

When Confidence and Competence Collide: Effects on Online Decision-Making Discussions

Liye Fu
Dept. of Computer Science
Cornell University
liye@cs.cornell.edu

Lillian Lee
Dept. of Computer Science
Cornell University
llee@cs.cornell.edu

Cristian Danescu-Niculescu-Mizil*
Dept. of Information Science
Cornell University
cristian@cs.cornell.edu

ABSTRACT

Group discussions are a way for individuals to exchange ideas and arguments in order to reach better decisions than they could on their own. One of the premises of productive discussions is that better solutions will prevail, and that the idea selection process is mediated by the (relative) competence of the individuals involved. However, since people may not know their actual competence on a new task, their behavior is influenced by their self-estimated competence — that is, their *confidence* — which can be misaligned with their actual competence.

Our goal in this work is to understand the effects of confidence-competence misalignment on the dynamics and outcomes of discussions. To this end, we design a large-scale natural setting, in the form of an online team-based geography game, that allows us to disentangle confidence from competence and thus separate their effects.

We find that in task-oriented discussions, the more-confident individuals have a larger impact on the group's decisions even when these individuals are at the same level of competence as their teammates. Furthermore, this unjustified role of confidence in the decision-making process often leads teams to under-perform. We explore this phenomenon by investigating the effects of confidence on conversational dynamics. For example, we take up the question: do more-confident people introduce more ideas than the less-confident, or do they introduce the same number of ideas but their ideas get more uptake? Moreover, we show that the language people use is more predictive of a person's confidence level than their actual competence. This also suggests potential practical applications, given that in many settings, true competence cannot be assessed before the task is completed, whereas the conversation can be tracked during the course of the problem-solving process.

Keywords

confidence; decision-making; conversations; linguistic; teams; overconfidence; synergy; group dynamics; collaboration; ideas

*Corresponding senior author.

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW 2017, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4913-0/17/04. <http://dx.doi.org/10.1145/3038912.3052681>



1. INTRODUCTION

It is hard to argue with Griffin and Tversky's [15] opening statement, "The weighing of evidence and the formation of belief are basic elements of human thought".¹ Griffin and Tversky go on to say that one of the major findings from psychology, philosophy and statistics regarding how to evaluate evidence and assess confidence is that "people", ranging from the ordinary to the expert, "are often more confident in their judgments than is warranted by the facts".² At the same time, many studies have also found underconfidence to be rampant, as well; for an entry into that literature, see Erev et al. [12], or, for an intriguingly titled study, [27]. Thus, confidence and competence are often misaligned.³

What effect does this misalignment of confidence and competence have on

- (a) group decision-making, and
- (b) the dynamics of group discussions?

With respect to (a), for example, who has more influence on the group's decision: those who are overconfident but under-competent, or those who are competent but less confident? Similarly, with respect to (b): do the overconfident but under-competent tend to dominate the entire conversation, or do the more-competent eventually have more say? Studies on (a) [30, 34, 41] have generally shown that the overconfident have greater clout and that overconfidence leads to sub-optimal outcomes; but these experiments have been relatively small in scale or involve simulations. Topic (b) appears understudied.

In contrast, *the work we present in this paper represents a first investigation into both the conversational and decision-making impact of confidence-competence misalignment at a large scale and in an online natural setting.* Our experimental platform is an online collaborative geography puzzle game that has been played by over 10,000 teams; in the game, players have potentially different "views" of a given location and try to determine that location by pooling their information. The system records each individual's guess before they interact with other teammates, each individual's confidence in their guess, the entire team conversation, and the accuracy of the final answer that the team settles on. (See Section 2.2 for more details.)

In our data, more than half the users indicate a confidence level that does not match their competence level, with these miscalibra-

¹This is not to say these are *all* the elements of human thought.

²See also, among others, [22, 23], the latter giving rise to the phrase "Dunning-Kruger effect".

³"One of the painful things about our time is that those who feel certainty are stupid, and those with any imagination and understanding are filled with doubt and indecision. I do not think this is necessary." — Bertrand Russell

tion errors being committed both by people with low competence and by people with high competence. (See Section 3 for more details.) We thus have a substantial number of instances of under- and over-confidence to analyze in the context of group decision-making discussions.

Our findings with respect to point (a) — effects on group decision-making — are that: beyond competence, confidence gives people additional control over team decisions (Section 4.1). While this result is in line with those of [41], we additionally explicitly examine the impact of degrees of confidence differences between group members, controlling for correctness. We see that *when the most competent individual in a team is less confident than the least competent one, the team generally performs worse compared to teams with the same composition in terms of correctness but in which individuals have accurate confidence estimates*. This suggests that teams where misalignment exists between individuals’ relative confidence and relative competence fail to reach their potential in terms of performance (Section 4.2).

We are thus naturally led to delve into the novel research topic (b) — the interplay between confidence and the dynamics of group discussions — to help understand what happens in a group conversation that leads to the effects just described. *For instance, we examine the role confidence plays in the idea-selection process: do more-confident people introduce more ideas than the less-confident, or do they introduce the same number of ideas but their ideas get more uptake?* We find that in fact, more-confident individuals tend to introduce more ideas in discussions, even when there are not any more competent than their teammates (Section 4.3).

Finally, as a practical contribution, we show that there are linguistic cues that are predictive of a person’s confidence level, suggesting that interfaces for group decision-making could potentially make use of confidence information to help in the collaboration process (Section 5).

Terminological note: from “competence” to “correctness”. The term “competence” is used in the related psychology literature to indicate a person’s actual observed performance in a given task, and we would have preferred to be consistent with that literature. However, in this paper, we are discussing both “confidence” and “competence”, and the two words unfortunately (and perhaps ironically) sound so similar as to potentially introduce confusion. We therefore use the term “correctness” instead of “competence” in most of the remainder of this paper.

