.128)	(.125)	(.132)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Peer effect	-.031	.087	.476***
	(.112)	(.140)	(.149)
Number of entrepreneurs: 779			

Notes: Estimated network-based peer effects for the benchmark setup ((LF) and $\tau = \infty$) on three different outcomes: the *submission* decision, *extensive* business quality, and *intensive* one. The estimation is jointly conducted for the three subsamples: v-within and v-across for large countries and v-across for all small countries, restricted to those entrepreneurs who completed Milestone 0 on time. We control for all baseline individual characteristics, using as well their second-neighbor values as instruments in a corresponding two-stage OLS. For a full account of the estimated coefficients, see Tables A10-A12 in the Appendix. Standard errors are clustered at the group level, which gives rise to 44 clusters. The number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

On the other hand, the subsample va-SC is the one where peer interaction displays the most cross-country heterogeneity. Therefore, as suggested by (b) and (d), we find that the effect on submission is not significant whereas that on intensive quality is positive and highly significant. Again, the positive and significant effect on the extensive margin must be interpreted as a consequence of the similar such effect obtained for intensive quality.

Finally, the subsample va-LC appears to embody, in a sense, the worse of the two former cases. The interaction is not homogeneous enough to support submission, nor is it sufficiently diverse to promote innovation. In both cases – submission and intensively measured quality – peer effects are found to be positive but non-significant (even at the 10% level). Naturally, this has the consequence of producing as well an insignificant effect on the extensive margin.

The former heuristic discussion suggests the need to explore, systematically, the type of communication unfolding in each of the treatment subarms of our experiment. Much could be learned from such analysis, integrating it with the study of the various incentives and protocols involved in the design of peer interaction mechanisms. Some preliminary work along these lines is described in Section 5, where we summarize the NLP methodology widely used for that purpose and apply it to the analysis of how the pattern and content of communication relate to individual characteristics, networking, and performance. As we shall discuss, some interesting patterns already arise from this initial effort, but substantial further research is needed to achieve the full potential.

We end this section with a remark bearing on the effect of interaction and competition on trust.

Remark **1 (Interaction and trust)** *In item (a) above, we highlight general trust as one of the leading consequences that familiarity, and hence easier communication, can generate in a peer group. The relationship between communication and general trust has been prominently discussed in the management literature, in particular concerning the functioning of virtual teams (see e.g. Sarker et al. (2011)). There is, however, another aspect of trust that also seems quite relevant for networks of entrepreneurs that is not considered by this literature. It concerns the fact that, when peers are entrepreneurs, they are also possible competitors. This, of course, would run against the incentives to share "sensitive" information.*

We conjecture that such a consideration may be at work in our setup. Specifically, it could underlie the fact that in the SCS context (under virtual-across interaction) peer effects are significant on intensive quality, while they are not so under virtual-within interaction. A plausible conjecture for this may be that, even if in the latter context some general trust arises because of easier communication, the content of such information may be general rather than specific. That is, it may avoid concrete details that can be used by peers, which in this context are much more likely to be potential competitors as well. This could partly explain why peer effects under virtual-within interaction are not significant while they are so when the groups are very fragmented along nationality lines. In the latter case, the likelihood that a peer is also a real competitor are much more remote.

4.3 Robustness

In this section we check the robustness of our results to extensions of our framework in three different directions. First, in Subsection 4.3.1, we show that adding network composition effects to the econometric specification considered in Section 4.2 does not alter the essence of our results. Second, in subsection 4.3.2, we explore the implications of allowing the maximum communication lag τ to be finite, possibly relatively short. Third, in Subsection 4.3.3 we study whether our results are affected by accounting for amount of information flowing both ways between every pair of entrepreneurs, i and j.

4.3.1 Network-composition effects

As a counterpart of our analysis of peer-composition effects for the treatment effect (cf. Subsection 3.3), here we investigate whether analogous composition effects arise concerning some baseline characteristics of network peers. To this end we rely on the general econometric specification posited in Section 4.2.2. For the reasons discussed there, a full-fledged joint consideration of all possible such effects would not be useful, so we decide to focus our analysis on the same two variables that were considered for the treatment analysis: baseline quality of peers and their business experience (or lack thereof). These two variables are also the two ones where the possibility of network composition effects seem most plausible and consequential.

Formally, the econometric model to be studied is:

$$\boldsymbol{y} = \alpha + G\boldsymbol{y}\,\beta + X\boldsymbol{\gamma} + \xi\, G\boldsymbol{z} + \varepsilon \tag{5}$$

$$G\boldsymbol{y} = \varrho + G^2 X\boldsymbol{\lambda} + X\boldsymbol{\delta} + \varsigma\, G\boldsymbol{z} + \nu, \tag{6}$$

where the notation used is just as in 4.2.1, with $\boldsymbol{z}$ listing the data on the peer characteristic (one at a time, as opposed to the full set of baseline co-variates) whose composition effect is estimated. The parameters ξ and ς are the corresponding coefficients in the first and second stage of the estimation procedure, respectively. Obviously, the above system represents the specialization of the general specification in which all except one of the components in each of the vectors $\boldsymbol{\xi}$ and $\boldsymbol{\varsigma}$ in (3)-(4) is postulated to be equal to 0.

Tables 11 and 12 below display the results for each of the two baseline characteristics being considered, business quality and experience, respectively. For each case, we report only the outcome-based peer effect, the own effect associated to the characteristic in question, and the corresponding network composition effect.

Table 11: **Network-composition effects: baseline business quality**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Peer effect	.569***	.177	.101
	(.192)	(.187)	(.302)
Own baseline quality	.071***	.055*	.015
	(.025)	(.030)	(.039)
Peers' baseline quality	-.018	-.059	-.003
	(.035)	(.051)	(.053)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Peer effect	.363**	.115	.547***
	(.153)	(.168)	(.140)
Own baseline quality	.308***	.302***	.186
	(.098)	(.109)	(.138)
Peers' baseline quality	-.066	-.097	-.116
	(.154)	(.173)	(.206)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Peer effect	.006	.096	.503***
	(.127)	(.148)	(.155)
Own baseline quality	.148*	.219***	.210**
	(.085)	(.078)	(.091)
Peers' baseline quality	-.067	-.017	-.164
	(.108)	(.128)	(.170)
Number of entrepreneurs: 779			

Notes: Estimated network-composition effects for the same network setup, underlying conditions, subsamples, and methodological approach as in Table 10, except that one additional regressor is added: the network-weighted average of peers' baseline qualities. For comparison, the estimated coefficient for the corresponding own variable is also included, besides that for the core peer effect. As usual, the number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

Several observations are in order. First, we note that the pattern of outcome-based peer effects are unchanged in both cases, as compared to the results reported in Section 4.2.2 (cf. Table 10). The same applies to own effects (cf. Tables A10-A12 in the Appendix). Thus, we find that incorporating network composition effects into the econometric specification has no effect on the previous estimates – in particular, those of outcome-based peer effects. The key reason, of course, is apparent from Tables 11 and 12 themselves. They show that composition-based effects are not significant if they are considered in combination with the outcome-based peer effects. We may conclude, therefore, that the former are subsumed in the latter – or, in other words, that peer effects in our context are channeled through (endogenous) outcomes rather than (exogenous) individual characteristics.

Table 12: **Network-composition effects: business experience**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Peer effect	.545***	.016	-.072
	(.148)	(.146)	(.321)
Own running business	-.011	.083*	-.042
	(.049)	(.042)	(.077)
Peer running business	-.029	-.002	.173
	(.100)	(.094)	(.127)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Peer effect	.335***	.073	.483**
	(.117)	(.132)	(.209)
Own running business	-.004	.465**	-.318
	(.189)	(.185)	(.290)
Peer running business	-.083	-.072	.272
	(.387)	(.368)	(.637)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Peer effect	-.002	.067	.509**
	(.103)	(.161)	(.210)
Own running business	-.009	.322**	-.363
	(.132)	(.149)	(.325)
Peer running business	-.256	.096	-.166
	(.319)	(.394)	(.643)
Number of entrepreneurs: 779			

Notes: Estimated network-composition effects for the same network setup, underlying conditions, subsamples, and methodological approach as in Table 10, except that one additional regressor is added: the network-weighted average of the peers' dummies indicating whether they currently run their own business. For comparison, the estimated coefficient for the corresponding own dummy is also included, besides that for the core peer effect. As usual, the number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

4.3.2 Parametrizing the communication lag

As explained in Subsection 4.1, our benchmark setup as one having the communication/interaction lag unbounded, i.e. $\tau = \infty$. In view of the the starkness of this assumption, we have analyzed how the results – in particular, the estimated peer effects – vary as the value of τ varies within a *finite* range. The outcome of this robustness exercise is diagrammatically described in Figures 6-8. Each individual figure shows, for ***one*** of the three dimensions considered here (submissions, the extensive quality margin, or the intensive one), how the estimated peer-effect coefficient and its corresponding confidence interval at 5% significance level change as a function of τ. This is displayed in three separate panels (a)-(c) for the following three different interaction scenarios: virtual-within for large countries, virtual-across for large countries, and virtual-across for small countries.

It is in the nature of the procedure used in network construction that higher values of τ must lead to a denser network. This in turn implies that, as displayed in Figures 6-8, the coefficient estimates and corresponding confidence intervals progressively stabilize as τ grows. Besides this simple observation, two are the main features we want to highlight from these diagrams:

(a) For a value of $\tau = 80$ (i.e. a maximum communication lag of 80 "intermediate" sentences in the chat's timeline) the estimated peer coefficient reaches in every case (i.e. for each of the three outcomes and for all three interaction scenarios) a value very close to the coefficient estimated in Table 10 for $\tau = \infty$. Indeed, we find (although, for the sake of focus, we do not show it in the diagrams) that for all values $\tau \geq 8$ the estimation results remain essentially unchanged around those obtained for the benchmark setup with $\tau = \infty$.

(b) Even for lower values of τ – in particular, it is enough that $\tau \geq 40$ – the statistical significance of the coefficients achieved in the benchmark setup already obtains in the present parametrized setup, with a slight exception: the case of extensive quality in large countries where such significance strictly consolidates only for $\tau \geq 80$ (c.f. panel (a) of Figure 7).

The former two points are reassuring. They indicate that our results are quite robust to how the network is constructed, at least concerning the "immediacy condition" required for effective interaction/communication. Another robustness test on network construction is conducted in the following subsection.

4.3.3 Two-sided communication flows

As suggested in Subsection 4.1, one may argue that the influence that an entrepreneur i exerts on another one j might depend on the stream of messages flowing in both directions. To capture this idea, we presently consider the following variation on the link-formation rule (LF) considered in the benchmark setup.

($\widehat{\mathbf{LF}}$): (a) A *directed link* $i \to j$ exists as postulated in (a) of (LF).

(b) When a link $i \to j$ exists, its weight is identified as in (b) of (LF). In contrast, however, the (symmetric) *aggregate interaction flow* across i and j is defined to be the aggregate weight over all links $i \to j$ ***and*** $j \to i$.

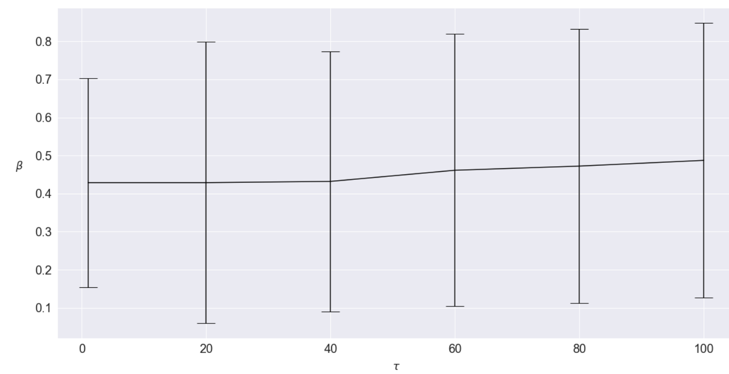

(a) Submission: virtual-within interaction, large countries

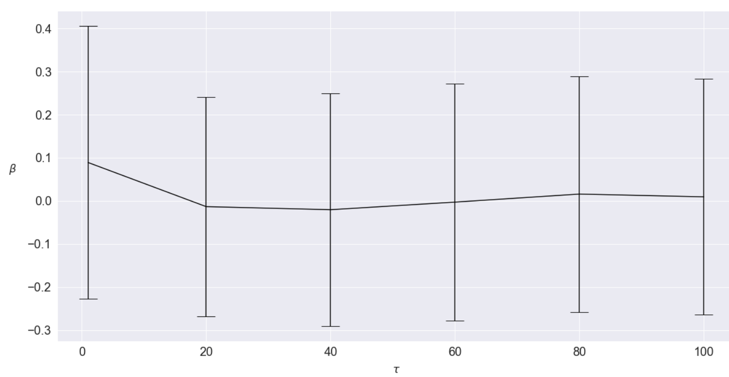

(b) Submission: virtual-across interaction, large countries

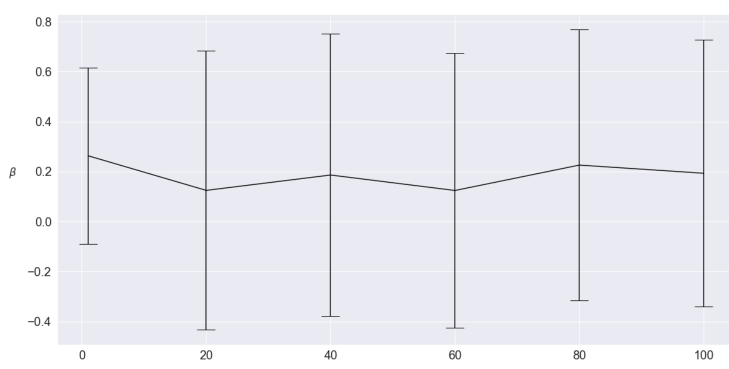

(c) Submission: virtual-across interaction, small countries

Figure 6: The diagram shows how the peer effects applied to ***submission*** depend on the upper bound τ on the maximum lag between constituent messages defining a link. The coefficient estimate and 95% confidence intervals are traced for each $\tau \in [1, 100]$ in three cases: the Large-Country Sample under the two treatment arms (virtual-within and virtual-across) and the Small-Country Sample under the virtual-across arm.

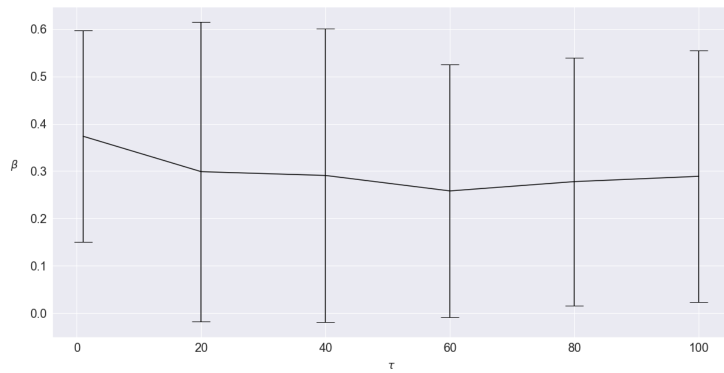

(a) Extensive quality margin: virtual-within interaction, large countries

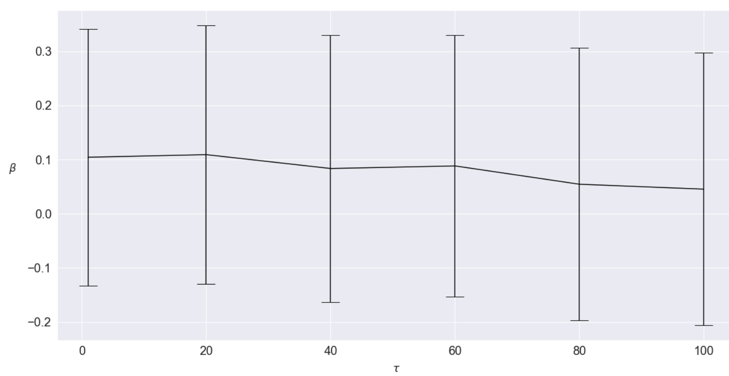

(b) Extensive quality margin: virtual-across interaction, large countries

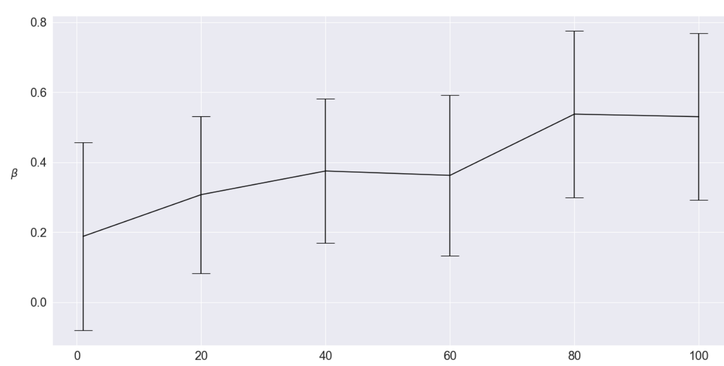

(c) Extensive quality margin: virtual-across interaction, small countries

Figure 7: The diagram shows how the peer effects applied to the ***extensive quality margin*** depend on the upper bound τ on the maximum lag between constituent messages defining a link. The coefficient estimate and 95% confidence intervals are traced for each $\tau \in [1, 100]$ in three cases: the Large-Country Sample under the two treatment arms (virtual-within and virtual-across) and the Small-Country Sample under the virtual-across arm.

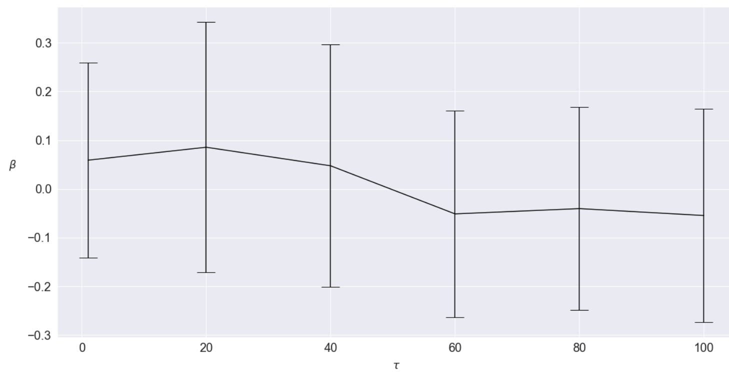

(a) Intensive quality margin: virtual-within interaction, large countries

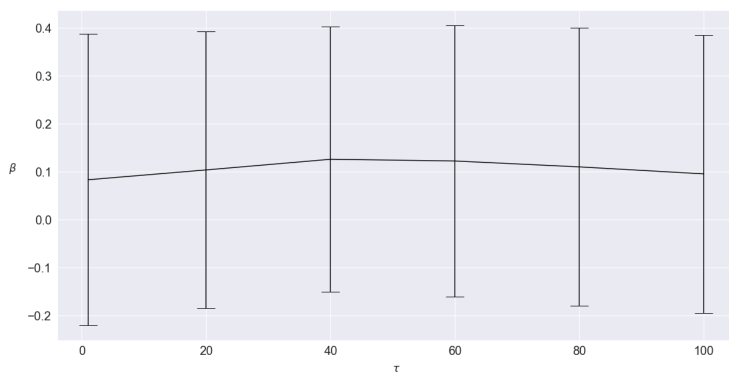

(b) Intensive quality margin: virtual-across interaction, large countries

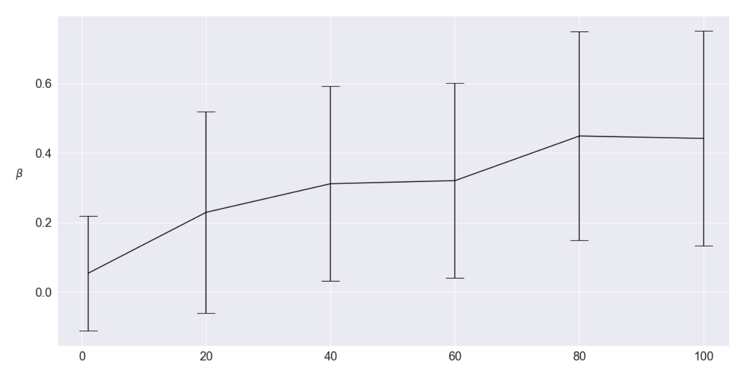

(c) Intensive quality margin: virtual-across interaction, small countries

Figure 8: The diagram shows how the peer effects applied to the ***intensive quality margin*** depend on the upper bound τ on the maximum lag between constituent messages defining a link. The coefficient estimate and 95% confidence intervals are traced for each $\tau \in [1, 100]$ in three cases: the Large-Country Sample under the two treatment arms (virtual-within and virtual-across) and the Small-Country Sample under the virtual-across arm.

As in the benchmark setup, we focus for simplicity on the case where the interaction lag $\tau = \infty$. Then, on the basis of (b) above – and, in particular, its revised *bidirectional* formulation – we are led to a (symmetric) matrix of interagent influence

$$\hat{M} = (\hat{m}_{ij})_{i,j=1}^{n} = M + M^{\top} \tag{7}$$

where M is the matrix of directed interaction constructed in Subsection 4.1 under formulation (LF) and $M^{\top}$ stands for the transpose of M. Normalizing the matrix $\hat{M}$ over rows we arrive at the matrix $\hat{G} = (\hat{g}_{ij})_{i,j=1}^{n}$ in which, as a counterpart of (2), we have:

$$\hat{g}_{ij} = \frac{\hat{m}_{ij}}{\sum_{j=1}^{n} \hat{m}_{ij}}. \tag{8}$$

With this adjacency matrix (not typically symmetric) now representing the pattern of influence across entrepreneurs, we re-estimate the peer effects through the same econometric model used in Section 4.2, with G replaced by $\hat{G}$. The results are displayed in Table 13 below.

Table 13: **Network peer effects, bidirectional influence**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Control for ***all*** baseline info.	.869***	-.327	-.014
	(.135)	(.256)	(.395)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Control for ***all*** baseline info.	.472***	-.045	.495***
	(.140)	(.182)	(.170)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Control for ***all*** baseline info.	.123	-.075	.362**
	(.135)	(.167)	(.176)
Number of entrepreneurs: 779			

Notes: Estimated network-based peer effects for the bidirectional setup that differs from the benchmark setup is considered in Table 10 only in that the link-formation rule ($\widehat{\text{LF}}$) substitutes the former (LF). (Hence, in particular, we continue to focus on a context with $\tau = \infty$.) As before, the peer effects are estimated for three different outcomes: the *submission* decision, *extensive* business quality, and *intensive* one. Furthermore, the estimation is separately conducted for the the following three subsamples: v-within and v-across for large countries and v-across for all small countries, restricted to those entrepreneurs who completed Milestone 0 on time. We also control for all baseline individual characteristics, using as well their second-neighbor values as instruments in a corresponding two-stage OLS. Standard errors are clustered at the group level, which gives rise to 44 clusters. The number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

Comparing Tables 13 and 10, we find that the qualitative pattern of (significant) peer effects is exactly the same under ($\widehat{\text{LF}}$) and (LF), although the magnitude of the corresponding coefficients differs somewhat. This happens despite the fact that the matrix M is highly non-symmetric, and therefore it displays an influence structure that is quite different from that of $\hat{M}$. More precisely, we have computed that the correlation between the entries of M and $M^{\top}$ is 0.39, hence quite below 1. Thus, even though the communication across agents is quite "conversational" (being the aforementioned correlation significantly higher than 0), it is also far from being balanced between peers. As we have seen, however, such imbalances do not change the qualitative pattern of peer effects in the entrepreneur peer network. In this sense, we can regard our benchmark analysis as robust to a reconsideration of the network construction that accounts not only for the communication density of the party that initiates the conversation but also of the party that reacts to it.

5 The semantic analysis

This section is organized as follows. First, in Subsection 5.1 we provide a description of the annotation and machine-learning procedures that – as developed by the field of Natural Language Processing (NLP) – have been used to extract "meaning" from the large body of communication undertaken by entrepreneurs. Then, in Subsection 5.2, we turn to identifying the patterns of how such stream of semantic content is related to individual or group characteristics and their corresponding entrepreneurial performance. Finally, in Subsection 5.3, the objective is to bring the NLP analysis to bear on the study of the peer network. This involves combining the information on social connections we already have with the knowledge we shall gain on the semantic content that flows through those connections. Such an approach will then help us understand whether alternative types of content have different effects on how the network operates – most importantly, on the strength of peer effects. At the end of the present section, the point to be hopefully gathered is that, even though the semantic analysis we report here is still preliminary, the NLP methodology being used has much to contribute to a deeper understanding of how peer networks operate. In turn, this should allow a significant improvement of their design and a consequent increase in overall performance.

5.1 Natural Language Processing: the annotation setup

At a basic level, semantic analysis aims to undertake a pair of twin tasks:
(a) classify the data along text-inherent semantic categories (for example business-relatedness),
(b) provide insights into the structure, relationships, and social embedding of the semantic content.
In our experimental context, we have over 140,000 sentences written by entrepreneurs throughout the duration of the intervention, each of them being conceived (and also labeled) as a separate message. This, of course, amounts to much more than what can be feasibly analyzed manually by even a relatively large team of researchers. Thus, instead, we have instead relied on the NLP techniques, by now quite standard, which are succinctly described in what follows.

The first step has been to use machine-learning classifiers to infer semantic category labels for the entire communication data available. This entailed having five (human) coders/annotators go through a large sample of the text in order to identify and label a number of specified categories. Among the categories considered, the following three proved to be the most informative:

- ***Business content***: a *binary label* in $\{0, 1\}$, indicating whether the message concerns business or some professional activity.
- ***Sentiment***: measured on a five-point scale $\{-2, -1, 0, 1, 2\}$ that stand, respectively, for the following set of *ordered labels*: strongly negative, critical, neutral, encouraging, and enthusiastic.

- ***Audience scope***: a *binary label* in $\{0,1\}$, specifying whether the message is, respectively, addressed to a single person or to a larger group.

Overall, we had the coders label a random sample of 10,500 of messages out of the whole collection of them available. A subset of 500 messages were annotated by all five coders, so that we could check the agreement among them. Such an agreement varied between categories, but was high overall as measured by a value of 0.67 for the statistic κ, commonly used in NLP.[25] Then, using a Bayesian model of annotation as in Hovy *et al.* (2013), we calculated the *individual* reliability of each coder, which was again high throughout. This indicates that the categories and labels used may be reliably applied to train a classifier.[26]

To train the machine-learning classifier, we used the random forest algorithm (see Breiman (2001)), which fits several decision tree models on the data and aggregates over them. We then estimated the performance of the model via 5-fold cross-validation, in order to get a sense of its ability to predict, in a reliable manner, the inferred labels on unseen data (i.e. alternative one-fifth samples extracted from the whole data set available, using the remaining four fifths to train the classifier). In this case, the reliability of label prediction is typically measured by the F_1 score, which is the harmonic mean of two conflicting desiderata: precision and recall, each of these calculated per class label.

Specifically, ***precision*** $\mathcal{P}$ is defined as

$$\mathcal{P} = \frac{TP}{TP + FP}$$

where TP is the total weight that the classifier associates to the true positives and FP stands for that associated to false positives, while ***recall*** $\mathcal{R}$ is given by

$$\mathcal{R} = \frac{TP}{TP + FN}.$$

To understand these measures, note that if a model were to classify everything as positive for a particular label it would get a perfect recall for it – after all, it would find all true positives while not producing any false negatives). However, such a model would obviously be useless, since its precision would be very low. Thus, in practice, we want to balance precision and recall and the F_1 score does exactly that by taking the harmonic mean the two of them:

$$F_1 = \frac{2\,\mathcal{P}\,\mathcal{R}}{\mathcal{P} + \mathcal{R}}.$$

The F_1-performance of the machine-learning classifier was found to be high in the training sample ($F_1 \geq 0.8$) for our three leading categories: business content (denoted by b), sentiment (s), and audience scope (a). This led us to extend it to the whole collection of messages written by the entrepreneurs, thus assigning to every message a corresponding weight for every label in each of the three categories in $X = \{b, s, a\}$.

In the end, therefore, the classifier assigns to every pair of

[25]This statistic measures inter-coder agreement. It is given by the ratio $\frac{A_o - A_e}{1 - a_e}$, where A_o is the observed agreement and A_e is the agreement that could be expected by pure chance. The indicated value of 0.67 is generally interpreted as a good extent of agreement.

[26]Heuristically, coder reliability is a measure of internal consistency of the annotation conducted by each separate coder. In this sense, therefore, it hinges on the internal consistency of each coder, in contrast with the cross-coder consistency measured by the aforementioned statistic κ.

- a message m in the set M of all messages sent throughout the experiment,
- a label r in the set L_x of labels associated to each category x,

a weight/probability $\zeta(r,m) \in [0,1]$. This weight indicates the strength/confidence with which the classifier posits/predicts that label r is present in message m. These weights are normalized, for each message m and each category x, so that $\sum_{r\in L_x} \zeta(r,m) = 1$. Overall, the vector

$$\boldsymbol{\chi} = \left\{ \left[(\zeta(r,m))_{r\in L_x} \right]_{x\in X} \right\}_{m\in M} \tag{9}$$

compactly describes the *primitive data* on which the semantic analysis is to be undertaken.

5.2 Communication patterns

In this section we undertake the analysis of the data described in (9). The endeavor is organized in three parts. First, in Subsection 5.2.1, we introduce the different semantic distributions, individual- and message-based, that underlie the analysis. Then, Subsection 5.2.2 carries out the core of our semantic analysis by studying how the labels of the different categories interact and how they are related to both individual characteristics and measures of performance. Finally, in Subsection 5.2.3, we turn to identifying the terms that best signal our three categories of interest: sentiment, business content, and audience scope.

5.2.1 Semantic distributions

Our semantic analysis adopts two complementary perspectives. One of them is entrepreneur-based and associates, to each entrepreneur, the average semantic content of the messages she sends. In contrast, the alternative perspective is message-based and therefore more "granular": it focuses on the semantic content of each message, independently of the entrepreneur who sends it. We introduce each of these alternative approaches in turn.

(a) Entrepreneur-based distributions:
To construct the entrepreneur-based distributions, we must associate, for each entrepreneur $i \in N$ and each of the three categories of interest (business focus, sentiment, and audience scope), a corresponding scalar index that suitably summarizes the overall (average) weights displayed by the messages sent by that entrepreneur in that category. To this end, first we define, for each category $x \in X$, the average x-intensity displayed by any given m as follows:

$$\omega_x^m = \sum_{r\in L_x} \zeta(r,m)\, r. \tag{10}$$

Then, for every particular entrepreneur $i \in N$, we compute her average index for this category:

$$\bar{\omega}_x^i = \frac{1}{|M_i|} \sum_{m\in M^i} \omega_x^m \tag{11}$$

with a population average given by

$$\mu_x^N = \frac{1}{|N|} \sum_{i \in N} \bar{\omega}_x^i. \tag{12}$$

Naturally, the population (node-based) distribution $\{\bar{\omega}_x^i\}_{i \in N}$ over each category x induces a density over the corresponding category index. Figure 9 represents such density functions for each of our three leading semantic categories (business content, sentiment, and audience scope). For visual clarity, here and throughout this section we present continuous counterparts of the density functions that rely on the non-parametric smoothing procedure called *Kernel Density Estimation* (KDE).[27]

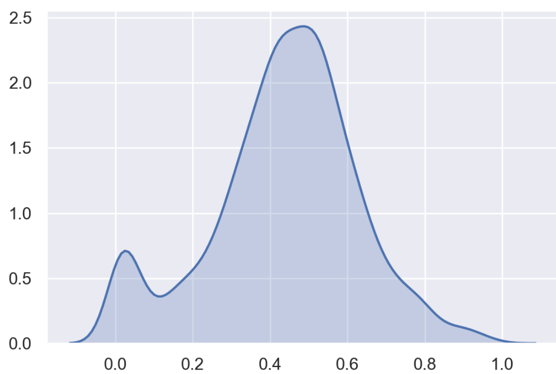

(a) Business

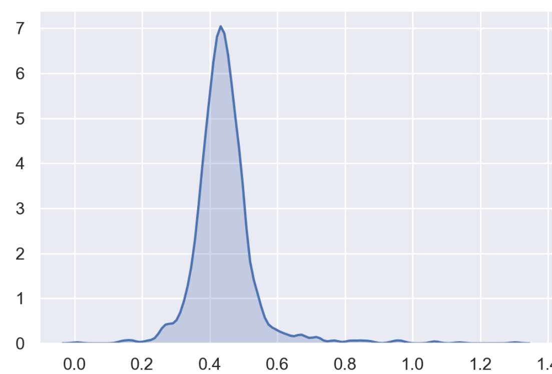

(b) Sentiment

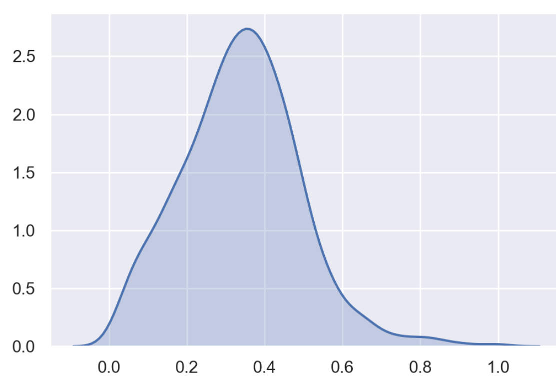

(c) Audience

Figure 9: ***Node-based*** **densities over semantic categories: business content, sentiment, audience scope.**

[27]This procedure, originating in the work of Rosenblatt (1956) and Parzen (1962), can also be viewed as a way of estimating an underlying continuous density from a (discrete) sample density. From the different versions of this procedure considered in the statistical literature, we consider the Gaussian implementation (in which the smoothing is conducted through a Standard Normal) with a so-called "bandwidth" given by the so-called Scotts Rule – see Scott (1992).

Figure 9 shows that the density of both (entrepreneur-average) sentiment and audience scope are unimodal, the former more narrowly concentrated around the mode than the latter. Instead, the corresponding density over entrepreneurs' business orientation is bimodal with two modes (local peaks) around 0 and 0.5, the latter being substantially more prevalent than the former. This implies that whereas the population is relatively homogeneous in the former two categories, in the latter there is some significant heterogeneity, dichotomous but asymmetric. This dichotomy indicates that, concerning the extent of business message content, the population is split into two subgroups: a smaller subgroup with essentially no business content and a larger one with a significant business content in their messages. These two separate subgroups are quite homogeneous, in contrast with the other two categories where it is the the population as a whole that displays such an homogeneity. A brief account of basic mean and dispersion statistics for each case is provided in Table 15.

Table 14: **Semantic node-based statistics**

	mean	standard deviation
business content	0.44	0.19
sentiment	0.44	0.09
audience focus	0.34	0.15

Given the substantial homogeneity displayed in Figure 9 at the entrepreneur level, one may wonder whether there is a a similar counterpart homogeneity at the message level – or, to put it differently, whether entrepreneurs also span only a narrow diversity across their own individual. To answer this question, we turn our consideration to the distributions defined at the message level.

(b) Message-based distributions:
From (10) we can directly consider the message-based distribution $\{\omega_x^m\}_{m\in M}$ for each category $x \in X$, whose average is:

$$\bar{\omega}_x^M = \frac{1}{|M|} \sum_{m\in M} \omega_x^m. \tag{13}$$

The induced densities for the index associated to each of our three leading semantic categories are displayed in Figure 10.