2. EXPERIMENTAL SETUP

2.1 Large-scale online discussion setting

To study the role confidence plays in the dynamics of online discussions and decision-making processes, we need a natural setting where we can gather participants’ confidence labels and observe details of the group decision-making process. With these constraints in mind, we design an experimental platform in the guise of an online team-based world exploration game, StreetCrowd [31].⁴

StreetCrowd is played in teams of at least two players, and is built around a geographic puzzle: the system chooses some spot on Earth, and players are given access (only) to ground-level Google StreetView images of that spot, which they can navigate to try to determine where the spot is. Each game consists of two stages:

⁴This game has been approved by the Cornell University’s IRB (Protocol 1504005555) and is currently live at <http://streetcrowd.us/start>; an unaffiliated gameplay video is available at <https://youtu.be/Lp47w-1Wsn0>.

Table 1: Excerpt from a chat during the team phase. The true answer for this puzzle is France. Player 1 has guessed the country correctly, but is only sure the location is in Europe. Player 2 and 3 guessed New Zealand and Nepal respectively. During the chat, players attempt to convince each other, while making intermediate team marker placements during the process. In this example, player 1 tried to persuade the team to guess France whereas the other two players argued for New Zealand.

| | |
|------------------|---|
| <i>player 1:</i> | its in the alps somewheres |
| <i>player 2:</i> | Crazy amount of sheep |
| <i>player 2:</i> | I put new zealand |
| <i>player 3:</i> | my sources say no |
| <i>player 1:</i> | well maintained mnt road says europe |
| <i>player 2:</i> | good point |
| <i>player 3:</i> | new zealand is good shout |
| <i>player 2:</i> | the land looks a bit arid though |
| <i>player 2:</i> | thats why i thought nz |
| <i>player 1:</i> | well NZ is like europe in that way but dont notice any of their typical vegetation |
| <i>player 3:</i> | nepal |
| <i>player 2:</i> | I think the road is too nice to be in nepal |
| <i>player 1</i> | <i>has moved the team marker [to France]</i> |
| <i>player 2:</i> | and yeah, nz is like europe but not in europe lol |
| <i>player 2:</i> | lol |
| <i>player 2</i> | <i>has moved the team marker [to New Zealand]</i> |

- **Solo phase – exploration but no discussion.** There is no communication between players at this stage. Each player has three minutes to independently explore the surrounding environment and find information that may help them locate the spot, e.g., geographical landmarks, apparent climate, whether cars drive on the left-hand or the right-hand side of the road, and so on. The player makes an *individual guess* by placing a marker on an interactive world map. To proceed to the team phase, the player is prompted for a reason (in a few words and a confidence level for her guess).
- **Team phase – discussion but no exploration.** At this point, further navigation is disabled, but players in a team are placed in a chat room where they exchange ideas and opinions in an attempt to decide on a single *team guess*. They have access to a shared map and to a marker that each one of them can move. When all players have agreed on the position of the marker on the map, or the time limit is reached, the game ends. (See Table 1 for an example team discussion.)

In this setting, we can generate a virtually infinite number of puzzles with known correct answer — the location L_{true} used to generate the respective StreetView. We can also directly measure how correct a guess G is based on the the distance to the true location $\text{dist}(G, L_{\text{true}})$ (we use great-circle distance). We also record a discrete level of correctness based on how precise the guess is: we use reverse geocoding to compare the coarse geographical location of the guess with that of the correct answer (e.g., do they both fall in the same country?), and categorize guesses into the four levels of correctness, which we call *guess precision*, depicted in Table 2.

Table 2: Numerical correctness levels based on guess precision.

| guess precision | correctness level |
|----------------------------------|-------------------|
| wrong continent | 1 |
| correct continent, wrong country | 2 |
| correct country, wrong region | 3 |
| correct region | 4 |

Table 3: Possible choices for self-estimating guess precision and corresponding numerical confidence levels. For validation, we also show average zoom levels used by players declaring the respective confidence levels. (At zoom level 1, the entire world is displayed, whereas at zoom level 10, players are looking at roughly the city level.)

| estimated precision | confidence level | zoom level |
|-------------------------|------------------|------------|
| “Could be anywhere” | 1 | 4.3 |
| “I know the continent!” | 2 | 4.6 |
| “I know the country!!” | 3 | 5.4 |
| “I know the region!!!” | 4 | 7.8 |

Over the past year, StreetCrowd has accumulated nearly 12,600 team games (comprising about 30,000 individual-phase games) involving more than 5,800 unique players. We discard games in which individual players fail to make a guess, the team does not agree on a final guess, or no communication occurs. We also heuristically filter out games that may involve cheating and games for which we could not identify the coarse location of a player’s guess automatically.

2.2 Measuring confidence

Studies of confidence in the psychology literature have largely considered three possible operationalizations of confidence [28]:

1. Confidence as a self-estimation of (a concrete measure of) performance, that is, how well one did or will do. Examples include a student’s estimate of her score on an exam [5, 8], or an estimate of the time it takes to finish a particular task [6].
2. Confidence as the level of certainty in one’s answer. This type of confidence can be extracted by, for example, asking participants to indicate how sure they are of their answer on an ordinal scale [29].
3. Confidence as placement of one’s ability relative to others. Studies that use this definition of confidence may ask participants to estimate their performance percentile with respect to all the other participants [25], e.g., “top 5%”.

Employing operationalization 1. In our work, we focus on the first operationalization of confidence, because it concerns a direct comparison between an individual’s actual performance and her estimate of her performance. Misalignment between the two can then be interpreted in terms of *overconfidence* — when the individual’s estimated performance is above her actual performance — and *underconfidence* — when the estimated performance is under the actual performance. For instance, player 1 in Table 1 exemplifies underconfidence as she has guessed the correct country, yet she only estimated her correctness to be at the continent level.