There are two interesting points to be gathered from Figure 10. First, we observe that the substantial homogeneity found in Figure 9 when the data were aggregated at the entrepreneur level in fact masks significant heterogeneity across the whole body of messages exchanged by the population. This suggests that, even though entrepreneurs do send messages that vary in semantic content in all three categories, the range of communication undertaken by each individual entrepreneur tends to display a quite similar balance among business and non-business focus, positive and neutral sentiment, or general and individual audience. As we shall see in Subsection 5.3, a similar feature arises in our analysis of semantically weighted networks but in a substantially stronger manner – i.e. not just when the messages are aggregated at the level of each entrepreneur but also when it is applied at the finer (and therefore more demanding) level of the ***effective*** volume of interaction flowing through the links established between each ***pair of entrepreneurs***.

The second point to be highlighted highlight from Figure 10 is that, for each of the three categories being considered, the NLP analysis is able to identify a number of modes (local density peaks) equal to the labels being respectively considered in each case: two of them for the dichotomous categories of business

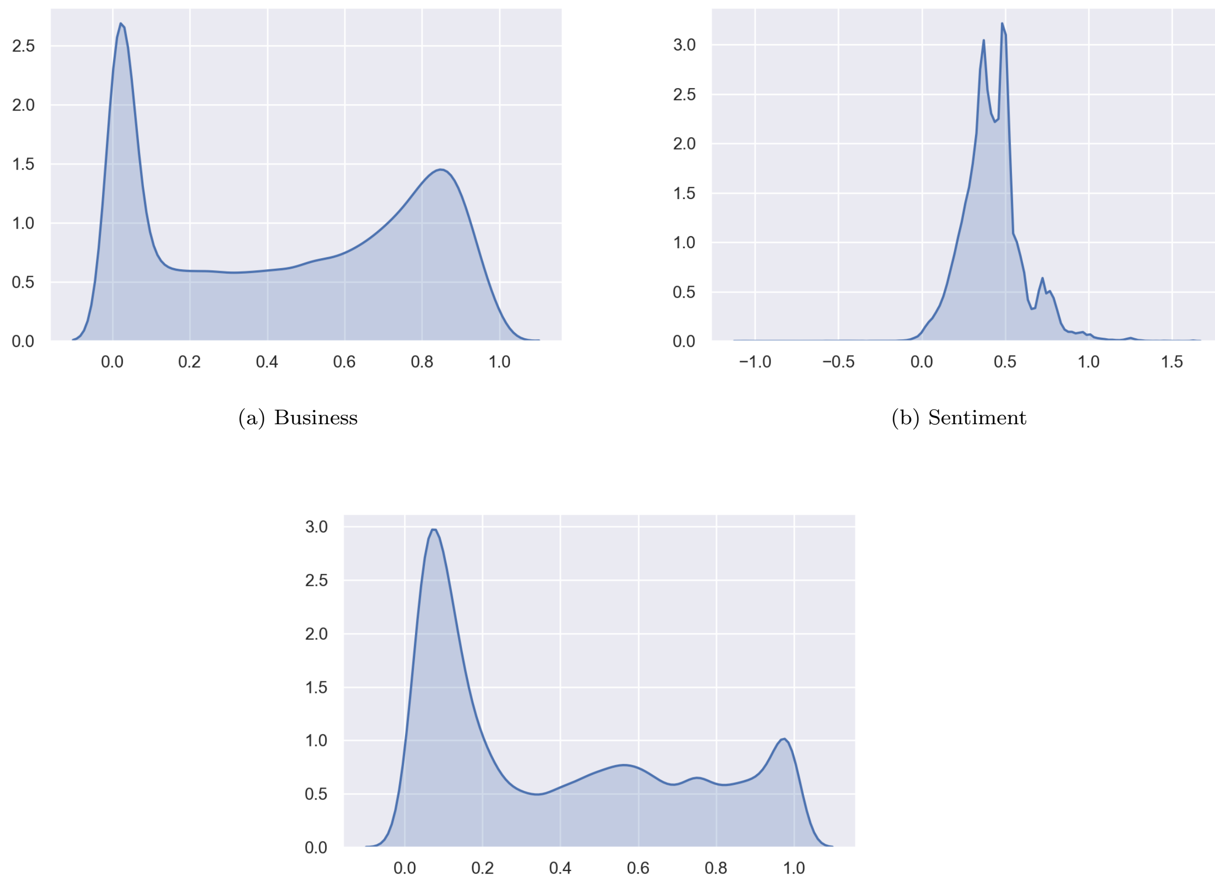

(a) Business

(b) Sentiment

(c) Audience

Figure 10: ***Message-based*** **densities over semantic categories: business content, sentiment, audience scope.**

and audience, and three for the multilateral one of sentiment. This, in a sense, provides a reassuring source of support for the methodological approach pursued. For business and audience the two modes identified are far apart, which indicates that the NLP classifier delivers very sharp message identification in this case. Instead, the induced identification is less clear-cut for sentiment, the three distinct levels singled out being relatively close to each other in this case These levels are to be associated to the labels 0, 1, and 2 since, as explained, negative instances are essentially non-existent in the data. Thus, as expected from the initial annotation process, the higher mode (associated to 2) turns out to be much infrequent than the other two. As before, to end our discussion of the semantic message-based distributions, we provide in Table 15 a brief account of basic mean and dispersion statistics for our three leading categories.

Table 15: **Semantic message-based statistics**

	mean	standard deviation
business content	0.45	0.34
sentiment	0.43	0.18
audience focus	0.39	0.33

5.2.2 Conditional relationships

Preliminaries and notation

The distributions defined in the preceding subsection are, of course, marginal distributions of an underlying joint distribution that involves not only the three categories (business, sentiment, and audience) but also entrepreneurs' exogenous (e.g. baseline) characteristics as well as the final outcomes (concerning, in particular, submission and quality performances). Naturally, we are most interested in identifying some of the patterns of cross-dependence displayed by the induced conditional distributions. Specifically, we shall focus on the following cases.

(i) The relationship between every two pair of different categories, x and x', as induced by their joint (marginal) distributions. Extending previous notation, these distributions will be denoted by $\{\bar{\omega}^i_{xx'}\}_{i\in N}$ when the data is aggregated at the entrepreneur level, and $\{\omega^m_{xx'}\}_{m\in M}$ when they are given at the message level.

(ii) The relationship between a certain category x and some exogenous entrepreneur characteristic φ. The corresponding joint distribution will be denoted by $\{\bar{\omega}^i_{x\varphi}\}_{i\in N}$.

(iii) The relationship between a certain category x and some entrepreneur performance measure π. The corresponding joint distribution will be denoted by $\{\bar{\omega}^i_{x\,\pi}\}_{i\in N}$.

In some of the previous cases – in particular, when two different categories x and x' are involved as in (i) – we shall find it interesting to compute the corresponding Pearson correlation coefficients. These coefficients can be computed either in an entrepreneur-based or a message-based manner, which will be respectively denoted by $\rho^N(x, x')$ and $\rho^M(x, x')$. They will provide a compact scalar measure of whether, say, business focus and neutral sentiment are associated, or whether messages that are targeted for a general audience tend to have a weak business content and/or a highly positive sentiment.

In other cases, however, no such correlations can be meaningfully defined – in particular, when qualitative variables are involved. Then, it will be useful to compare the distributions themselves, as they are associated to two alternative values of the variable ϑ of interest – a variable that can either concern a baseline characteristic (e.g. gender), assigned treatment arm (say, virtual across) or performance (e.g. submission). To be precise, let $\tilde{\vartheta}$ and $\hat{\vartheta}$ represent two alternative such values of a qualitative variable, e.g. male and female, or submission and not. For simplicity suppose that, as in the examples, the whole population of entrepreneurs can be partitioned in terms of two values and denote the two corresponding classes as $\tilde{N} \equiv \{i \in N : \varphi^i = \tilde{\vartheta}\}$ and $\hat{N} \equiv \{i \in N : \varphi^i = \hat{\vartheta}\}$. Then, the comparison between the corresponding distributions, $\{\bar{\omega}^i_x\}_{i\in\tilde{N}}$ and $\{\bar{\omega}^i_x\}_{i\in\hat{N}}$ for any given category x can be used to understand how the entrepreneurs' communication on that category is associated to the qualitative variable under consideration – be it some exogenous characteristic, a given treatment arm, or a particular performance measure of the entrepreneur population.

When the conditioning variables are discrete variables (often dichotomous, as in the former illustrations), the former construction is straightforward. The procedure is not so clear, however, if the aim is to compare distributions that reflect co-dependence between two categories, x are x' – e.g. business content and sentiment, or sentiment and audience. For, in this case, the NLP classifier assigns, for any category $x \in X$ and every message m, a continuous weight $\zeta(r, m) \in [0, 1]$ for each $r \in L_x$. Thus, in order to undertake an analogous co-dependence analysis in this case, we shall proceed as follows. Consider any two categories, $x, x' \in X$, and suppose the focus is on how different values of x affect the distribution of x'. In order to remain in the simple realm of dichotomous comparisons, we shall split the population distribution over x into the set of individuals whose value for x is above some statistic of the original distribution (we will use the median) and those whose value is below it. Then, denoting by N_x^+ and N_x^- the respective subsets, a direct adaptation of the previous approach simply leads to a comparison of the distributions $\{\bar{\omega}^i_{x'}\}_{i \in N_x^+}$ and $\{\bar{\omega}^i_{x'}\}_{i \in N_x^-}$.

In what follows, the analysis is organized into the three kinds of relationships among variables that may be considered in our context:

(a) cross-semantics relationships,

(b) semantics-characteristic relationships,

(c) semantics-performance relationships,

each of these cases three separately addressed in turn. Before turning to it, we simply note that an approach fully analogous to that described can be used when the data is studied at the finer message level. In this case, the partitions to be used must be defined, *mutatis mutandis*, on the whole set of message M.

(a) Cross-semantics relationships

Naturally, a first approach to studying the cross-semantics (or cross-category) conditional relationships is to compute their Pearson correlation coefficients – both at the message and node/entrepreneur levels – which have been denoted by $\rho^M(x, x')$ and $\rho^N(x, x')$. They are presented in Tables 16 and 17 below.

Table 16: **Cross-category correlations: message-based approach**

	business content	positive sentiment	private audience
business content	1		
positive sentiment	-0.477950	1	
private audience	-0.157837	0.243414	1

Notes: The table presents the correlations across categories at the message level.

Table 17: **Cross-category correlations: entrepreneur-based approach**

	business content	positive sentiment	private audience
business content	1		
positive sentiment	-0.479776	1	
private audience	-0.017388	0.095115	1

Notes: The table presents the correlations across categories at the entrepreneur/node level.

A first interesting observation is that the pattern of correlations is very similar at both the message and entrepreneur levels. We also see that business content and sentiment are negatively correlated, and quite strongly so, which suggests that most messages have a rather one-dimensional focus and the same applies to individual entrepreneurs. On the other hand, we find a mild negative correlation (at both levels) between business content and audience, an indication that messages with significant business content tend to aim, weakly, for a relatively wide reach. Finally, we note that, as one would expect, when messages display a strong (positive) sentiment, they tend to be directed to individual peers rather than the whole peer group at large.

A more thorough understanding of the relationship among the different categories can be gathered from a detailed analysis of the corresponding distributions. As before, this analysis can be conducted at the message-based or entrepreneur-based levels. For the sake of focus, we first undertake here a message-based analysis, while relying instead on an entrepreneur-based one in the other two cases (category-characteristic and category-performance) in which focusing on the entrepreneur as the "unit of analysis" seems most natural.

We start by studying how the extent of business content embodied by the messages, or their audience focus (general or private), is related to the (positive) sentiment intensity they contain. The corresponding KDE-smoothed densities are graphically displayed in the two panels of Figure 11. They show how these densities – defined over business and audience – depend, in a purely *statistical* sense, on the sentiment exhibited by those messages. As already explained, this is done by comparing two different distributions: one conditional on the sentiment level of messages m being above the median (or, using former notation, conditioned on $m \in M_s^+$) and another analogous distribution conditioned on messages lying below the median (i.e. $m \in M_s^-$). That is, more formally, for each case (business and audience), we compare the following pairs of distribution:

(i) $\{\omega_b^m\}_{i \in M_s^+}$ and $\{\omega_b^m\}_{i \in M_s^-}$

(ii) $\{\omega_a^m\}_{i \in M_s^+}$ and $\{\omega_a^m\}_{i \in M_s^-}$

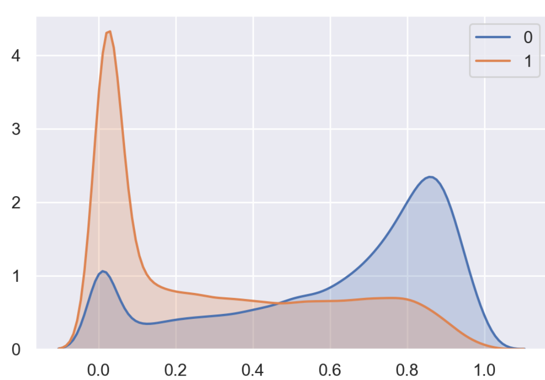

(a) Business content of messages conditional on sentiment

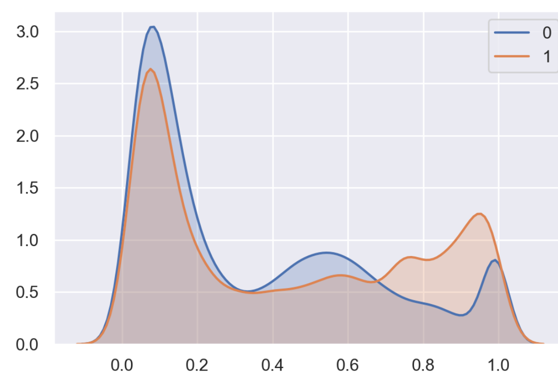

(b) Audience focus of messages conditional on sentiment

Figure 11: The diagrams compare the densities over the business content of messages (panel (a)) and their audience focus (panel (b)), conditional on messages having a sentiment weight higher than the median (indicator 1) or below it (indicator 0).

Figure 11 shows that the density over business content conditional on low sentiment dominates in the first-order stochastic dominance (FOSD) sense that conditional of high sentiment. In contrast, concerning audience focus, we find that higher sentiment leads to a higher private focus (also relying on the strong FOSD criterion). Even though these conclusions are unambiguously implied by the diagrams shown above, one must bear in mind that the densities depicted there are smooth counterparts of the underlying discrete densities. This suggests conducting explicitly a rigorous statistical test of the aforementioned conclusions.

Among the various tests that exist in the statistical literature, we shall rely on one of the most standard: the (non-parametric) Kolmogorov-Smirnov (KS) test,[28] as adapted by Heathcote *et al.* (2010) to test ***FOSD relationships*** between pairs of probability distributions. Essentially, this approach takes any two sample cumulative distribution functions (cdf's) $\tilde{F}, \tilde{G} : [0,1] \to [0,1]$ and compares them by conducting simultaneously two tests on the following pair of null hypotheses:

$(\mathbf{H}_0')\ \ \forall z \in [0,1],\ F(z) \le G(z)$

$(\mathbf{H}_0'')\ \ \forall z \in [0,1],\ G(z) \le F(z)$

where F and G are the underlying (unknown) distributions for the reference population. Thus, while $\mathbf{H}_0'$ states that F dominates G in the FOSD sense, $\mathbf{H}_0''$ asserts the converse relationship. According to the KS test, the statistic to be considered in each case is given by the corresponding maximum (algebraic) differences between the two sample cdf's – that is, $\Delta_{\tilde{F}\tilde{G}} \equiv \max_{z\in[0,1]} \tilde{F}(z) - \tilde{G}(z)$ and $\Delta_{\tilde{G}\tilde{F}} \equiv \max_{z\in[0,1]} \tilde{G}(z) - \tilde{F}(z)$, respectively. To be more precise, let $\alpha \in (0,1)$ be a selected significance level for the test, while p_0' and p_0'' denote the p-values associated to the nulls $\mathbf{H}_0'$ and $\mathbf{H}_0''$, respectively. Then, if the first null is not rejected (i.e. $p_0' \ge \alpha$) but the second is rejected (i.e. $p_0'' < \alpha$) the KS test does **not** reject the hypothesis (or, heuristically, accepts the possibility) that F FOSD-dominates G.

For future reference, let us also advance that, in some cases, we shall also be interested in applying the standard KS test, as applied to test whether two cdf's, F and G, can be accepted as equal. That is, in this case the null hypothesis is:

$(\mathbf{H}_0)\ \ \forall z \in [0,1],\ F(z) = G(z)$

For this test, the sample statistic to be used is $\Delta_{|FG|} \equiv \max_{z\in[0,1]} |\tilde{F}(z) - \tilde{G}(z)|$ – i.e. the *sup*-norm applied to the sample cdf's. The corresponding p-value of the test will be denoted by p_0.

To proceed systematically on the different types of questions addressed in this section, we shall find it useful to rely on the following mnemonic notational conventions. Let u, v be indices that, depending on the context, may stand for two categories x and x', or a category x and an entrepreneur characteristic φ, or a category x and a performance measures π. Then, $\boldsymbol{p}_{u/v} = (p_0, p_0', p_0'')$ will denote the three dimensional vector of p-values associated to the hypotheses-testing exercises described above – $\mathbf{H}_0$, $\mathbf{H}_0'$, and $\mathbf{H}_0''$, respectively – corresponding to the cdf's F and G specified as follows:

- F is the cdf of the variable u conditional on a low value of v,
- G is the cdf of the same variable u conditional on a high value of v.

[28]See Kolmogorov (1933) and Smirnov (1939).

When v is dichotomous, the interpretation of what is a high or low level will be immediate; instead, when v is a continuous variable, it will be interpreted as suggested before: “high” will mean above the median and “low” below it.

With these conventions in place, we may return to the cases considered in Figure 11, with u standing for the business and audience categories b and a, and v being the sentiment category s. Then, for the business content of messages conditional on sentiment, we find that

$$\boldsymbol{p}_{b/s} = (0.0\,,1.0\,,0.0) \tag{14}$$
$$\boldsymbol{p}_{a/s} = (0.0\,,0.0\,,1.0). \tag{15}$$

These p-values provide formal support to our former discussion on how the business content of messages and their audience focus are oppositely related to the sentiment level exhibited by those messages (negatively in the first case, positively in the second). A diagrammatic counterpart of this twin statement is found in Figure 12, which displays a comparison between the corresponding cdf’s in each case.