To capture this notion of (over)confidence in StreetCrowd, we compare guess precision with *estimated precision*, the four levels

of which are given in Table 3. At the end of the solo phase, players are prompted to indicate how good they think their guess is by clicking on one of the four estimated-precision buttons, whose labels are shown in the leftmost column of Table 3. The correspondence between the correctness levels in Table 2 and confidence levels in Table 3 is deliberate, because *the signed difference between the numerical confidence level and the numerical correctness level gives us a degree of misalignment*: a larger confidence level relative to correctness level (e.g., “I know the country” but the guess was really “correct continent, wrong country”) indicates overconfidence; when the confidence level is smaller than the correctness level, we have underconfidence; the two values being equal indicates correct calibration.

To check if people are indicating their self-evaluated confidence reasonably — as opposed to clicking randomly or primarily based on how the buttons are placed in the interface — we compare the guess-precision levels with how much the users zoomed in on the map before marking their guess. Higher zoom levels reflect a greater level of detail. We expect players who believe their guess is more precise to zoom more, and indeed in our data we do observe that players with higher confidence levels tend to have higher zoom levels on average (Spearman’s $\rho = 0.30$, p -value < 0.001 , Table 3).⁵

The other operationalizations. To compare with relevant literature, in an earlier version of StreetCrowd, we also evaluated confidence according to the second operationalization listed above: how certain an individual is in her answer. We collected these labels from about 1,800 games by prompting the individual to indicate the level of certainty in her guess on a scale from 0% to 100% using a slider. (We excluded users who did not move the slider at all.) However, we abandoned this design in favor of the estimated-precision approach outlined above: offering users the relatively objective comparison points of continents, countries, and within-country regions⁶ was preferable to trying to interpret what, say, “65%” means or deciding whether one person’s “65%” could definitely be considered more confident than another person’s “60%”.

In this work we do not consider the third operationalization of confidence — relative placement among others — which requires the individuals to jointly consider other people’s correctness as well as their own, for two reasons: (1) this would require post-hoc surveys of the participants after playing the game, which would be harder to collect in a natural setting, and (2) it would be hard to distinguish whether misalignments between estimated and actual performance are caused by participants misjudging their own ability, or misjudging the ability of the others.

3. CONFIDENCE MISCALIBRATION

As mentioned in the Introduction, a person’s confidence does not necessarily align with her actual objective ability. Earlier studies suggest that this miscalibration follows systematic patterns driven by the level of competence of the individuals. For instance, several researchers [23, 29] argue that one needs a certain level of competence in order to reasonably calibrate one’s confidence, and that as a result, non-competent individuals tend to overestimate their performance. Indeed, an extensive series of studies from the field of education dating back to the 1930’s [36] (see [5] for a comprehensive survey) consistently indicate that when estimating their grades, “weak” students tend to overestimate their grades, whereas “strong” students have a (less pronounced) tendency to underesti-

⁵We use Spearman’s ρ throughout our study for reporting correlations. Kendall’s τ gives similar qualitative result.

⁶We take states returned by Google and Bing geocoding services as regions.

mate their grades. In this section, we use our online setting to analyze the nature of this systematic bias in confidence miscalibration in more detail, and at a much larger scale than was previously possible. This builds the backdrop against which we later examine its consequences on the dynamics and outcomes of decision-making discussions.

We examine misalignments with respect to two different common operationalizations of confidence introduced earlier in Section 2.2:

Certainty in one’s answer. For the subset of the data where players indicated their confidence via a slider, we group players by their correctness level (as listed in Table 2). For each such correctness group, we compute the correlation between the actual distances between each participant’s initial guess and the true location of the puzzle, $dist(G, L_{true})$, and their indicated level of confidence. We find that for players whose location was on the wrong continent, their self-reported confidence level has no significant correlation with distance between their guess and the true location (Spearman’s $\rho = -0.006$, p -value = 0.88), whereas in contrast, the correlation is significant for all other better-performing groups. In fact, the correlation is the strongest for the group of players who hit the correct region ($\rho = 0.41$, $p < 0.001$). This confirms that one needs a certain level of competence to reasonably evaluate one’s performance.

Self-estimation of performance. Since the previous operationalization of confidence-as-certainty is not rooted in an objectively measurable outcome, misalignments cannot be directly translated into under- or over-confidence. For this reason, for the remainder of this paper, we operationalize confidence as a self-estimation of performance, a common interpretation in the psychology and education literature [5, 23, 28]. This allows us to non-ambiguously map misalignment between the estimated performance and the actual performance directly to under- and over-confidence.

More than half of the users fail to correctly estimate their performance. Figure 1 breaks this failure down by correctness level, and shows that indeed, in agreement with [23, 29], individuals that are more wrong are more likely to overestimate their performance (blue dashed line), whereas increasing correctness corresponds to increasing underconfidence.

We note here two experimental limitations that are intrinsic to any experimental study of confidence misalignment. First, there is a boundary effect in that individuals achieving maximum performance on a task cannot be overconfident: e.g., a student that receives an exam grade of A+ cannot produce an estimate of a higher grade; symmetrically, individuals that achieve minimum performance can not be underconfident. Second, some individuals can appear to be under-confident simply by chance: e.g., a student answering a multiple choice question can select the correct answer by chance and simply appear to be more skillful than they really are. This effect does not apply symmetrically to over-confidence: a student knowing the answers to a multiple choice question is much less likely to select the incorrect answer by chance and appear less skillful than in reality. This asymmetry alone could explain why underconfidence can appear to be more common than overconfidence. Acknowledging that such limitations naturally apply to our setting as well, we refrain from making any claims regarding the relative propensity of under- and over-confidence from the data depicted in Figure 1.