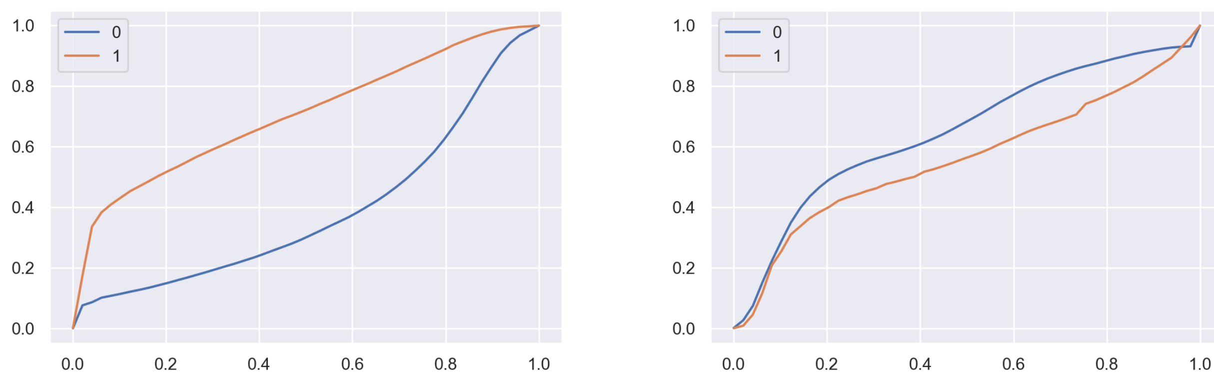

(a) Message business conditional on sentiment

(b) Message audience conditional on sentiment

Figure 12: The diagrams compare the cdf’s defined over the business content of messages (panel (a)) and their audience focus (panel (b)), conditional on messages having a sentiment weight higher than the median (indicator 1) or below it (indicator 0).

Next, we conduct an analogous exercise for the four other types of cross-relationship among categories that still remain to be considered:

- business content and sentiment conditional on audience,
- sentiment and audience conditional on business content.

The situation is illustrated graphically in Figures 13 and 14 in terms of the corresponding cdf’s, while the corresponding p-values are as follows.

For the conditional dependence on the business of the other two categories we obtain:

$$\boldsymbol{p}_{s/b} = (0.0\,, 0.99\,, 0.0) \quad (16)$$
$$\boldsymbol{p}_{a/b} = (0.0\,, 0.0\,, 0.0) \quad (17)$$

and for the dependence on audience focus:

$$\boldsymbol{p}_{b/a} = (0.0\,, 0.03\,, 0.0) \quad (18)$$
$$\boldsymbol{p}_{s/a} = (0.0\,, 0.0\,, 0.99). \quad (19)$$

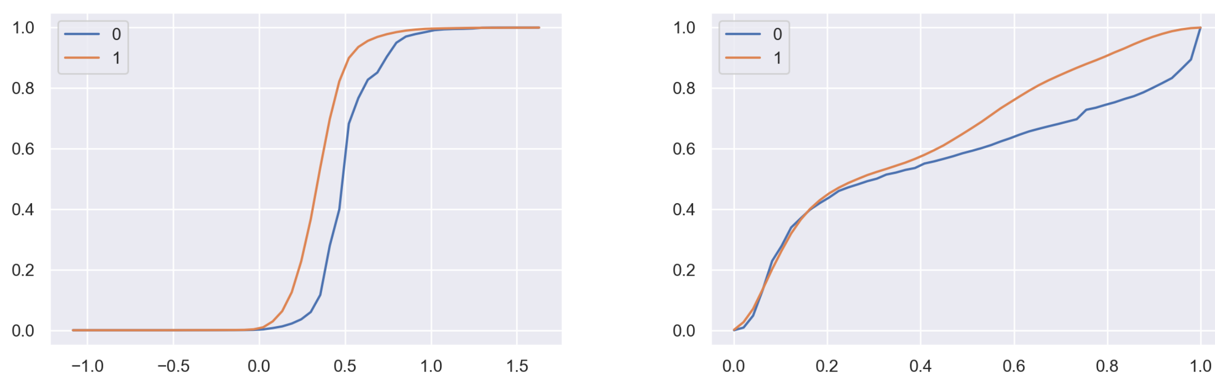

(a) Message sentiment conditional on business content (b) Message audience conditional on business content

Figure 13: The diagrams compare the cdf's defined over the sentiment of messages (panel (a)) and their audience focus (panel (b)), conditional on messages having a weight over business content higher than the median (indicator 1) or below it (indicator 0).

Combining all of the above, the main conclusions can be summarized as follows.

1. In everyone of the six cases (category pairs), we can *reject* that the conditioning category does *not* affect the corresponding conditional distribution.

2. We find support (cannot reject) the following statements/hypotheses:

   (a) high business content is conditionally associated to low/neutral sentiment in the FOSD sense, in both directions.

   (b) positive sentiment is conditionally associated to individual message focus in the FOSD sense, in both directions.

3. We find no support for a FOSD relationship between business content and audience focus.

The above conclusions do not exhaust the range of interesting points that can be derived from a more thorough analysis of the data. As an illustration, we may want to study further the reciprocal relationship between different categories and consider, for example, the important pair given by business content and

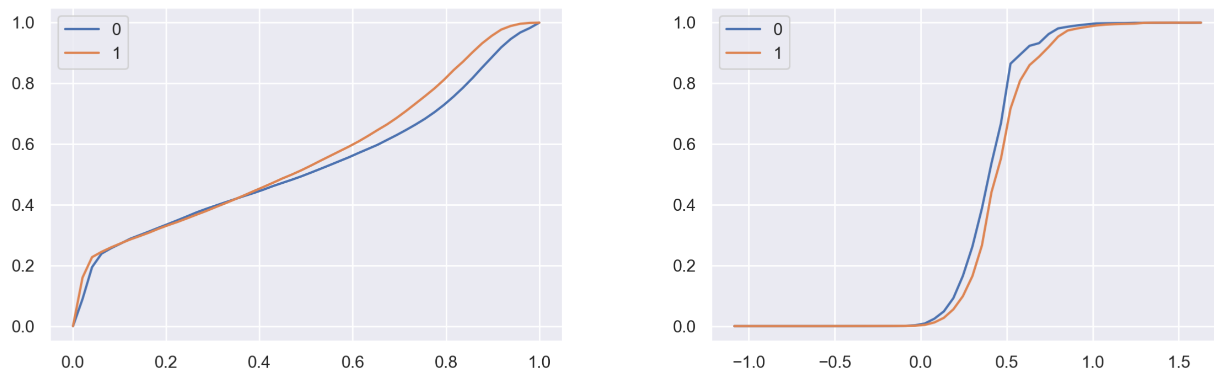

(a) Message business content conditional on audience focus

(b) Message sentiment conditional on audience focus

Figure 14: The diagrams compare the cdf's defined over the business content of messages (panel (a)) and their sentiment (panel (b)), conditional on messages having a weight on individual audience above the median (indicator 1) or below it (indicator 0).

sentiment. As stated in 2a above, we know that both categories display an inverse two-way relationship. It is interesting, however, to understand better this relationship by comparing the conditional densities in either direction. When the conditioning event is high/low sentiment, the corresponding densities on business content were depicted in Figure 11a and display a rough symmetry in each case (high or low sentiment). In contrast, the reciprocal case is quite different, as found in Figure 15 below. While, as anticipated, we have that the distribution associated to low sentiment dominates (in the FOSD sense) that associated to high sentiment, the induced distribution in each case are hardly symmetric. Specifically, we find that the density associated to high sentiment is unimodal and sharply centered around a "middle average," whereas the one associated to low sentiment inherits the three modes that characterize the unconditional sentiment distribution (cf. Figure 10b).

(b) Semantics-performance relationships

Now we study how the semantic patterns relate to our leading performance measures: submission, extensive quality, and intensive quality. As shown in what follows, again we find that their are strong FOSD relationships between the semantic content of messages and many of those measures of performance. We follow in this respect a parallel approach to that pursued before: first we present the KS tests that address the question through the usual hypothesis-testing methods, then complementing this statistical methods with a diagrammatic representation of the corresponding cdf's that describe the phenomenon in a more intuitive manner. Note that, in this case, since our different measures of performance are defined at the entrepreneur level, the analysis is undertaken at that level as well. Thus, in particular, the semantic content of messages is aggregated at the entrepreneur level, as explained in the preliminary discussion opening this section.

The p-values that embody the KS testing of the hypotheses of interest are displayed in Table 18, for each the the three aforementioned measures of entrepreneur performance and our three leading semantic categories: business content, sentiment, and audience focus. Here, we adapt previous notation and have F_0 and F_1 stand for the cdf's conditional on low and high performance, respectively, where "low" and "high"

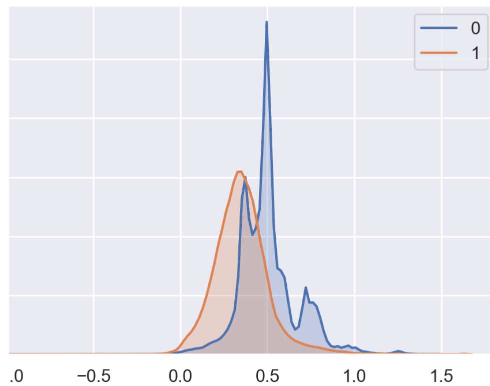

Figure 15: The diagram compares the density over the sentiment content of messages, conditional on messages having a business weight higher than the median (indicator 1) or below it (indicator 0).

are defined relative to the median if the performance is not dichotomous (i.e. for extensive and intensive quality). We observe a very clear-cut pattern. On the one hand, from the left column in that table we note that the hypothesis that the business content, sentiment, and audience of entrepreneurs' messages is unrelated to performance – in any of the three dimensions considered: submission, extensive quality, or intensive quality – can be strongly rejected. Such a relationship, however, is very different for each of the three semantic categories, as the middle and center columns of Table 18 show and we explain next.

First, for business content, since the KS test supports the hypothesis that F_1 FOSD dominates F_0 for each of the performance dimensions, the implication is that the subpopulation of high-performing entrepreneurs display a strong tendency to send messages with a high business content. This statement is illustrated in Figure 16.

Second, for sentiment, the conclusion is even stronger (the p-values are significantly higher) but the subpopulation of entrepreneurs that display a high sentiment in their messages are those who are low-performing since the hypothesis supported in this case is that F_0 FOSD dominates F_1. A graphical reflection of this conclusion is found in Figure 17.

Third, concerning audience, we have the statistically strongest pattern, which is analogous to that displayed by business content. That is, high-performing entrepreneurs produce messages whose target is individual rather than general. The diagrammatic counterpart is shown in Figure 18.

Table 18: **Semantics-performance relationship: p-values on distributional hypotheses**

Null hypotheses	$F_0 = F_1$	$F_0 \leq F_1$	$F_1 \leq F_0$
Business content			
Submission	0.0	0.0	0.18
Quality: extensive margin	0.0	0.0	0.29
Quality: intensive margin	0.0	0.01	0.55
Sentiment			
Submission	0.0	0.34	0.0
Quality: extensive margin	0.0	0.83	0.0
Quality: intensive margin	0.03	1.0	0.03
Audience focus			
Submission	0.0	0.0	0.98
Quality: extensive margin	0.0	0.0	0.96
Quality: intensive margin	0.0	0.0	0.98

Notes: The table displays the p-values of the three types of null hypotheses indicated in each column, where F_0 and F_1 are the cumulative distribution functions (cdf's) of each semantic category (business content, sentiment, and audience focus) *conditional* on each of the entrepreneur-based performance measures (submission, extensive quality, and intensive quality) being below or above their median value, respectively. In the case of the submission outcome (a dichotomous variable), F_0 and F_1 are interpreted to be the cdf's conditional on "no submission" or "submission," respectively.

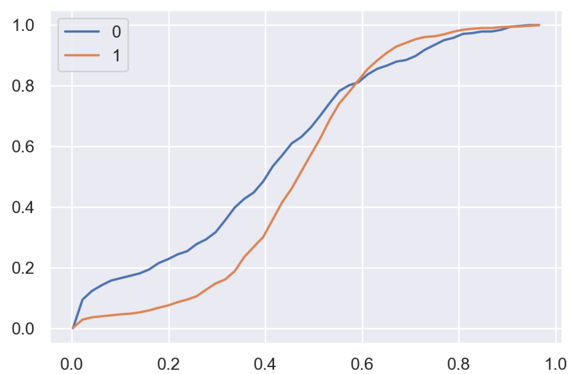

(a) Avge. business focus of entrepreneur cond. on submission

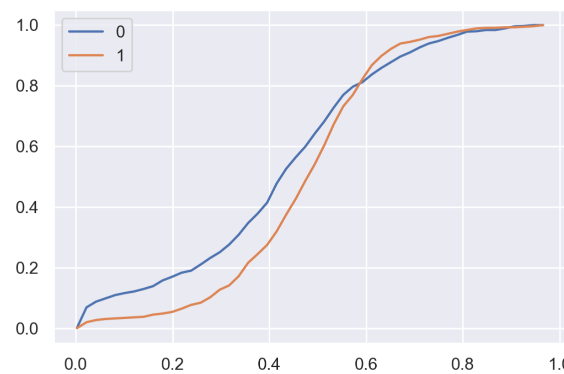

(b) Avge. business focus of entrepreneur cond. on extensive quality

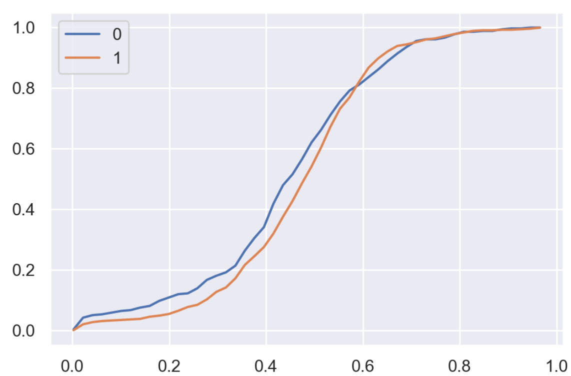

(c) Avge. business focus of entrepreneur cond.on intensive quality

Figure 16: The diagrams compare the cdf's defined over the business content of messages, aggregated at the level of entrepreneurs and conditional on their performance. On performance, we consider our three leading measures: submission (panel (a)), extensive quality (panel (b)) and intensive quality (panel (c)). In each of these cases, high performance is signaled by the indicator 1 and is interpreted as a value above the median; instead, lower performance is labeled by 0 and identified with a value below the median. For the dichotomous submission variable, 0 and 1 identify the cdf's conditional on "no submission" or "submission," respectively.

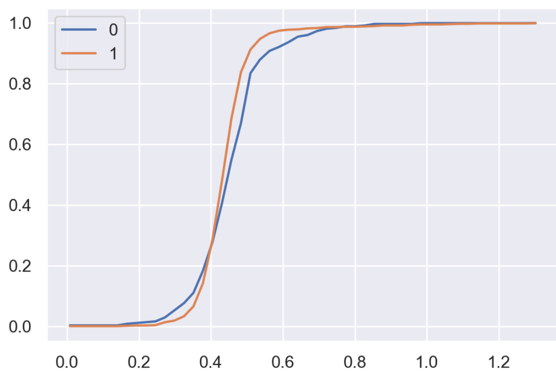

(a) Avge. sentiment of entrepreneur cond. on submission

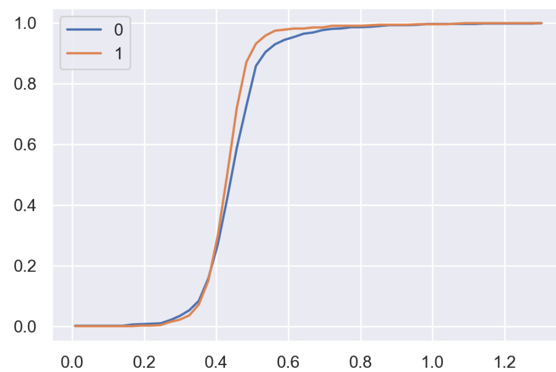

(b) Avge. sentiment of entrepreneur cond. on extensive quality

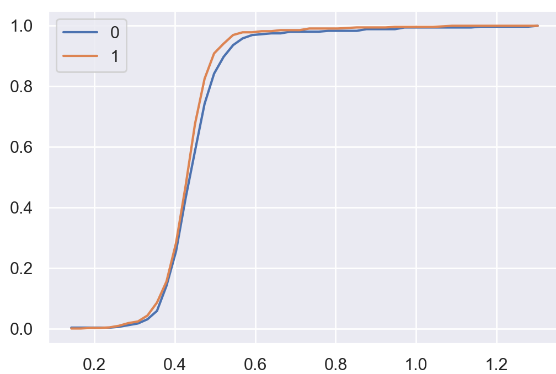

(c) Avge. sentiment of entrepreneur cond. on extensive quality

Figure 17: The diagrams here are a counterpart of those in Figure 16 with the cdf's defined over the sentiment of entrepreneurs' messages.

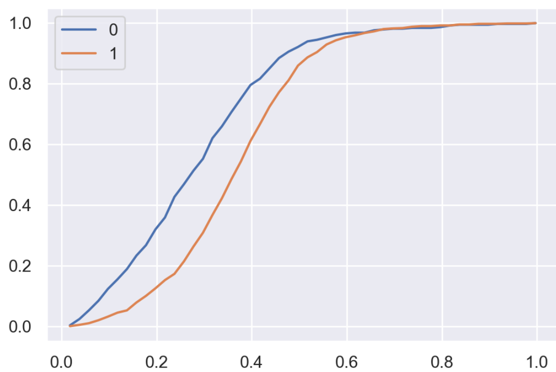

(a) Avge. audience focus of entrepreneur cond. on submission

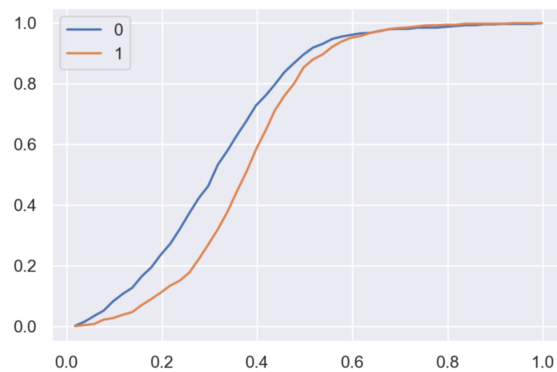

(b) Avge. audience focus of entrepreneur cond. on extensive quality

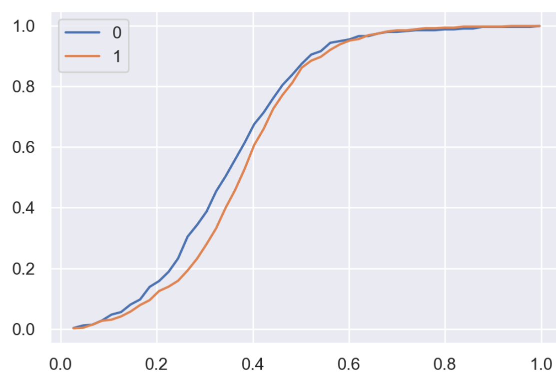

(c) Avge. audience focus of entrepreneur cond. on extensive quality

Figure 18: The diagrams here are a counterpart of those in Figure 16 and 17 with the cdf's defined over the audience focus of entrepreneurs' messages.

(c) Semantics-characteristic relationships

Finally, we explore how different exogenous characteristics of the entrepreneurs (baseline information or assigned treatment) affect their pattern of communication. Again we focus our analysis on our three leading semantic categories, and for each of them investigate whether certain baseline characteristics are associated to them in a statistical (distributional) sense. The results derived from KS tests applied to a set of selected baseline characteristics, are spelled out in Table 19. In analogy to former notation, F_0 and F_1 stand, respectively, for the cdf's conditional on low and high values of the characteristic where "low" and "high" are defined relative to the median if the characteristic is not dichotomous. If the characteristic is dichotomous, the subindices 1 and 0 are indicators of the characteristic being present or absent, respectively.

Table 19: **Semantic relationship to indiv. characteristics: p-values on distributional hypotheses**

Null hypotheses	$F_0 = F_1$	$F_0 \leq F_1$	$F_1 \leq F_0$
Business content			
Age	0.03	0.02	0.97
Female	0.0	0.99	0.0
Has business	0.62	0.42	0.37
Baseline quality	0.0	0.0	0.96
M0 complete	0.0	0.0	0.76
Virtual-across	0.05	0.86	0.08
Sentiment			
Age	0.09	1.0	0.054
Female	0.08	0.08	0.99
Has business	0.53	0.55	0.62
Baseline quality	0.53	0.73	0.53
M0 complete	0.0	0.39	0.01
Virtual-across	0.75	0.72	0.60
Audience focus			
Age	0.68	0.63	0.44
Female	0.67	0.83	0.38
Has business	0.73	0.56	0.77
Baseline quality	0.13	0.11	0.98
M0 complete	0.0	0.0	1.0
Virtual-across	0.97	0.81	0.75

Notes: The table displays the p-values of the three types of null hypotheses indicated in each column, where F_0 and F_1 are the cumulative distribution functions (cdf's) of each semantic category (business content, sentiment, and audience focus) *conditional* on each of the individual (entrepreneur) characteristic being below or above their median value, respectively. In the case of dichotomous characteristics, F_0 and F_1 are interpreted to be the cdf's conditional on the characteristic mentioned being absent or present, respectively.

Here we find it more useful to frame the comparisons in a way somewhat different from before: rather than grouping the results by category, we do it by characteristic. One of the remarkable features that arises from the results is that the treatment arm (vitual-within *vs* virtual across) does not have any significant effect on how the three semantic categories are used by the entrepreneurs. This follows from the corresponding

p-values in Table 19 and is intuitively from Figure 19, where we show the cdf's for each of the three semantic categories and treatment arm characteristic.

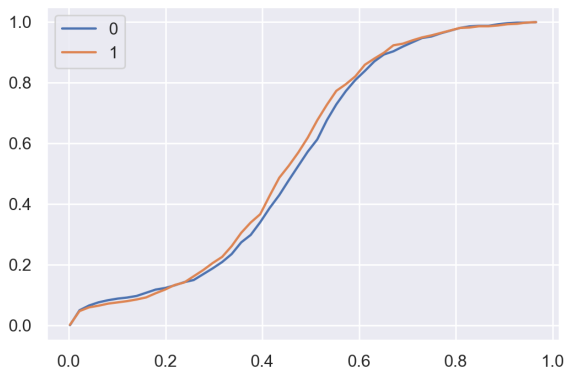

(a) Avge. business focus of entrepreneur cond. on v-across

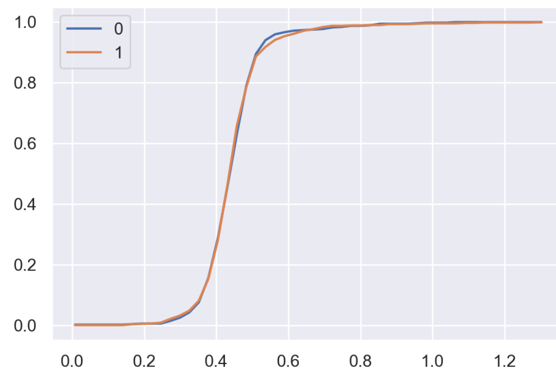

(b) Avge. sentiment of entrepreneur cond. on v-across

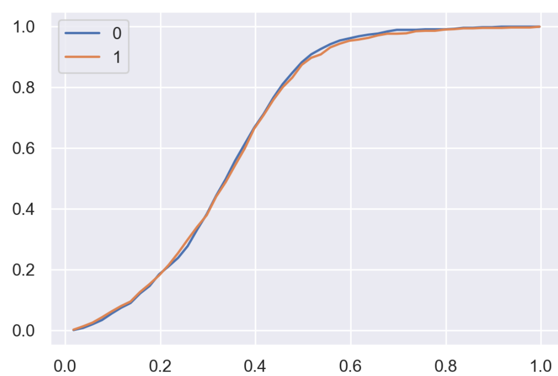

(c) Avge. audience focus of entrepreneur cond. on v-across

Figure 19: The diagrams compare the cdf's defined over our three leading semantic categories – business content (panel (a)), sentiment (panel (b)), and audience focus (panel (c)), aggregated at the entrepreneur level and conditioned on the treatment arm. The indicator 0 stands for the virtual-within arm and the indicator 1 for the virtual-across arm.

In contrast, an individual characteristic that does have a clear effect on the semantic categories used by entrepreneurs is whether they completed Milestone 0 on time (recall Subsection 2.5). In Table 19 we find that those entrepreneurs who did so wrote messages with significantly higher business content and private audience than those that did not, the latter showing instead significantly higher sentiment. (Note, however, that the latter conclusion requires a significance level that is not too high.) Again we illustrate these conclusions by showing in Figure 20 the corresponding cfds.

Finally, we highlight the only other cases where significant effects are present. They pertain to the business category and its dependence on age, gender, and baseline quality. Specifically, the p-values in

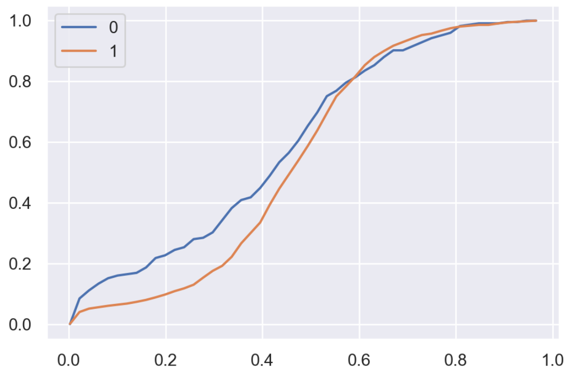

(a) Avge. business focus of entrepreneur cond. on M0 complete

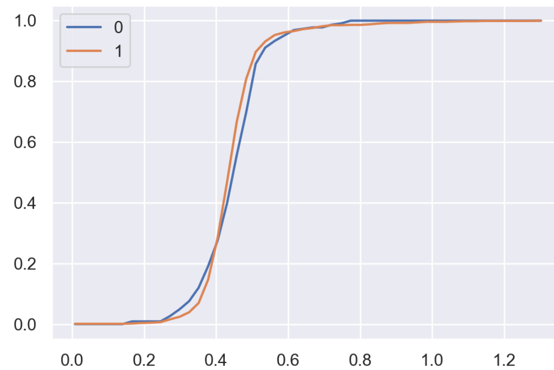

(b) Avge. sentiment of entrepreneur cond. on M0 complete

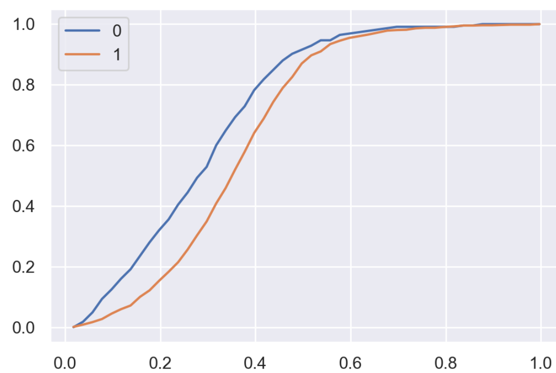

(c) Avge. audience focus of entrepreneur cond. on M0 complete

Figure 20: The diagrams compare the cdf's defined over our three leading semantic categories – business content (panel (a)), sentiment (panel (b)), and audience focus (panel (c)), aggregated at the entrepreneur level and conditioned on whether the entrepreneur completed Milestone 0 on time. The indicator 1 stands for on-time completion, while the indicator 0 corresponds to the opposite case.

Table 19 indicate that older people and those with a high baseline quality wrote messages with higher business content, while females used messages with a lower such content.[29] These conclusions are graphically illustrated in Figure 21.

[29]If the significance level $\alpha \geq 0.1$, then we may accept the hypothesis that females write instead messages with higher sentiment. Instead, gender does not have any clear impact on audience focus at reasonably moderate significance levels.

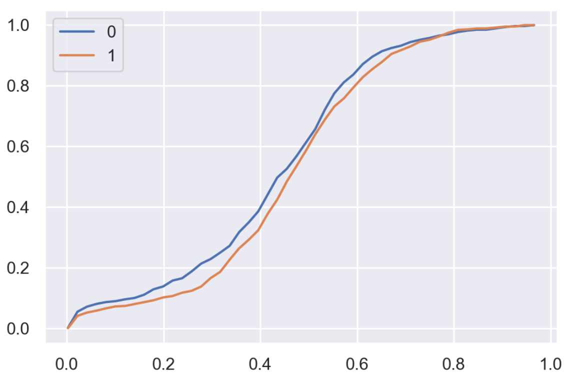

(a) Avge. business focus of entrepreneur cond. on age

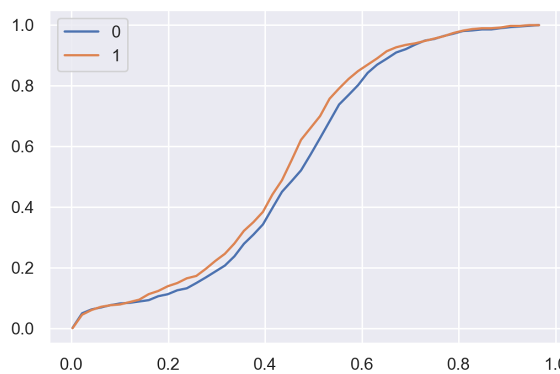

(b) Avge. business focus of entrepreneur cond. on gender

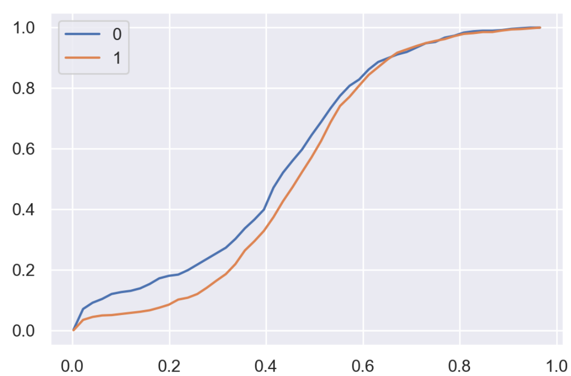

(c) Avge. business focus of entrepreneur cond. on baseline quality

Figure 21: The diagrams compare the cdf's defined over the business content of messages, aggregated at the entrepreneur level, and conditioned on age (panel (a)), gender (panel (b)), and baseline quality (panel (c)). For the age and baseline quality the indicator 1 identifies the set of values above the median, while the indicator 0 corresponds to those below the median. For gender, the indicator 1 signals females and 0 males.

5.2.3 Semantic indicators

We close our analysis of the communication unfolding in our experiment with a conceptually simple, and mostly illustrative, exercise. The primary objective of it is simply to provide a somewhat informal account of the main semantic indicators (terms and phrases) associated to the features (categories, baseline characteristics, and performance) studied in Subsection 5.2.2. To this end, we use the methodology generally known as randomized logistic regression – or stability selection, as called by Meinshausen and Buehlmann 2009. This is a bootstrap-based approach that repeatedly fits a logistic regression[30] with LASSO penalization on random subsets of the data.

More concretely, in our case we use as indicators words and two-word phrases that occur at least 10 times in our data, and use the target feature (e.g., business content) as dependent variable. The penalty term of the LASSO regressions encourages sparse models, i.e., coefficient vectors with many zeroes. Thus indicators that receive positive non-zero coefficient in a large fraction of all fitted specifications can be considered robust and give us a good understanding of the corresponding category. In particular, if one applies this procedure to a binary target category, the importance of individual terms with respect to each category can be assessed by inspecting the magnitude of the coefficients corresponding to each term.

Such an approach, however, presupposes that one has already selected the independent variables. In the case of text, however, the vocabulary size – and hence the number of independent variables – can be massive yet their individual occurrences sparse (for example, some words occur only once even in large corpora). It becomes impossible, therefore, to know a priori which variables to select. To address the problem, there are several methods that could be potentially used to select a subset of informative variables:

(1) Variables with non-significant Wald statistics can be removed. This, however, requires a large data set and not even then may lead to the identification of the best fitting model.

(2) One could conceivably explore all possible models on the entire data set and pick the one with the best fit, using some measure of fitness. Yet, due to the large number of possible models, this is, for a computational viewpoint, prohibitively extensive and in any case not robust to overfitting.

(3) The model can be estimated on the entire data set using an l_1 penalty term, the problem then being that this method is highly reliant on data composition.

In view of the somewhat unsatisfactory features induced by each of the three above methods, here we shall pursue the following *enrichment* of one proposed in (3):

(4) We estimate the model on the entire data set with an l_1 penalty term but, in order to address the indicated sensitivity to data composition, we collect *repeated random subsets of the data*, fitting a l_1 LASSO model to each realization and aggregating the correspondingly estimated coefficient vectors. In the end, those variables that frequently receive a high coefficient are selected.

[30] A logistic regression is a generalized linear regression model, where the log odds of the response probability is a linear function of the independent variables. The coefficients of these variables are estimated via maximum likelihood and then interpreted accordingly.

Intuitively, therefore, what we do is to explore a range of randomly generated conditions and pick the variables that are typically informative across the range of sample data sets induced. In practice, we fit from 200 to 1,000 independent logistic-regression models on random subsets of the data, each sampled with replacement. For each of these models, we use a different weight for the l_1 regularization parameter that controls how aggressively the estimation tends to drive regression coefficients down to 0. Finally, the regressors (indicators) that have a score larger than 0.5 (i.e., they received a non-zero coefficient weight in over 50% of the fitted models) are declared to be informative.

The methodology just outlined has been used to associate indicators (characteristic terms) for two different kinds of variables: *category labels* and *individual characteristics.* We summarize our findings for each of these cases in turn.

(a) Category indicators

- ***Business content***:
  - *High*: For messages high on business content, the top-10 most indicative terms are: *product, project, business, market, idea, need, customer, service, plan*, and *cost.*
  - *Low*: In contrast, for those messages that fare low on business content, the set of most significant terms includes: greetings (*hello, hi*), references to the token system (*token, leaderboard, status, giveusername*) and direct messages (*thanks, username, ok, great idea*).

- ***Sentiment***:

  Recall that sentiment is measured on a 5-point Likert scale (from -2 to +2). However, for the overwhelming majority of cases, the score is either 0 (neutral sentiment) or 1 (encouragement). We focus, therefore, on these two cases.
  - *Neutral*: Interestingly, many of the top-10 most indicative terms for neutral sentiment show a strong connection/overlap with the business category: *business, status, channel, market, help, product, start, need.*
  - *Encouragement*: Instead, for positive sentiment, we find terms associated to social interaction (*welcome, thank, username*), and positive judgment (*nice, good, great, awesome, idea, wish, appreciate*).

- ***Audience***:
  - *Individual scope*: Focus on a particular individual is typically reflected by the the use of *usernames* and other addressing terms (e.g. *bro*, for "brother"), as well as by references to the token system (*token, giveusername, deserve*), and by discussion of project-related terms (*idea, business, project*) and interaction verbs (*tell, mean*)
  - *General scope*: In contrast, messages targeting a general audience display a heavy use of community or group-wide terms (*leaderboard, status*), generic interaction terms (*hello, thanks, ok, welcome, yes*), and mentions of origin (e.g. *Nigeria* or *Africa*).