4. CONFIDENCE IN TEAM INTERACTION

The promise of group over individual deliberation lies in the intuitive notion that the more people we bring together, the more ideas for solving a problem will be generated, and the more ideas generated, the more likely one of them will provide an excellent




Figure 1: Misalignment between competence and confidence. For each correctness level: the fraction of players who are underconfident, at right level of confidence, and overconfident. More competent individuals are less likely to be overconfident.

solution [19, 26, 39]. Ideally, this process could lead to a collective performance that surpasses what individuals can achieve on their own, i.e., the interaction will have a *synergistic* effect.⁷

In practice, however, this “sum greater than the parts” promise is far from guaranteed [31]. The process through which solutions surface is complex and relies on constant negotiation between the participants, and “process losses” [35] can also occur. While in an ideal setting the more competent individuals would have a larger say, in many realistic scenarios the actual level of competence on a new task is unknown. The decision-making process is then driven in part by self-estimations of competence (which, as we have discussed above, can be heavily misaligned with actual competence), and decisions made through heuristics relying on such estimates can be worse than desired [2]. In this section we examine what are the consequences of confidence (mis)alignment on the decision-making process, on the degree of resulting synergy and on the overall dynamics of the discussion.

4.1 Confidence and relative influence

When a team makes a decision, different individuals can contribute to the proposed solution to different degrees. Here we examine the role confidence plays in regulating the relative influence different team members have on the final decision.

Relative influence. In our setting, we can estimate the influence on the team decision by considering how much closer to her individual guess the player has managed to get the team guess to be, relative to the other player. More formally, we define influence of a given player p_1 (with answer G_{p_1}) on a team answer G_{team} relative to p_2 (with answer G_{p_2}) by:

$$inf(p_1, p_2) = \frac{dist(G_{p_2}, G_{team})}{dist(G_{p_2}, G_{team}) + dist(G_{p_1}, G_{team})},$$

noting that this measure of *relative influence* is agnostic to how correct the solution is.

⁷See also the related concepts of collective intelligence [40] and assembly bonus effect [10].




Figure 2: The larger the difference in correctness between the best and the worst player, the more relative influence the best player has on the final decision. A difference in correctness of 0 corresponds to the case when all team members were at the same level of correctness; in this case the “best” and the “worst” player have equal contributions (i.e., relative influence of 0.5), as expected since their roles are randomly assigned. Throughout, error bars indicate standard error.

For each team, we pair the player who achieved the highest level of correctness with the player that has the lowest level. Importantly, in cases in which all players have the same level of correctness, best and worst players are arbitrarily chosen. For each team, we compute how much more control the best player has over the team decision relative to the worse player, i.e., $\text{inf}(p_{\text{best}}, p_{\text{worst}})$.⁸

Controlling for correctness. As previously discussed, in order for the team discussions to be productive, it is desirable to have the more competent member take more control and guide the team in their decisions. It is comforting to see that, indeed, the best player has a greater relative influence on the team solution than the worse player in the team: Figure 2 shows that the larger the difference in correctness between the best player and worst player, the more relative influence the best player has on the final decision.

Given that confidence and correctness are, at least to some extent, correlated ($\rho = 0.46$, p -value < 0.001), it is important to control for the difference in correctness in order to disentangle the effect of confidence. To this end, we group games based on the four possible values of the difference in correctness and we compare the effect of confidence within these correctness-controlled groups.

Effects on relative influence. By comparing across correctness-controlled groups (columns) in Figure 3, we observe that on average the more confident the best player is relative to the worst player, the greater her influence on the final decision. Notably, this trend also holds even when the team members are equally correct (see first column, where “best” and “worst” roles are randomly assigned; $\rho = 0.21$, p -value < 0.001). *Beyond competence, confidence gives people unjustifiably greater control in the decision-making process.*

4.2 Confidence and team performance

When combined with the observation that non-competent individuals are likely to overestimate their competence (Section 3), the result we just discussed suggests that misalignment between the competence and confidence of team members can prevent teams from achieving their potential. Driven by overconfidence, the least

⁸For clarity we focus our analysis on the best and worst players in the team. The majority ($> 60\%$) of the games only have two players.




Figure 3: The more confident the best player is relative to the worst player, the more influence (darker color) she has on the decision, even after controlling for difference in correctness (comparing across columns). Spearman’s correlation coefficients are computed using all data points within a group and are indicated, together with corresponding significance levels, for each column and row. Throughout, buckets with less than 15 instances are discarded. Statistical significance levels are indicated as: $*$: $p < 0.05$, $$: $p < 0.01$, $***$: $p < 0.001$.**

competent individuals can take control of the decision process, with harmful effects on the team outcome.

Measuring team synergy. While many studies use correctness of the team answer as an indication of team performance [9, 14], this type of measure neglects differences in initial team composition: when a team of A students is doing better than a team of C students, it does not automatically mean they have collaborated better. Here we focus instead on measuring how much a team improves over the potential of its members.

To measure such improvement, or synergy, for a team of n players, we first compare the distance between the team guess G_{team} and the true answer L_{true} to the average of distances between each player’s individual guess to the true location. This gives the team a constructiveness score [31], which provides a control for the initial individual performance of the players in the team:

$$c_{\text{avg}} = \frac{1}{n} \sum_{i=1}^n \text{dist}(G_{p_i}, L_{\text{true}}) - \text{dist}(G_{\text{team}}, L_{\text{true}}).$$

In our setting, however, a team whose average guess is already close to the true location will have smaller maximum possible c_{avg} compared to teams that are farther from the correct answer. This can happen either when a puzzle is relatively easy or when a team is composed of particularly strong players. To account for this, we normalize the average score of the individuals in the team to obtain a team synergy score:

$$s_{\text{avg}} = \frac{c_{\text{avg}}}{\frac{1}{n} \sum_{i=1}^n \text{dist}(G_{p_i}, L_{\text{true}})}.$$

Higher values indicate greater team improvement over the individual performance of the team members. We note that this measure can take negative values if the team’s performance is worse than the average performance of its members.

Effects of confidence misalignment. To test our intuition that misalignment between competence and confidence can be harmful to the decision-making process, we first consider the 227 teams where




Figure 4: When there is a real distinction between the best and the worst player (difference in correctness > 0), the more confident the best player is relative to the worst player, the greater synergy the team achieves (darker color), controlling for difference in correctness (comparing across columns).

the worst player is overconfident to the point that she is more confident than the best player. We find that 37% of these teams have negative synergy, compared to only 26% of the rest of the teams (2-proportion z-test p -value < 0.001).