(b) Characteristic indicators

We find that the term-based regression analysis can also be used to single out specific terms preferentially associated to different baseline characteristics, endogenous behavioral features, or treatment conditions. Rather than providing an exhaustive account of these matters, we close this subsection with two mere illustrations.

First, it is well-documented that language embodies gender differences – c.f. Coates (2015). In our case, these gender-based differences arise both in language use as well as in topics. Concerning the first aspect, one of the main distinctive features is the use of different addressing terms: men use terms like *bro* or *boss*, whereas women use the *names* of the persons being addressed, and the terms *dear* or *girl*. Topically, on the other hand, when women talk about business they use more words pertaining to education and people (*think, child, like, school, work, know, woman, food*), whereas men focus more on products and services (*customer, product, service, business, plan, challenge, service channel, model*).

Second, we also detect a difference in the way entrepreneurs communicate under the different virtual-treatment conditions. Entrepreneurs in the virtual-across scenario refer much more to their own countries, as indicated by the prevalence of terms such as *Zimbabwe, Rwanda, Cameroon, Botswana*, or simply *country*. Instead, unsurprisingly, a common feature found in the virtual-within scenario is the prevalence of individual words from local native languages, such as *Swahili* or *Igbo*.

5.3 Semantically weighted networks

In this section, we integrate the semantic analysis and the network information to attempt an understanding of how different types of communication may operate through different subnetworks and what are the network implications of different types of communication. For the sake of focus, we restrict to the two dimensions/categories that seem most pertinent in this respect: business content and sentiment.

On the basis of the corresponding message (and hence link) weights associated to each category – as obtained from the annotation procedure described in Section 5.1) – we construct separate business-weighted and sentiment-weighted networks. The procedure is just as the one used in Subsection 4.1, except that in the present case the weights computed in each case are those of the respective category. In this manner, a "business(-weighted) network" is constructed where the communication accounted for is business-oriented, as well as another "sentiment(-weighted) network" where the communication flows are assessed in terms of their sentiment content.

An analysis parallel to that conducted in Section 4.1 for the benchmark setup leads to the estimation of semantically-weighted peer effects that are displayed in Tables 20 and 21. These are the counterpart of the peer effects that were reported in Table 10 where, in contrast, every sort of content was attributed the same weight.[31]

Tables 20 and 21 show that, either through business-weighted or sentiment-weighted links alike, peers can

[31] The semantically-weighted peer effects under bidirectional influence can be found in Tables A13-A14 in the Appendix. As it happens for the original scenario with uniform communication weight, peer effects under semantic weighting are robust to the consideration of bidirectional links.

Table 20: **Network peer effects, business network**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Control for ***all*** baseline info.	.521***	.010	-.009
	(.176)	(.137)	(.288)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Control for ***all*** baseline info.	.302**	.059	.598***
	(.132)	(.137)	(.120)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Control for ***all*** baseline info.	-.037	.081	.557***
	(.120)	(.153)	(.144)
Number of entrepreneurs: 779			

Notes: The table estimates the network-based peer effects channeled through the business-projected network on three different outcomes: the *submission decision* (see Table 2), the extensive effect on business quality (as measured in Table A2), and business quality measured in the *intensive margin* (see Table 4) for the business-projected network. The network-construction procedure is as described in Subsection 4.1, with each link weighted by its corresponding business-orientation index. As in the counterpart Table 10, no upper bound on the communication lag is contemplated (i.e. $\tau = \infty$). Standard errors are clustered at the group level, which gives rise to 44 clusters. The number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

Table 21: **Network peer effects, sentiment network**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Control for ***all*** baseline info.	.530***	.016	.099
	(.172)	(.141)	(.295)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Control for ***all*** baseline info.	.322**	.057	.503***
	(.132)	(.131)	(.146)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Control for ***all*** baseline info.	-.040	.087	.445***
	(.119)	(.150)	(.147)
Number of entrepreneurs: 779			

Notes: The table estimates the network-based peer effects channeled through the sentiment-projected network on three different outcomes: the *submission decision* (see Table 2), the extensive effect on business quality (as measured in Table A2), and business quality measured in the *intensive margin* (see Table 4) for the business-projected network.The network-construction procedure is as described in Subsection 4.1, with each link weighted by its corresponding sentiment-orientation index. As in the counterpart Table 10, no upper bound on the communication lag is contemplated (i.e. $\tau = \infty$). Standard errors are clustered at the group level, which gives rise to 44 clusters. The number of stars (*, **, ***) codes for statistical significance at ($p < 0.1$, $p < 0.05$, and $p < 0.01$), respectively.

exert significant influence. This seems a remarkable feature of the data, reinforced by the observation that, in fact, the pattern and magnitude of the peer effects channeled through business- or sentiment-weighted links are very similar to those obtained for our benchmark network, (c.f. Table 10). We find, therefore, that peer effects are very similar across the three cases.

At first glance, the above observations appear to be puzzling. For, in line with the conclusions obtained in Subsection 5.2.2 on the negative correlation across the business and sentiment categories at the message level, we also find negative correlation ***at the link level*** between the business- and sentiment-weighted links – the Pearson correlation is -0.44. In principle, this might suggest that, in effect, the business- and sentiment-weighted networks might be quite different (and so would be, consequently, the induced peer effects). This, however, is ***not*** the case, as Table 22 shows.

Table 22: **bilateral interaction flows: benchmark, business & sentiment networks.**

	benchmark	business	sentiment
benchmark	1		
business	0.820780	1	
sentiment	0.950685	0.698501	1

Notes: The table provides, for the benchmark, business, and sentiment networks, the correlation among the (directed) interaction flows from any entrepreneur i to some other j, obtained as the weight aggregation of all links $i \to j$.

Table 22 indicates that, at the level of *aggregate* communication flows *across peers*, the normalized (relative) link weights obtained by considering either business-oriented messages or sentiment-laden ones (as well as, for that matter, those ignoring semantic content) are very highly correlated. This contrasts with the aforementioned negative correlation between the two categories at the *message* and *link* levels. A natural interpretation of this state of affairs is in line with the point made in Subsection 5.2.1: even though business-oriented messages are often weak in sentiment content and *vice versa*, individuals tend to use a quite balanced set of the two types in the overall ***communication*** they establish with any ***given peer***. In the present case, however, the conclusion is substantially stronger in that such a semantic balance applies bilaterally across pairs of peers and for the communication actually flowing between them – i.e. not just to the messages generally sent by an entrepreneur, which in general may not be received – recall (LF) in Section 4.1. This suggests, in other words, that peer communication typically is multidimensional, either consciously (i.e. purposefully) or not.

6 Summary, conclusions, and ensuing research

In this document we have presented and discussed the results of a large experiment conducted in the African continent. Its objective has been to shed light on how peer networks, operating under alternative conditions, can help promote entrepreneurship in developing countries. To this end, we conducted an RCT with three treatment arms that differ in how peer interaction is conducted: face-to-face; "virtually within" (i.e. in groups of the same nationality); and "virtually across" (in nationally heterogeneous groups). In addition to such peer interaction, which was enjoyed only by treated individuals, the whole population – control and treatment alike – followed an online course for the duration of the experiment (two and a half months).

We have studied the problem from three different perspectives: treatment impact, network peer effects, and semantic NLP analysis. The main conclusions of this analysis can be summarized as follows.

First, we focus on the treatment effect – that is, the question of whether access to peer networking has an effect on entrepreneurship. The answer, of course, must depend on how one assesses "entrepreneurship." Three performance criteria are used: the submission of a business proposal; the quality of the submitted proposal (as graded by a panel of experts), which we call intensive quality; a combination of the former two, which we call extensive quality. We find that virtual-within interaction yields a positive and significant treatment effect in all three respects, while face-to-face only on submission, and virtual-across interaction in none. This is the first instance in which we identify a non-monotonic effect of interaction diversity (here spanning the geographical dimension) on entrepreneurial performance.

Second, we turn to the study of how peer effects operate through the network of participating entrepreneurs. This requires some preliminary steps, such as (a) suitably "extracting" the entrepreneur network from the interaction data, and (b) addressing some delicate identification issues. Our analysis then identifies positive and significant peer effects in two cases: on submission, under virtual-within interaction in large countries; on intensive quality, under virtual-across interaction in small countries. This adds an interesting complementary perspective on the former treatment analysis, highlighting the way in which the network operates and eventually succeeds, or fails, in producing a positive impact on individual and overall performance.

Third, in order to better understand how network and treatment effects interplay, we look into the "black box" of peer communication, as it unfolds over time throughout the experiment. To this end, we have conducted an extensive semantic (NLP) analysis of the vast amount of messages exchanged by peers in the virtual-treatment arms. Preliminary results unveil interesting patterns on how the business focus, sentiment, and target audience of messages interplay. They also provide a quite clear-cut understanding on how those semantic categories relate to exogenous characteristics (e.g. gender, age, or experience) and endogenous variables (entrepreneurial performance). This NLP-based approach has also been applied to the network analysis, by "projecting" the peer networks onto different semantic categories. None of these results turn out to be highly surprising. However, the fact that they are sharp and intuitive suggests the potential of a more thorough semantic research of peer interaction, a task we plan to carry out in the near future.

More generally, there are three further avenues of future research that seem important to us at this point. First, we need to develop a theoretical framework that helps better understand the wide range of conclusions obtained. Second, we want to build upon the insights derived from our experiment to enrich the design of follow-up studies that expand on a number of dimensions, e.g. the type and flexibility of interaction allowed (verbal communication, joint projects, etc.) or the range of feasible interventions (on incentives, norms, etc.). Finally, we want to test the external validity of the present African-based RCT in other environments, geographically or otherwise defined.

We believe that any significant advance along the lines proposed may play an important role in the design of distributed systems of innovation – for example, in suggesting what ranges of diversity should be spanned because they are fruitful and what other are to be avoided for being detrimental. Such an enhanced understanding of the problem should also be able to provide "food-for-design" advice, highlighting what incentives, protocols, trust-supporting and coordination mechanisms are most effective. In turn, it should benefit practitioners as well in setting up cost-effective and *scaleable* policies for the promotion of entrepreneurship, and hence economic development.

References

[1] Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2016): "Networks, shocks, and systemic Risk," in Bramoullé, Y., A. Galeotti, and B. Rogers, *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[2] Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitaliy (2017): "Tackling youth unemployment: evidence from a labor market experiment in Uganda," *Economic Organisation and Public Policy Discussion Papers* eopp64, Suntory and Toyota International Centres for Economics and Related Disciplines, London School of Economics.

[3] Attanasio, G., A. Kugler, and C. Meghir (2011): "Subsidizing vocational training for disadvantage youth in Colombia: evidence from a random trial," *American Economic Journal: Applied Microeconomics* **3**, 188-220.

[4] Bala, V. and S. Goyal (1998): "Learning from neighbors," *Review of Economic Studies* **65**, 595-621.

[5] Beaman, L. (2016): "Social networks and the labor market," in Bramoullé, Y., A. Galeotti, and B. Rogers, *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[6] Bloch, F. , M. Jackson, and P. Tebaldi (2017): "Centrality Measures in Networks," available at SSRN: https://ssrn.com/abstract=2749124 or http://dx.doi.org/10.2139/ssrn.2749124.

[7] Breiman, Leo (2001): "Random Forests," M*achine Learning* **45**, 5-32.

[8] Breza, E. (2016): "Field experiments, social networks, and development," in Bramoullé, Y., A. Galeotti, and B. Rogers, *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[9] Bramoullé, Y., H. Djebbari, B. Fortin (2009): "Identification of peer effects through social networks," *Journal of Econometrics* **150**, 41-55.

[10] Bramoullé, Y., A. Galeotti, and B. Rogers (2016): *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[11] Burt, R. S. (1992): *Structural Holes: The Social Structure of Competition*, Cambridge: Harvard University Press.

[12] Burt, R. S. (2004): "Structural holes and good ideas," *American Journal of Sociology*, **110**, 349-399.

[13] Cabrales, A., D. Gale, and P. Gottardi (2016): "Financial contagion in networks," in Bramoullé, Y., A. Galeotti, and B. Rogers, *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[14] Cai, J. and A. Szeidl (2017): "Interfirm Relationships and Business Performance," forthcoming in *Quarterly Journal of Economics*.

[15] Carrell, S., Sacerdote, B., and West, J. (2013): "From natural variation to optimal policy? the Importance of endogenous peer group formation," *Econometrica* **81**, 855-882.

[16] Coates, J. (2015): "Women, men and language: A sociolinguistic account of gender differences in language," London: Routledge.

[17] Dahl, M. S. and C. . R. Pedersen, (2005): "Social networks in the R & D process: The case of the wireless communication industry around Aalborg, Denmark," *Journal of Engineering and Technological Management* **22**, 75-92.

[18] Davidsson, P. and B. Honig (2000): "The role of social and human capital among nascent entrepreneurs," *Academy of Management Proceedings*, B1-B6.

[19] Dayan, M., M. Ozer, and H. Almazrouei (2017): "The role of functional and demographic diversity on new product creativity and the moderating impact of project uncertainty," Industrial Marketing Management **61**, 144-154.

[20] De Giorgi, G., M. Pellizzari, and S. Redaelli (2010). "Identification of social interactions through partially overlapping peer groups," *American Economic Journal: Applied Econometrics* **2**, 241-275.

[21] DeMarzo, P.M., D. Vayanos and J. Zwiebel (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions", *The Quarterly Journal of Economics*, **118**, 909-968.

[22] Dessein W. and Andrea Prat (2016): "Attention in organizations," in Bramoullé, Y., A. Galeotti, and B. Rogers, *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

[23] Fafchamps, M. and S. Qinn (2016): Networks and Manufacturing Firms in Africa: Results from a Randomised Experiment, *World Bank Economic Review*, lhw057, https://doi.org/10.1093/wber/lhw057.

[24] Fafchamps, M. and C. Woodruff (2017): "Identifying Gazelles: Expert Panels vs. Surveys as a Means to Identify Firms with Rapid Growth Potential", in collaboration with Chris Woodruff, World Bank Economic Review **31**, 670-686.

[25] Franke J. and T. Ozturk (2015): "Conflict networks," *Journal of Public Economics* **126**, 104-113.

[26] Edquist, C., M.-L. Eriksson, H. Sjögren, (2000): "Collaboration in product innovation in the East Gothia regional system of innovation," *Enterprise and Innovation Management Studies* **1**, 37-56.

[27] Golub, B. and M. O. Jackson (2010): "Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds", *American Economic Journal: Microeconomics* **2**, 112-149.

[28] Gonzalez-Uribe, J. and M. Leatherbee (2018): "The effects of business accelerators on venture performance: evidence from Start-Up Chile," *Review of Financial Studies* **31**, 1566-1603.

[29] Granovetter, M. (1973): "The strength of weak ties," *American Journal of Sociology* **78**, 1360-1380.

[30] Hacamo I. and K. Kleiner (2018): "Belief in yourself: Do peers promote or discourage entrepreneurship?," working paper, Kelley School of Business, Indiana University,

[31] Hagedoorn, J. (2002): "Inter-firm R&D partnerships: an overview of major trends and patterns since 1960," Research Policy 31, 477-492.

[32] Heathcote A., S. Brown, E.J. Wagenmakers, A. Eidels (2010): "Distribution-free tests of stochastic dominance for small samples," *Journal of Mathematical Psychology* **54**, 454-463

[33] Heijnen, P. and A. R. Soetevent (2018): Price competition on graphs, *Journal of Economic Behavior and Organization* **146**, 161-179.

[34] Hovy, D., T. Berg-Kirkpatrick, A. Vaswani, and E. Hovy. (2013): "Learning whom to trust with MACE," in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1120-1130.

[35] Jackson, M. (2008): *Social and economic networks*, Princeton: Princeton University Press.

[36] Jackson, M., T. Rodriguez-Barraquer and Xu Tan (2012): "Social capital and social quilts: network patterns of favor exchange," *American Economic Review* **102**, 1857-97.

[37] Kolmogorov, A. (1933): "Sulla determinazione empirica di une legge di distribuzione," *Giornale dell'Istituto Italiano degli Attuari* **4**, 83-91.

[38] König, M.D, D. Rohner, M. Thoenig and F. Zilibotti (2017): "Networks and Conflicts: Theory and Evidence from Rebel Groups in Africa," *Econometrica* **85**,10931132.

[39] J. Lerner and U. Malmendier (2013): "With a little help from my (random) friends: Success and failure in post-business school entrepreneurship," *Review of Financial Studies* **26**, 2411-2452.

[40] Letaifa, S. B. and Y. Rabeau (2013): "Too close to collaborate? How geographic proximity could impede entrepreneurship and innovation," *Journal of Business Research* **66**, 207-2078.

[41] Lippert, S. and G. Spagnolo (2011): "Networks of relations, word-of-mouth communication, and social capital," *Games and Economic Behavior* **72**, 202-217.

[42] Lungeanu A. and N. S. Contractor (2015): "The effects of diversity and network ties on Innovations: The emergence of a new scientific field," *American Behavioral Scientist* **59**, 548-564.

[43] Manski, C. F. (1993): "Identification of endogenous social effects: The reflection problem," *The Review of Economic Studies* **60**, 531-542.

[44] McKenzie, D. and D. Sansone (2017): Man vs. machine in predicting successful entrepreneurs: evidence from a business plan competition in Nigeria," Wolrld Bank Policy Group, Policy Research Working Paper 8271.

[45] McKenzie, D. and C. Woodruff (2014): "What are we learning from business training and entrepreneurship evaluations around the developing world?," *The World Bank Research Observer* **29**, 48-82.

[46] Meinshausen, N. and P. Bühlmann (2015): "Stability selection," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**, 417-473.

[47] Karlan, D., M. Mobius, T. Rosenblat, and A. Szeidl (2009): "Trust and Social Collateral," *Quarterly Journal of Economics* **124**, 1307-1331.

[48] Nanda R. and J. B. Sorensen (2010): "Workplace peers and entrepreneurship," *Management Science* **56**, 1116-1126.

[49] Parzen, E. (1962): "On estimation of a probability density function and mode," *The Annals of Mathematical Statistics* **33**, 1065-1076.

[50] Powell, W.W., K.W. Koput, L. Smith-Doerr (1996): "Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology," *Administrative Science Quarterly* **41**, 116-45.

[51] Rosenblatt, M. (1956): "Remarks on some nonparametric estimates of a density function," *The Annals of Mathematical Statistics* **27**, 832-837.

[52] Sarker , S., M. Ahuja , S. Sarker, and S. Kirkeby (2011): "The role of communication and trust in global virtual teams: A social network perspective," *Journal of Management Information Systems* **28**, 273-310.

[53] Scott, D.W. (1992): *Multivariate Density Estimation: Theory, Practice, and Visualization*, New York, Chicester: John Wiley & Sons.

[54] Smirnov, N. V. (1939): "Sur les carts de la courbe de distribution empirique," *Matematiceskii Sbornik NS* **6**, 3-26.

[55] Tavassoli and Carbonara (2014): "The role of knowledge variety and intensity for regional innovation," *Small Business Economics* **43**, 493-509.

[56] Vega-Redondo, F. (2007): *Complex Social Networks*, Econometric Society Monograph Series, Cambridge: Cambridge University Press.

[57] Zhao, L. and J. Aram (1995): "Networking and growth of young technology-intensive ventures in China," *Journal of Business Venturing* **10**, 349-370.

Appendix: supplementary material

In this Appendix we include the tables that check the characteristic auxiliary tables that complement those included in the main text.

A.1 Balance checking

In Table A1-A3 below, we check the balance across individual baseline characteristics in the randomization described in Subsection 2.4. This balance check is separately performed for each of our three samples: The Uganda Sample (UgS), the Large-Country Sample (LCS) and the Small-Country Sample (SCS).

Table A1: **Balance check, Uganda sample**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Control		Face to Face		Virtual Within				
	Mean	S.D.	Mean	S.D.	Mean	S.D.	p-value (control = f2f)	p-value (control =vwithin)	p-value (within =control =f2f)
Observations	189		189		190				
p-value Joint F test*							0.31	0.62	
Panel A. Stratification Variables									
Female	0.35	0.48	0.35	0.48	0.35	0.48			
Has Business	0.63	0.48	0.63	0.48	0.63	0.49			
Course first milestone on time	0.51	0.50	0.51	0.50	0.52	0.50			
Panel B. Demographics									
Age	30.60	8.11	30.02	8.01	29.48	8.37	0.49	0.17	0.41
University Complete	0.86	0.35	0.92	0.27	0.89	0.31	0.05	0.35	0.14
Married	0.36	0.48	0.32	0.47	0.30	0.46	0.47	0.19	0.42
Panel C. Business Idea									
Has Idea about Existing Business	0.39	0.49	0.39	0.49	0.35	0.48	0.97	0.36	0.58
Has Idea about New Business	0.49	0.50	0.51	0.50	0.58	0.50	0.77	0.08	0.17
Has no Idea yet	0.12	0.32	0.10	0.30	0.07	0.26	0.59	0.15	0.31
Sector: Agriculture	0.23	0.42	0.28	0.45	0.26	0.44	0.24	0.43	0.48
Sector: Services	0.12	0.33	0.16	0.37	0.11	0.31	0.25	0.58	0.24
Sector: Technology	0.15	0.36	0.12	0.33	0.17	0.38	0.46	0.48	0.36
Sector: Manufacture	0.08	0.27	0.07	0.25	0.10	0.30	0.67	0.48	0.55
Sector: Social Entrepreneurship	0.20	0.40	0.16	0.37	0.15	0.36	0.36	0.17	0.39
Sector: Retail	0.03	0.18	0.04	0.20	0.05	0.21	0.57	0.44	0.72
Has written business plan	0.57	0.50	0.61	0.49	0.54	0.50	0.45	0.57	0.43
Participated in business competition	0.31	0.46	0.35	0.48	0.32	0.47	0.35	0.82	0.64
Has employees	0.39	0.49	0.38	0.49	0.39	0.49	0.74	0.85	0.87
Panel D. Financial Access									
Saves at a Bank	0.81	0.39	0.84	0.37	0.85	0.36	0.48	0.26	0.52
Got Bank Loan for business	0.12	0.32	0.15	0.36	0.08	0.27	0.35	0.22	0.10
Prefers equity debt to loans	0.57	0.50	0.50	0.50	0.60	0.49	0.17	0.59	0.14
Prefers loans to equity debt	0.14	0.35	0.17	0.38	0.10	0.30	0.39	0.25	0.13
Prefers either equity or loans	0.26	0.44	0.30	0.46	0.26	0.44	0.34	1.00	0.55
Panel E. Labor Market Outcomes									
Reservation Wage (in USD)	1510	2008	1199	1596	1333	1947	0.10	0.39	0.25
Years of work experience	4.87	3.30	4.77	3.24	4.53	3.22	0.76	0.30	0.58
Has a job	0.54	0.50	0.60	0.49	0.54	0.50	0.22	0.97	0.37
Panel F. Networks									
Number of People discuss business	4.79	3.25	4.39	3.33	4.84	3.43	0.22	0.90	0.33
Prefers to discuss with different sector	0.13	0.33	0.10	0.30	0.09	0.29	0.41	0.31	0.58
Prefers to discuss with different gender	0.15	0.36	0.15	0.36	0.13	0.33	0.99	0.45	0.68
Prefers to discuss with different country	0.21	0.41	0.16	0.37	0.19	0.39	0.23	0.67	0.47
Panel G. Personality Traits									
Risk Aversion (choice among 6 lotteries)	3.48	2.00	3.38	2.06	3.46	1.96	0.64	0.90	0.88
Trust Measure (0 to 10)	4.70	2.69	5.12	2.96	4.69	2.51	0.14	0.99	0.24
Position in your country: current (0 to 10)	4.79	1.81	4.57	1.58	4.97	1.71	0.24	0.31	0.07
Position in your country: expected (0 to 10)	7.65	1.59	7.39	1.55	7.48	1.56	0.14	0.29	0.31
Position in your country: desired (0 to 10)	8.85	1.68	9.00	1.52	8.95	1.66	0.41	0.63	0.71

Notes: The randomization was stratified on gender, having a business and submitting the first milestone of the course on time. The table uses variables from the online application form (March-May 2017). Columns 7-9: p-values for tests of equality of means obtained from a regression of each variable on treatment controlling for randomization strata with robust standard errors. For the orthogonality test we replaced missing values with zeros and included dummies for missing variables with missing values.

Table A2: **Balance check, large-country sample**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Control		Virtual Within		Virtual Across				
	Mean	S.D.	Mean	S.D.	Mean	S.D.	p-value (control =vwithin)	p-value (control =vacross)	p-value (within =control =across)
Observations	1111		1111		1111		0.98	0.98	
p-value Joint F test*									
Panel A. Stratification Variables									
Ghana	0.13	0.34	0.13	0.34	0.13	0.34			
Kenya	0.12	0.32	0.12	0.32	0.12	0.32			
Nigeria	0.61	0.49	0.61	0.49	0.61	0.49			
South Africa	0.09	0.28	0.09	0.28	0.09	0.28			
Tanzania	0.06	0.23	0.05	0.23	0.05	0.23			
Female	0.31	0.46	0.31	0.46	0.31	0.46			
Has Business	0.65	0.48	0.65	0.48	0.65	0.48			
Course first milestone on time	0.56	0.50	0.56	0.50	0.55	0.50			
Panel B. Demographics									
Age	31.38	7.74	31.26	7.47	31.26	7.76	0.69	0.70	0.90
Unversity Complete	0.82	0.39	0.81	0.39	0.81	0.39	0.78	0.60	0.88
Married	0.38	0.49	0.39	0.49	0.40	0.49	0.50	0.32	0.60
Panel C. Business Idea									
Has Idea about Existing Business	0.34	0.47	0.37	0.48	0.35	0.48	0.09	0.49	0.23
Has Idea about New Business	0.59	0.49	0.56	0.50	0.59	0.49	0.06	0.82	0.11
Has no Idea yet	0.07	0.26	0.07	0.26	0.06	0.24	0.76	0.37	0.42
Sector: Agriculture	0.25	0.44	0.24	0.43	0.26	0.44	0.64	0.62	0.62
Sector: Services	0.17	0.38	0.19	0.39	0.17	0.37	0.41	0.77	0.53
Sector: Technology	0.14	0.35	0.13	0.33	0.14	0.34	0.36	0.92	0.60
Sector: Manufacture	0.10	0.30	0.12	0.32	0.10	0.30	0.23	0.78	0.29
Sector: Social Entrepreneurship	0.09	0.29	0.09	0.29	0.10	0.30	0.83	0.65	0.80
Sector: Retail	0.05	0.21	0.05	0.22	0.05	0.22	0.51	0.51	0.74
Has written business plan	0.58	0.49	0.62	0.49	0.60	0.49	0.06	0.40	0.17
Participated in business competition	0.36	0.48	0.38	0.49	0.38	0.49	0.34	0.23	0.45
Has employees	0.36	0.48	0.34	0.47	0.39	0.49	0.25	0.15	0.04
Panel D. Financial Access									
Saves at a Bank	0.93	0.25	0.94	0.24	0.93	0.26	0.69	0.70	0.73
Got Bank Loan for business	0.09	0.29	0.08	0.27	0.09	0.28	0.24	0.49	0.51
Prefers equity debt to loans	0.44	0.50	0.45	0.50	0.43	0.50	0.56	0.77	0.67
Prefers loans to equity debt	0.16	0.36	0.16	0.37	0.16	0.37	0.71	0.82	0.93
Prefers either equity or loans	0.35	0.48	0.34	0.47	0.36	0.48	0.62	0.59	0.58
Panel E. Labor Market Outcomes									
Reservation Wage (in USD)	1664	2173	1685	2204	1642	2147	0.76	0.84	0.87
Years of work experience	5.31	3.24	5.23	3.19	5.25	3.25	0.58	0.69	0.85
Has a job	0.54	0.50	0.54	0.50	0.55	0.50	0.87	0.73	0.88
Panel F. Networks									
Number of People discuss business	4.66	3.41	4.53	3.33	4.55	3.31	0.36	0.44	0.62
Prefers to discuss with different sector	0.10	0.30	0.09	0.29	0.11	0.31	0.27	0.83	0.35
Prefers to discuss with different gender	0.12	0.32	0.11	0.31	0.11	0.32	0.61	0.81	0.88
Prefers to discuss with different country	0.17	0.38	0.16	0.37	0.14	0.35	0.58	0.05	0.13
Panel G. Personality Traits									
Risk Aversion (choice among 6 lotteries)	3.45	2.06	3.39	2.03	3.41	2.07	0.44	0.62	0.74
Trust Measure (0 to 10)	4.74	2.73	4.75	2.77	4.78	2.82	0.92	0.69	0.92
Position in your country: current (0 to 10)	4.85	1.59	4.81	1.60	4.89	1.65	0.47	0.60	0.49
Position in your country: expected (0 to 10)	7.94	1.57	7.91	1.53	7.96	1.56	0.67	0.77	0.84
Position in your country: desired (0 to 10)	9.09	1.55	9.05	1.55	9.08	1.58	0.63	0.92	0.87

Notes: Randomization was stratified on country, gender, having a business and submitting the first milestone of the course on time. The table uses variables from the online application form (March-May 2017). Columns 7-9: p-values for tests of equality of means obtained from a regression of each variable on treatment controlling for randomization strata with robust standard errors. For the orthogonality test we replaced missing values with zeros and included dummies for missing variables with missing values.

Table A3: **Balance check, small-country sample**

	(1)	(2)	(3)	(4)	(5)
	Control		Virtual Across		
	Mean	S.D.	Mean	S.D.	p-value (control =vacross)
Observations	529		528		
p-value Joint F test*					.36
Panel A. Stratification Variables					
Northern Africa	.06	.24	.06	.24	
Western Africa	.17	.37	.17	.37	
Eastern Africa	.57	.5	.57	.5	
Southern Africa	.03	.17	.03	.17	
Middle Africa	.18	.38	.18	.38	
Female	.31	.46	.32	.47	
Has Business	.58	.49	.58	.49	
Course first milestone on time	.56	.5	.57	.5	
Panel B. Demographics					
Age	30.15	7.9	30.65	7.99	.29
Unversity Complete	.78	.41	.8	.4	.42
Married	.35	.48	.4	.49	.06
Panel C. Business Idea					
Has Idea about Existing Business	.45	.5	.46	.5	.84
Has Idea about New Business	.47	.5	.45	.5	.58
Has no Idea yet	.08	.28	.09	.29	.54
Sector: Agriculture	.22	.42	.21	.41	.6
Sector: Services	.16	.37	.16	.37	.98
Sector: Technology	.12	.33	.13	.33	.8
Sector: Manufacture	.06	.24	.08	.26	.41
Sector: Social Entrepreneurship	.18	.38	.18	.38	.87
Sector: Retail	.03	.17	.05	.21	.14
Has written business plan	.56	.5	.56	.5	.85
Participated in business competition	.4	.49	.36	.48	.15
Has employees	.32	.47	.33	.47	.54
Panel D. Financial Access					
Saves at a Bank	.81	.39	.83	.37	.39
Got Bank Loan for business	.1	.3	.11	.31	.87
Prefers equity debt to loans	.4	.49	.4	.49	.99
Prefers loans to equity debt	.17	.37	.18	.38	.69
Prefers either equity or loans	.37	.48	.37	.48	.95
Panel E. Labor Market Outcomes					
Reservation Wage (in USD)	1509.43	1921.24	1597.59	2038.58	.44
Years of work experience	4.6	3.3	5.07	3.44	.02
Has a job	.52	.5	.56	.5	.17
Panel F. Networks					
Number of People discuss business	4.79	3.54	4.73	3.37	.8
Prefers to discuss with different sector	.13	.34	.12	.33	.62
Prefers to discuss with different gender	.17	.38	.13	.33	.02
Prefers to discuss with different country	.24	.43	.23	.42	.7
Panel G. Personality Traits					
Risk Aversion (choice among 6 lotteries)	3.61	2.1	3.45	2.00	.21
Trust Measure (0 to 10)	4.99	2.87	4.92	2.72	.69
Position in your country: current (0 to 10)	4.6	1.64	4.88	1.68	.01
Position in your country: expected (0 to 10)	7.68	1.57	7.81	1.48	.26
Position in your country: desired (0 to 10)	9.12	1.39	9.06	1.46	.46

Notes: Randomization was stratified on region, gender, having a business and submitting the first milestone of the course on time. The table uses values of the variables collected in the online application form (March-May 2017). Columns 4-6: p-values for tests of equality of means obtained from a regression of each variable on treatment controlling for randomization strata with robust standard errors. For the orthogonality test we replaced missing values with zeros and included dummies for missing variables with missing values.

A.2 The core treatment effects: full set of controls

Here we reproduce most of the regressions reported in Section 3 including as controls the full set of baseline characteristics obtained through the initial survey.

Table A4: **Replicating Table 3 (extensive margin), with controls**

	(1)	(2)
	Quality. Full Sample	Quality. M0 Sample
Panel A. Uganda Sample (UgS)		
Face-to-face treatment	0.336^{***}	0.374^{*}
	(0.112)	(0.208)
Virtual-within treatment	0.209^{**}	0.286
	(0.100)	(0.191)
Control Mean	0.88	1.47
p-value face-to-face = virtual within	0.17	0.61
Number of Entrepreneurs	568	291
Panel B. Large-Country Sample (LCS)		
Virtual-within treatment	0.118^{***}	0.247^{***}
	(0.046)	(0.076)
Virtual-across treatment	0.030	0.056
	(0.051)	(0.086)
Control Mean	1.03	1.58
p-value virtual-within = virtual-across	0.04	0.01
Number of Entrepreneurs	3,333	1,848
Panel C. Small-Country Sample (SCS)		
Virtual-across treatment	-0.074	0.003
	(0.073)	(0.112)
Control Mean	1.10	1.62
Number of Entrepreneurs	1,056	595

Notes: See Table 3. Regressions also control for all variables listed in the balance table and baseline quality (obtained from evaluations of a summary of proposals submitted before treatment assignment); we replace missing values in co-variates with zeros and include dummies for missing observations.

Table A5: **Replicating Table 4 (intensive Margin), with controls**

	(1) Intensive Margin. Full	(2) Intensive Margin M0 Sample
Panel A. Uganda Sample (UgS)		
Face-to-face treatment	0.381	0.198
	(0.251)	(0.244)
Virtual-within treatment	0.507***	0.576***
	(0.155)	(0.161)
Quality Control Mean if submitted	2.61	2.66
Quality Effect f2f	0.02	-0.04
Quality Effect VW	0.42	0.30
Number of entrepreneurs who submitted	221	179
Number of entrepreneurs	568	291
Panel B. Large-Country Sample (LCS)		
Virtual-within treatment	0.086	0.157**
	(0.072)	(0.070)
Virtual-across treatment	-0.040	-0.032
	(0.078)	(0.082)
Quality Control Mean if submitted	2.70	2.71
Quality Effect virtual within	0.05	0.18
Quality Effect virtual across	-0.02	0.01
Number of entrepreneurs who submitted	1,322	1,118
Number of entrepreneurs	3,333	1,848
Panel C. Small-Country Sample (SCS)		
Virtual-across treatment	0.026	0.081
	(0.111)	(0.116)
Quality Control Mean if submitted	2.65	2.65
Quality Effect virtual across	0.22	0.17
Number of entrepreneurs who submitted	409	351
Number of entrepreneurs	1,056	595

Notes: See Table 4. Regressions also control for all variables listed in the balance table and baseline quality (obtained from evaluations of a summary of proposals submitted before treatment assignment); we replace missing values in co-variates with zeros and include dummies for missing observations.