To explore the effect of misalignment more systematically, we compare the difference in confidence level of the best and the worst player, after controlling for correctness (Figure 4). Since we defined synergy as a symmetric team-based measure, we do not expect to observe any effect of confidence when the team members are equally correct (first column). In all other cases, we notice that *the larger the difference in confidence between the best and the worst player, the more productive the team discussion is: the player with the best solution is likely to lead the team towards better results*. In particular, teams where the confidence and correctness are misaligned such that the difference in correctness is larger than the difference in confidence — i.e., the cells above the main diagonal (where difference in confidence equals difference in correctness) — exhibit on average lower average synergy. This suggests that *misalignment in confidence can be harmful to group discussions*, preventing groups from reaching their potential in terms of performance.

4.3 How confidence mediates team interactions

The effects of confidence on team performance are intriguing, considering that we control for the actual competence of the individuals. We now turn to examine the mechanisms through which confidence mediates group interactions, focusing on a dimension of conversational dynamics that is particularly pertinent in the decision-making process: the introduction and discussion of ideas [31, 42].

We consider two possible hypotheses that could explain the observed effects of confidence-competence alignment:

- Hyp 1** More-confident individuals contribute more ideas. If the best individual is also the most confident one, the team benefits from having better quality ideas to consider in order to make steady progress to better solutions.
- Hyp 2** Ideas introduced by less-confident individuals are easier to discard. If competence and confidence are aligned, this can help the team focus on the ideas introduced by the more competent individuals.




Figure 5: The more confident the best player is, the more ideas she introduces (relative to the worst player). Horizontal dashed line indicates equal contribution of ideas by the best and worst player. Vertical dashed line indicates the case when best and worst individuals are equally confident. Significance level for Spearman’s correlation coefficient is indicated for each difference in correctness level (each line).

To gather valid ideas for a particular puzzle, we consider all nouns, proper nouns and adjectives that are not stopwords as candidate ideas. If any of these candidate ideas has been adopted (i.e., mentioned by at least two different participants) by any team attempting the respective puzzle it is included in the set of ideas for that puzzle. Example ideas discussed in the game illustrated in Table 1 include “sheep”, “alps” and “arid”. We also add to this puzzle-specific collection a set of geographical terms,⁹ such as “Nepal”, to provide better coverage of potential ideas. While we find this simple method effective for our particular setting, we acknowledge that more sophisticated methods for detecting, representing and tracking ideas are an important direction for future work (Section 7).

To investigate the effect of relative confidence on the number of new ideas a player introduces to a discussion, we compare the ratio between the number of ideas introduced by the best player and the number of ideas introduced by the worst player, controlling for difference in correctness. We find that the more confident the best player is relative to her teammate (larger difference in confidence), the more ideas she tend to introduce relative to their teammates (Figure 5). Particularly noticeable is the case where there is no difference in correctness between team players (difference in correctness = 0, p -value < 0.001): the more confident, but equally correct, individual still introduces more ideas.

It may be conceivable that the difference in number of ideas introduced may simply be explained by confident players talking more during the games. Yet, there is no clear correlations between number of words spoken and player confidence, indicating that it is the ideas in the chat messages that are more important.

Since our setting also allows players to suggest solutions by moving the team marker during the discussion, we also examine marker movements as an alternative means of understanding the dynamics of the decision-making process that is less sensitive to the noise inherent to extracting ideas from natural language. We observe similar trends as in the case of idea introduction: the more confident the best player is relative to her worst teammate, the more she contributes to team’s marker movements (Figure 6).

Together, these results support the first hypothesis linking confidence with conversational dynamics: confidence seems to regulate

⁹Source: Wikipedia’s glossary of geography terms.




Figure 6: The more confident the best player is, the more marker movements she makes relative to the worst player. Horizontal dashed line indicates equal marker placements.

the exchange of ideas in decision-making discussions, such that the more-confident individuals introduce more ideas.

The effect of confidence in filtering and selecting ideas in teams is less clear. We find no significant correlation between the level of confidence and the percentage of ideas that are eventually adopted by other team members.

5. LINGUISTIC CUES OF CONFIDENCE

So far we have discussed how confidence mediates team interaction and performance. However, in online chats, the only means that a person (or a system) has to judge or estimate the confidence of another is by picking up cues from what the other person says. Therefore, in this section we ask whether there are linguistic cues that are indicative of an individual’s confidence. Motivated by our previous results, we focus on identifying signals that can not simply be explained by actual competence.

We formulate this exploration as a paired prediction task: given a pair of equally competent teammates, can we predict from their language which one is more confident? If such predictive signals exist, it follows that confidence — rather than competence alone — can influence how people talk when collaborating.

Analyzing the language of the conversations is complicated by the fact that utterances tend to be short texts that are sprinkled with misspellings and abbreviations common in online chats; see Table 1 for an example chat and the last two columns of Table 4 for additional example utterances from individual players. To find particularly predictive language signals, we consider three types of features and analyze their strength in predicting confidence.

Word choice (bag of words, or BOW). We expect that an individual’s choice of words conveys hints as to her confidence level. In order to find predictive word choices that are generalizable to other, non-geography scenarios, we discard words that are tagged as nouns (e.g., “continent”), proper nouns (e.g., “Canada”), adjectives (e.g., “Portuguese”) and a set of manually identified adverbs and verbs.¹⁰ We also discard words that appear in only one puzzle. This filtering process leaves us a vocabulary size of 474 words.

Ideas. Figure 5 suggests that an individual’s confidence level is reflected in the number of ideas she brings into the discussion,

¹⁰A manual scan suggests that these three word types do indeed cover the majority of geography-related terms, which is our intent, although helpful domain-independent cues such as the adjective “sure” may unfortunately be excluded by this procedure as well.

where our notion of “ideas” was defined above in Section 4.3. Because it may be the case that *when* ideas are introduced may also indicate an individual’s confidence, in addition to (1) the total number of ideas introduced during the entire game, we also consider the number of ideas introduced at two special stages: (2) before the start of the discussion, when a player can provide a reason for her initial individual guess (recall that this explanation is made instantly available to the other teammates once the team phase commences); and (3) at the very early stages of the conversation (i.e., within each player’s first three utterances).

Additional indicators. We also examine the use of *hedges* and expressions of agreement, both of which are, intuitively, cues for lack of confidence. Hedges, which are expressions of uncertainty or lack of commitment, have been studied extensively in the linguistics literature (a classic reference is [24]), and identifying hedging is an active area of NLP research [13]. We adopt a list of hedging terms [37], such as “apparently” or “in my opinion”, created by manual curation of prior collections [13, 16, 18]. We consider agreement cues because their occurrence potentially indicates that the speaker was not too confident in their own opinions; we use the list of agreement cues from [31].

To set up our competence-controlled paired prediction task, we match the most and least confident players in a team (not allowing ties) only if the two players are equally correct in their individual guesses.¹¹ This results in 747 matched pairs from 216 unique puzzles. Pair features are simply obtained by taking the difference between the feature values for each of the two players, e.g., for idea features, we take the number of ideas introduced by the first player in the pair subtracted by the number of ideas introduced by the second player in the pair.

We use logistic regression and evaluate performance using a leave-one-puzzle-out approach: in each fold, we reserve instances from one puzzle as the test set and train on instances from all other puzzles; this further controls for effects of puzzle-specific cues. We report accuracies macro-averaged by puzzle, such that performance on each puzzle has equal weight regardless of how many instances it contains.

With our setup, both random guessing and using correctness as a predictor would yield 50% accuracy. The length of the utterances, i.e., the number of words a player speaks, is also found not to do better than random guessing.

BOW features can predict significantly better than chance or correctness alone (binomial test p -value < 0.001), with a macro-average accuracy of 58%. This suggests that there are indeed linguistic signals that are predictive of the difference in confidence between equally-competent teammates. We also note that adding the additional indicators does not bring noticeable improvement.

Notably, the three ideas features can on their own perform significantly better than chance or correctness (p -value < 0.001), achieving 60% macro-averaged accuracy. While all idea introduction features have positive coefficients (i.e., introducing more ideas signals more confidence), the order of relative importance of these features suggests that the early stages of the conversation are particularly indicative of relative confidence:

| feature type | coef. |
|---------------------------------------|-------|
| number of ideas in reason | 1.05 |
| number of ideas in first 3 utterances | 1.05 |
| number of ideas overall | 0.50 |

¹¹To account for data sparsity and simplify the task we group the two highest confidence levels together. We also considered the task of detecting absolute, rather than relative, confidence and obtained qualitatively similar results.

Table 4: Linguistic cues for confidence, placed in post-hoc categories. The “more-confident” and “less-confident” columns list words that are among the top 30 features in at least half the test folds; daggers mark words that are among the top 10 features in at least half the test folds. In the example utterances, word of the type represented by the row are shown in bold.

| category | more-confident | less-confident | more-confident example | less-confident example |
|---------------------|---|---|--|---|
| personal pron. | we | me, you [†] | i think we just go up in the mountains | looks like thailand to me . |
| interrogatives | which | how, what [†] , where, why | i found some writing which looked mandarin | what do you think ? |
| thinking (tense) | think | thought, thinking [†] | i think it is nearer to the ocean | i was thinking brazil too |
| hedges and modals | probably, maybe, might, should, can | guess, could | it might be too far north | lol i guess so |
| certainty | 100[%], actually, pretty [sure] | | i'm pretty sure this is a nordic country | |
| agreement | | too, sounds [good] | | ahh sounds good to me |
| negations | not [†] | no [†] , dont | its definetly not that side | i have no idea |
| coordinating conj. | and [†] , but | | building style and cars are al quite uk ish | |
| subordinating conj. | | though, as [†] | | though it could be a forgotten part of Italy |
| demonstrative pron. | these, there [†] , this | that [†] , thats | trust me on this one! | lets try that one |
| articles | the [†] | | look at the statues | |
| filler words | | ah, haha [†] | | no clue haha |
| others | it, is [†] , its [†] , by, from, in [†] , on [†] , somewhere [†] , found, said, saw [†] , far, move | for, about, have [†] , did [†] , do, lets, looks, makes, put, near, all, your | | |

5.1 Linguistic features of interest

In order to understand possible sources of confidence signals beyond the introduction of ideas, we now look at the BOW linear regression coefficients. Table 4 shows words that ranked among the top 30 features in at least half of the folds, grouped into post-hoc categories to emphasize several distinctions between the language used by more-confident vs. less-confident individuals.

In terms of the use of personal pronouns (first row of Table 4), more confident individuals are more likely to speak on behalf of the team, whereas less confident individuals tend to stick more with first and second person perspectives. Interestingly, this echoes in part findings showing that in small groups leaders are generally more other-focused, preferring first-person plural to first-person singular [7, 21]. In this regard, more confident individuals seem to assume an ad-hoc leadership role (even when they are equally-competent).

On the other hand, less confident players tend to consult others' opinions more, as indicated by the increased use of interrogatives (second row). We also note that “did” and “do” categorized under “others” are mostly used in questions involving such interrogatives.

When expressing thoughts (third row), less confident individuals tend to use past tense,¹² which might be an indication that they perceive themselves as less authoritative members in the dialogue, according to some studies in linguistics [20].

The patterns for features that are intuitively, or at least anecdotally, associated with confidence, such as hedging, certainty and

agreement, are listed in the second block of rows in Table 4. These patterns do not align completely with common intuition.

More confident individuals do indeed tend to use more phrases indicating certainty, and express less agreement with other's opinions. On the other hand, hedges and modal verbs, are not necessarily reserved to the less confident people. For instance, “might” and “probably” are associated with the more confident people; one hypothesis worth exploring further is that confident individuals make use of certain types of hedges as a strategy to gently persuade their partners [38].