Table A6: **Replicating Table 5 (marginal-quality effects), with controls**

	(1) Score = 1	(2) Score = 2	(3) Score = 3	(4) Score = 4	(5) Score = 5
Face-to-face treatment	-0.040	-0.057	0.044	0.048	0.005
	(0.047)	(0.059)	(0.045)	(0.054)	(0.007)
Virtual-within treatment	-0.088**	-0.184***	0.063	0.177***	0.032**
	(0.036)	(0.043)	(0.041)	(0.044)	(0.014)
Number of Entrepreneurs	179	179	179	179	179
Virtual-within treatment	-0.035**	-0.023**	0.008*	0.034**	0.016**
	(0.016)	(0.010)	(0.005)	(0.015)	(0.007)
Virtual-across treatment	0.009	0.005	-0.003	-0.008	-0.003
	(0.020)	(0.011)	(0.007)	(0.017)	(0.007)
Number of Entrepreneurs	1,118	1,118	1,118	1,118	1,118
Virtual-across treatment	-0.017	-0.013	0.006	0.020	0.004
	(0.026)	(0.019)	(0.010)	(0.028)	(0.006)
Number of Entrepreneurs	351	351	351	351	351

Notes: See Table 5. Regressions also control for all variables listed in the balance table and baseline quality (obtained from evaluations of a summary of proposals submitted before treatment assignment); we replace missing values in co-variates with zeros and include dummies for missing observations.

Table A7: **Replicating Table 7 (second-stage outcomes), with controls**

	(1)	(2)	(3)	(4)
	Reach Stage2	Evaluated Stage2	Investor interest	Quality Stage2
Panel A. Uganda Sample (UgS)				
Face-to-face treatment	0.073*	0.077	0.048	-1.522
	(0.043)	(0.578)	(.)	(.)
Virtual-within treatment	0.043	-0.046	-0.098	-1.265
	(0.030)	(0.305)	(.)	(.)
Control Mean	0.06	0.75	0.22	0.53
p-value face-to-face = virtual within	0.48	0.81	.	.
Number of Entrepreneurs	568	60	37	37
Panel B. Large-Country Sample (LCS)				
Virtual-within treatment	0.024**	0.032	-0.096	-0.091
	(0.011)	(0.053)	(0.065)	(0.150)
Virtual-across treatment	-0.001	0.080	-0.067	-0.014
	(0.012)	(0.053)	(0.066)	(0.151)
Control Mean	0.12	0.77	0.28	-0.01
p-value virtual-within = virtual-across	0.01	0.33	0.63	0.62
Number of Entrepreneurs	3,333	418	332	332
Panel C. Small-Country Sample (SCS)				
Virtual-across treatment	-0.014	-0.032	-0.068	0.099
	(0.018)	(0.094)	(0.100)	(0.373)
Control Mean	0.13	0.64	0.20	0.04
Number of Entrepreneurs	1,057	130	85	85

Notes: See Table 7. Regressions also control for all variables listed in the balance table and baseline quality (obtained from evaluations of a summary of proposals submitted before treatment assignment); we replace missing values in co-variates with zeros and include dummies for missing observations.

A.3 Peer composition: robustness check

Here we conduct two robustness checks on the peer composition effects estimated in Subsection 3.3: a placebo test for the control group, and an analysis of the surprise component.

Table A8: **Effects for peer composition (I): placebo for control group**

	(1)	(2)	(3)	(4)	(5)	(6)
	Submission	Quality Ext.	Quality Int.	Submission	Quality Ext.	Quality Int.
Panel A. Sample who submitted M0 on time in Large Countries						
Own Baseline Quality	0.085^{***}	0.298^{***}	0.093	0.085^{***}	0.296^{***}	0.096
	(0.021)	(0.067)	(0.070)	(0.021)	(0.067)	(0.069)
Average Peer Quality	0.014	0.083	0.065			
	(0.074)	(0.240)	(0.228)			
Share Peers with Business				0.091	0.404	0.439
				(0.154)	(0.513)	(0.522)
Number of Entrepreneurs	617	617	360	617	617	360
Panel B. Full Sample in Large Countries						
Own Baseline Quality	0.086^{***}	0.299^{***}	0.094	0.085^{***}	0.297^{***}	0.096
	(0.021)	(0.067)	(0.070)	(0.021)	(0.067)	(0.070)
Average Peer Quality	0.051	0.208^{*}	0.152			
	(0.038)	(0.118)	(0.177)			
Share Peers with Business				0.146	0.274	-0.232
				(0.149)	(0.473)	(0.641)
Number of Entrepreneurs	1,111	1,111	423	1,111	1,111	423

Notes: The table uses data for entrepreneurs in the control group (no interaction) in the Large-Country Sample (LCS). Groups are randomly created using the same procedure as for the interaction groups. In all regressions we include strata fixed effects and control for baseline level of quality obtained from evaluations of a summary of proposals submitted before treatment assignment; we replace missing values in baseline quality with zeros and include a dummy for missing observations. Standard errors are clustered at the group level (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively. The outcome in columns (1) and (4) is submission of the proposal, for columns (2) and (4) quality evaluation replacing missing values with zeros, for columns (2) and (4) quality evaluation dropping missing values. Average Peer Quality is the average of the baseline quality score of other (artificial) group members. Share Peer with Business is the share of other (artificial) group members that had a business at baseline. Panel A restrict the sample to those who submitted the milestone 0 (baseline proposal) before the randomization and for which we have a baseline quality score. Panel B uses the full sample and replaces by 0 the quality score of those who did not submit milestone 0 before the randomization.

Table A9: **Effects of peer composition (II): surprise component**

	(1)	(2)	(3)	(4)
	Quality Ext.	Quality Int.	Quality Ext.	Quality Int.
Panel A. Sample who submitted M0 on time in Large Countries				
Virtual-within interaction	0.186^{***}	0.169^{***}	0.180^{**}	0.146^{**}
	(0.068)	(0.058)	(0.073)	(0.067)
Own Baseline Quality	0.447^{***}	0.249^{***}	0.446^{***}	0.247^{***}
	(0.050)	(0.045)	(0.050)	(0.045)
Surprise Peer Quality	-1.231^{**}	-1.599^{***}		
	(0.468)	(0.365)		
Expected Peer Quality	0.488	-5.615		
	(5.196)	(5.124)		
Surprise Peers with Business			-0.240	-0.357^{*}
			(0.156)	(0.191)
Expected Peers with Business			1.100	-2.210
			(2.323)	(2.370)
Number of Entrepreneurs	1,231	758	1,231	758
Panel B. Full Sample in Large Countries				
Virtual-within interaction	0.092^{**}	0.125^{*}	0.090^{**}	0.117
	(0.043)	(0.071)	(0.043)	(0.070)
Own Baseline Quality	0.453^{***}	0.244^{***}	0.450^{***}	0.247^{***}
	(0.050)	(0.044)	(0.050)	(0.046)
Surprise Peer Quality	-0.327	-1.499^{***}		
	(0.342)	(0.502)		
Expected Peer Quality	4.053	-8.975		
	(4.238)	(5.346)		
Surprise Peers with Business			-0.124	-0.424^{**}
			(0.144)	(0.165)
Expected Peers with Business			0.661	0.937
			(1.250)	(1.841)
Number of Entrepreneurs	2,222	899	2,222	899

Notes: The table uses data for entrepreneurs in the virtual interaction arms in the Large-Country Sample (LCS), both for the virtual-within and virtual-across treatment arms. *Surprise Peer Quality* (*Surprise Peer with Business*) is the difference between average baseline quality of peers (share of peers with and ongoing business) and their corresponding expectations, the latter computed as the average across 1,000 realizations of the group assignment randomization. In all regressions we include strata fixed effects. We also control for own baseline quality, replacing missing values with zeros and adding a corresponding dummy. Standard errors are clustered at the group level (reported in parenthesis). The number of stars (*, **, ***) codes for statistical significance at (10%, 5%, and 1%), respectively. In columns (1) and (2) the outcome is the quality score obtained by the business proposals of the entrepreneurs when missing values are replaced by zeros, whereas in columns (3) and (4) the entrepreneurs with missing values are removed from the sample. Panel A restricts to those who submitted Milestone 0 on time (before the randomization) and have been assigned a baseline quality score. Instead, Panel B uses the full sample and imputes a value of 0 to those who did not submit Milestone 0 on time.

A.4 Peer effects: complete results

In Tables A10-A12 below, we provide a complete account of the regression results partially presented in Table 10.

Table A10: **Network peer effects, benchmark setup, complete results on submission**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Peer effect	.533***	.014	.098
	(.174)	(.14)	(.296)
Baseline quality	.071***	.053*	.015
	(.025)	(.031)	(.037)
Female	-.031	-.098*	.026
	(.051)	(.051)	(.071)
Age	.002	.003	.009
	(.002)	(.002)	(.004)
Language English	-.174	-.128	.16
	(.17)	(.237)	(.158)
Language Arabic	.333***	.101	-.047
	(.105)	(.217)	(.232)
Language French	.038	.209	-.05
	(.118)	(.139)	(.071)
Language Other	-.004	-.211**	-.028
	(.061)	(.098)	(.075)
University complete	.012	.036	.102
	(.045)	(.056)	(.068)
Eastern Africa	-.037	.097*	-.012
	(.047)	(.053)	(.096)
Middle Africa	-	-	-.063
	(-)	(-)	(.113)
Southern Africa	.067	-.027	-.021
	(.088)	(.141)	(.153)
Northern Africa	-	-	.349*
	(-)	(-)	(.191)
Has Business	-.012	.083**	-.027
	(.049)	(.041)	(.075)
Has Idea about New Business	-.064	-.037	-.155**
	(.05)	(.054)	(.073)
Has written business plan	.084*	.012	.056
	(.048)	(.044)	(.064)
Sector: Manufacture	.018	-.018	.046
	(.061)	(.068)	(.139)
Sector: Mining	-.149	.226	-
	(.275)	(.148)	(-)
Sector: Other	-.044	.082	.131
	(.062)	(.062)	(.086)
Sector: Retail	.102	-.154	.01
	(.116)	(.096)	(.184)
Sector: Services	-.009	-.077	.198*
	(.073)	(.073)	(.107)
Sector: Social Entrepreneurship	.018	.009	.124
	(.082)	(.062)	(.116)
Sector: Technology	.042	-.064	.073
	(.063)	(.072)	(.127)
Risk Aversion	.005	0	-.018
	(.008)	(.009)	(.022)
Trust Measure	-.003	.005	-.002
	(.008)	(.008)	(.01)
Number of people discuss business	.014**	.019***	.023***
	(.006)	(.007)	(.009)
Number Facebook friends	0	0	0
	(0)	(0)	(0)
Time spent on Facebook	-.001	0	-.001
	(.001)	(.001)	(.003)
Number Twitter friends	0	0	0
	(0)	(0)	(0)
Time spent on Twitter	.002	.001	-.002
	(.002)	(.001)	(.003)
Number of entrepreneurs: 1016			

Notes: A complete account of the regression results on submission partially presented in Table 10.

Table A11: **Network peer effects, benchmark setup, complete results on extensive quality**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Extensive margin			
Peer effect	.327**	.06	.537***
	(.132)	(.131)	(.137)
Baseline quality	.307***	.300***	.17
	(.098)	(.11)	(.123)
Female	-.082	-.316*	-.02
	(.172)	(.185)	(.247)
Age	.022**	.024***	.011
	(.011)	(.008)	(.013)
Language English	-1.636***	.853**	.304
	(.543)	(.388)	(.540)
Language Arabic	.008	2.605***	1.429**
	(.377)	(.450)	(.645)
Language French	.05	2.221***	-.161
	(.354)	(.482)	(.218)
Language Other	.197	-.740**	-.287
	(.17)	(.335)	(.346)
University complete	.08	.182	.664***
	(.173)	(.153)	(.233)
Eastern Africa	.055	.388**	-.558
	(.156)	(.166)	(.403)
Middle Africa	-	-	-.698
	(-)	(-)	(.449)
Southern Africa	.289	-.328	-.466
	(.399)	(.493)	(.742)
Northern Africa	-	-	-.323
	(-)	(-)	(.69)
Has Business	-.005	.463**	-.293
	(.189)	(.183)	(.262)
Has Idea about New Business	-.074	-.287	-.105
	(.157)	(.177)	(.269)
Has written business plan	.301*	.109	.253
	(.178)	(.142)	(.238)
Sector: Manufacture	-.059	-.209	.152
	(.257)	(.209)	(.459)
Sector: Mining	-1.446**	-.243	-
	(.561)	(.555)	(-)
Sector: Other	-.159	.282	.17
	(.202)	(.296)	(.399)
Sector: Retail	.13	-.262	.215
	(.458)	(.414)	(.570)
Sector: Services	-.066	-.051	.317
	(.272)	(.278)	(.355)
Sector: Social Entrepreneurship	.081	-.036	.161
	(.300)	(.262)	(.338)
Sector: Technology	.320	-.132	-.101
	(.265)	(.263)	(.480)
Risk Aversion	.011	.003	-.013
	(.033)	(.035)	(.064)
Trust Measure	-.053*	.029	.042
	(.028)	(.021)	(.036)
Number of people discuss business	.061***	.057**	.08**
	(.023)	(.024)	(.035)
Number Facebook friends	0	0	0
	(0)	(0)	(0)
Time spent on Facebook	-.003	-.001	-.002
	(.004)	(.003)	(.008)
Number Twitter friends	-.001**	.001*	0
	(0)	(0)	(.001)
Time spent on Twitter	.005	.001	-.008
	(.005)	(.005)	(.009)
Number of entrepreneurs: 1016			

Notes: A complete account of the regression results on the extensive quality margin, partially presented in Table 10.

Table A12: **Network peer effects, benchmark setup, complete results on intensive quality**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Intensive margin			
Peer effect	-.031	.087	.476***
	(.121)	(.150)	(.149)
Baseline quality	.144*	.218***	.192**
	(.084)	(.077)	(.086)
Female	-.025	-.123	-.118
	(.152)	(.147)	(.176)
Age	.02***	.019**	-.018*
	(.007)	(.008)	(.011)
Language English	-1.261**	1.172***	-.324
	(.602)	(.366)	(.466)
Language Arabic	-.641*	2.388***	1.566**
	(.356)	(.439)	(.700)
Language French	-.109	1.747***	.027
	(.466)	(.611)	(.252)
Language Other	.19	-.224	-.357
	(.191)	(.31)	(.263)
University complete	.003	.127	.608***
	(.194)	(.176)	(.209)
Eastern Africa	.183**	.118	-.522*
	(.073)	(.155)	(.308)
Middle Africa	-	-	-.646
	(-)	(-)	(.359)
Southern Africa	.239	-.256	-.313
	(.217)	(.422)	(.637)
Northern Africa	-	-	-1.429**
	(-)	(-)	(.565)
Has Business	-.012	.323**	-.382
	(.132)	(.148)	(.275)
Has Idea about New Business	.091	-.223*	.551**
	(.088)	(.126)	(.22)
Has written business plan	.154	.116	.367*
	(.155)	(.119)	(.19)
Sector: Manufacture	-.105	-.197	-.301
	(.245)	(.202)	(.378)
Sector: Mining	-1.482***	-.886	-
	(.456)	(.533)	(-)
Sector: Other	-.043	.091	-.347
	(.223)	(.269)	(.330)
Sector: Retail	-.088	.187	.318
	(.393)	(.488)	(.405)
Sector: Services	.029	.197	-.602*
	(.221)	(.256)	(.346)
Sector: Social Entrepreneurship	.068	-.103	-.278
	(.23)	(.263)	(.295)
Sector: Technology	.278	.042	-.411
	(.260)	(.184)	(.372)
Risk Aversion	.002	.007	.028
	(.023)	(.035)	(.039)
Trust Measure	-.05***	.019	.068**
	(.018)	(.018)	(.035)
Number of people discuss business	.027*	.009	.027
	(.016)	(.018)	(.024)
Number Facebook friends	0	0	0
	(0)	(0)	(0)
Time spent on Facebook	0	-.004	.003
	(.003)	(.003)	(.007)
Number Twitter friends	0	.001**	.001**
	(0)	(0)	(.001)
Time spent on Twitter	-.002	-.002	-.009
	(.004)	(.004)	(.009)
Number of entrepreneurs: 779			

Notes: A complete account of the regression results on the intensive quality margin, partially presented in Table 10.

A.5 Peer effects: semantically weighted networks under bidirectional influence

In Tables A13-A14 below, we provide the counterparts of Tables 20 and 21 for semantically weighted networks under bidirectional influence, as the latter is described for the benchmark setup in Subsection 4.3.3.

Table A13: **Network peer effects, business network under bidirectional influence**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Control for ***all*** baseline info.	.852***	-.285	-.409
	(.127)	(.352)	(.430)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Control for ***all*** baseline info.	.422***	-.018	.442**
	(.137)	(.177)	(.172)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Control for ***all*** baseline info.	.052	-.112	.400**
	(.130)	(.174)	(.188)
Number of entrepreneurs: 779			

Notes: The details are as in Table 20, except that the matrix $M + M'$ used to construct the bidirectional network (see Subsection 4.3.3) is based on a measure of *aggregate interaction flow* (see (LF) in Section 4.1) that relies on the link weights induced by the influence network projected on the business-content category.

Table A14: **Network peer effects, sentiment network under bidirectional influence**

	Virtual within, LCS	Virtual across, LCS	Virtual across, SCS
Proposal submission			
Control for ***all*** baseline info	.859***	-.341	.116
	(.133)	(.254)	(.385)
Number of entrepreneurs: 1016			
Quality: extensive margin			
Control for ***all*** baseline info.	.480***	-.043	.504***
	(.128)	(.181)	(.185)
Number of entrepreneurs: 1016			
Quality: intensive margin			
Control for ***all*** baseline info.	.112	-.065	.336*
	(.137)	(.163)	(.172)
Number of entrepreneurs: 779			

Notes: The details are as in Table 21, except that the matrix $M + M'$ used to construct the bidirectional network (see Subsection 4.3.3) is based on a measure of *aggregate interaction flow* (see (LF) in Section 4.1) that relies on the link weights induced by the influence network projected on the sentiment category.