There are also interesting differences in terms of the surface structure of the utterances (third block) with more confident individuals being more likely to use coordinating conjunctions and the definite article “the”, potentially in an attempt to bring more concrete evidence into the discussion. On the other hand, the utterances of less confident individuals are sprinkled with filler words.

This shallow post-hoc analysis of words signaling confidence (or lack thereof) can serve as a starting point for future work on more complex linguistic and conversational features for confidence detection. The fact that such predictive cues exist even in a setting that controls for competence underlines the role confidence has in shaping people's language in group discussions. Furthermore, this suggests that future interfaces for facilitating group decision-making could potentially extract confidence information from the language of the participants.

¹²In these chats, past continuous tense (“was thinking”) is more prevalent than present continuous tense (“am thinking”).

6. ADDITIONAL RELATED WORK

Confidence has been studied as a factor affecting human behavior. Horowitz [17] has shown that there is a significant negative correlation between a person's level of imitation behavior and the difference between the person's confidence relative to others. Confidence as approximated by poll scores is shown to correlate with how frequent candidates shift topics in political debates [32]. In the group opinion forming processes, highly confident individuals have been identified as one of the two major attractors of opinion [29], based on simulations on models derived from controlled experiments. However, most of these experiments merely present participants the opinions and confidence levels of others, without allowing any real interactions between them.

The correlation between confidence and correctness has been studied in many contexts. Studies relating the confidence of eyewitness to the actual correctness [4] and self-reported confidence to observed competence of junior medical officers [3] both indicate important problems with the use of confidence as a predictor of correctness. Miscalibration has been noted for a long time, with most studies focusing on overconfidence [1, 11, 23]. See Moore and Healy [28] for a reconciliation of existing studies on overconfidence.

7. CONCLUSIONS AND FUTURE WORK

Our main contribution in this paper is to investigate pre-existing and new research questions regarding the effects of confidence, particularly when misaligned with correctness, on group decision-making and group discussions, doing so at a large scale and in a natural online setting.

Our results, as well as the limitations of our setting, point naturally to several directions for future work:

Confidence prediction and miscalibration detection. In our explorations of possible linguistic cues of confidence, we simplified matters by binarizing confidence labels (more confident vs. less confident) and applying basic cue-extraction techniques. Given that we do discover that there are signals in language that capture confidence, one natural extension is to aim at more accurately predicting confidence at a finer-grained level, perhaps by applying more sophisticated natural language processing techniques. We could build on such work to address the more ambitious goal of detecting misalignment between confidence and correctness, with the eventual aim of facilitating group discussion by ameliorating the effects of under- and over-confidence on team conversations and outcomes.

Finer-grained measures of confidence and competence. In this work, we discretized players' confidence and competence using four coarse levels. As discussed in Section 3, this discretization allows us to identify clear cases of under- and over-confidence. In future work, finer-grained measures of confidence and competence might enable better controls, while providing a more detailed analysis of the extent of confidence-competence (mis)alignment effects in team interactions.

Dynamics of team discussions. We have looked into conversational dynamics in terms of idea introduction and selection, where we employed a simple heuristic to track ideas. Pursuing better methods of representing ideas, capturing idea introduction, identifying idea adoption, and further distinguishing good and bad ideas would be beneficial in further understanding the dynamics of discussions.

On a related note, in this work we have only considered the initial confidence of individual players. Studies have looked into how individual confidence combines to form eventual dyadic confidence in collaborative decisions [33], and it would be interesting to track

and understand how individuals' confidence in their own guess and in their team's working solution changes over time.

Causality. We have revealed correlations between confidence and influence in teams, as well as between confidence and team synergy arising from interaction. Our full control over the game interface gives us the ability to conduct experiments involving the introduction of stimuli that artificially manipulate individuals' confidence levels; this ability opens the possibility of studying not just correlation, but also causality.

Confidence in other decision-making contexts. Will our findings generalize beyond StreetCrowd? While we have striven to analyze our data in a non-domain specific fashion, studying the effects of confidence in team decision-making in other contexts, such as group work in Massive Open Online Courses, or in the process of eRulemaking, is an important direction for future work.

Acknowledgments. We are grateful to Eric Horvitz and Sendhil Mullainathan for inspiring discussions on the role of confidence in conversations. We would also like to thank Pantelis Analytis, David Byrne, Claire Cardie, Susan Fussell, Mala Gaonkar, Malte Jung, Dan Jurafsky, Jon Kleinberg, Anita W. Woolley, Justine Zhang, and the anonymous reviewers (all of whom were supremely competent!) for their insightful (and confident) comments, and Chenhao Tan for sharing the list of hedges. We are also grateful to all StreetCrowd players and to David Garay, Maheer Iqbal, Jinjing Liang, Vlad Niculae, and Neil Parker for participating in the game's development together with the first author. This work was supported in part by a Discovery and Innovation Research Seed Award from the Office of the Vice Provost for Research at Cornell.

8. REFERENCES

- [1] M. D. Alicke. Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology*, 49(6):1621, 1985.
- [2] D. Bang, R. Fusaroli, K. Tylén, K. Olsen, P. E. Latham, J. Y. Lau, A. Roepstorff, G. Rees, C. D. Frith, and B. Bahrami. Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and Cognition*, 26:13–23, 2014.
- [3] L. Barnsley, P. M. Lyon, S. J. Ralston, E. J. Hibbert, I. Cunningham, F. C. Gordon, and M. J. Field. Clinical skills in junior medical officers: A comparison of self-reported confidence and observed competence. *Medical Education*, 38(4):358–367, 2004.
- [4] R. K. Bothwell, K. A. Deffenbacher, and J. C. Brigham. Correlation of eyewitness accuracy and confidence: Optimality hypothesis revisited. *Journal of Applied Psychology*, 72(4):691, 1987.
- [5] D. Boud and N. Falchikov. Quantitative studies of student self-assessment in higher education: A critical analysis of findings. *Higher Education*, 18(5):529–549, 1989.
- [6] R. Buehler, D. Griffin, and M. Ross. Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, 67(3):366, 1994.
- [7] J. Cassell, D. Huffaker, D. Tversky, and K. Ferriman. The language of online leadership: Gender and youth engagement on the internet. *Developmental Psychology*, 42(3):436, 2006.
- [8] D. E. Clayson. Performance overconfidence: Metacognitive effects or misplaced student expectations? *Journal of Marketing Education*, 27(2):122–129, 2005.

- [9] D. Coetzee, S. Lim, A. Fox, B. Hartmann, and M. A. Hearst. Structuring interactions for large-scale synchronous peer learning. In *Proceedings of CSCW*, 2015.
- [10] B. Collins and H. Guetzkow. *A social psychology of group processes for decision-making*. Wiley, 1964.
- [11] D. Dunning, K. Johnson, J. Ehrlinger, and J. Kruger. Why people fail to recognize their own incompetence. *Current Directions in Psychological Science*, 12(3):83–87, 2003.
- [12] I. Erev, T. S. Wallsten, and D. V. Budescu. Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3):519–527, 1994.
- [13] R. Farkas, V. Vincze, G. Móra, J. Csirik, and G. Szarvas. The CoNLL-2010 shared task: Learning to detect hedges and their scope in natural language text. In *Proceedings of CoNLL: Shared Task*, 2010.
- [14] H. Friedberg, D. Litman, and S. B. F. Paletz. Lexical entrainment and success in student engineering groups. In *Proceedings of the IEEE Spoken Language Technology Workshop*, 2012.
- [15] D. Griffin and A. Tversky. The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3):411–435, 1992.
- [16] D. A. Hanauer, Y. Liu, Q. Mei, F. J. Manion, U. J. Balis, and K. Zheng. Hedging their bets: The use of uncertainty terms in clinical documents and its potential implications when sharing the documents with patients. In *Proceedings of the AMIA Annual Symposium*, 2012.
- [17] H. Horowitz. The relationship of confidence to competence and imitation behavior. *Journal of Psychology*, 63(2):235–247, 1966.
- [18] K. Hyland. *Hedging in scientific research articles*. Pragmatics and Beyond New Series. John Benjamins Publishing, 1998.
- [19] M. H. Jackson and M. S. Poole. Idea-generation in naturally occurring contexts. *Human Communication Research*, 29(4):560–591, 2003.
- [20] B. Johnstone. ‘He says ... so I said’: Verb tense alternation and narrative depictions of authority in American English. *Linguistics*, 25(1):33–52, 1987.
- [21] E. Kacewicz, J. W. Pennebaker, M. Davis, M. Jeon, and A. C. Graesser. Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33(2):125–143, 2014.
- [22] A. Koriat, S. Lichtenstein, and B. Fischhoff. Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2):107–118, 1980.
- [23] J. Kruger and D. Dunning. Unskilled and unaware of it: How difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, 77(6):1121, 1999.
- [24] G. Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2, 1975.
- [25] R. P. Larrick, K. A. Burson, and J. B. Soll. Social comparison and confidence: When thinking you’re better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes*, 102(1):76–94, 2007.
- [26] P. R. Laughlin, E. C. Hatch, J. S. Silver, and L. Boh. Groups perform better than the best individuals on letters-to-numbers problems: effects of group size. *Journal of Personality and Social Psychology*, 90(4):644, 2006.
- [27] S. Lichtenstein and B. Fischhoff. Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance*, 20(2):159–183, 1977.
- [28] D. A. Moore and P. J. Healy. The trouble with overconfidence. *Psychological Review*, 115(2):502–517, 2008.
- [29] M. Moussaïd, J. E. Kämmer, P. P. Analytis, and H. Neth. Social Influence and the Collective Dynamics of Opinion Formation. *PLOS One*, 8(11):e78433, 2013.
- [30] M. A. Neale and M. H. Bazerman. The effects of framing and negotiator overconfidence on bargaining behaviors and outcomes. *Academy of Management Journal*, 28(1):34–49, 1985.
- [31] V. Niculae and C. Danescu-Niculescu-Mizil. Conversational markers of constructive discussions. In *Proceedings of NAACL-HLT*, 2016.
- [32] V. Prabhakaran, A. Arora, and O. Rambow. Power of confidence: How poll scores impact topic dynamics in political debates. In *Proceedings of the ACL Joint Workshop on Social Dynamics and Personal Attributes*, 2014.
- [33] J. P. Schuldt, C. F. Chabris, A. W. Woolley, and J. R. Hackman. Confidence in dyadic decision making: The role of individual differences. *Journal of Behavioral Decision Making*, 2015.
- [34] J. A. Sniezek and R. A. Henry. Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, 43(1):1–28, 1989.
- [35] I. Steiner. *Group Processes and Productivity*. Social Psychology. Academic Press, 1972.
- [36] F. Sumner. Marks as estimated by students. *Education*, 32:429, 1932.
- [37] C. Tan and L. Lee. Talk it up or play it down? (Un)expected correlations between (de-)emphasis and recurrence of discussion points in consequential U.S. economic policy meetings. Presented in Text as Data, 2016.
- [38] C. Tan, V. Niculae, C. Danescu-Niculescu-Mizil, and L. Lee. Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Proceedings of WWW*, 2016.
- [39] J. S. Valacich, A. R. Dennis, and J. F. Nunamaker. Group size and anonymity effects on computer-mediated idea generation. *Small Group Research*, 23(1):49–73, 1992.
- [40] A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688, 2010.
- [41] P. Zarnoth and J. A. Sniezek. The social influence of confidence in group decision making. *Journal of Experimental Social Psychology*, 33(4):345–366, 1997.
- [42] J. Zhang, R. Kumar, S. Ravi, and C. Danescu-Niculescu-Mizil. Conversational flow in Oxford-style debates. In *Proceedings of NAACL*, 2016.