Judgment and Decision Making, Vol. 11, No. 2, March 2016, pp. 168–173

“Isn’t everyone like me?”: On the presence of self-similarity in strategic interactions

Ariel Rubinstein* Yuval Salant†

Abstract

We propose that in strategic interactions a player is influenced by *self-similarity*. Self-similarity means that a player who chooses some action X tends to believe, to a greater extent than a player who chooses a different action, that other players will also choose action X. To demonstrate this phenomenon, we analyze the actions and the reported beliefs of players in a two-player two-action symmetric game. The game has the feature that for “materialistic” players, who wish to maximize their own payoff, there should be negative correlation between players’ actions and the beliefs that they assign to their opponent choosing the same action. We first elicit participants’ preferences over the outcomes of the game, and identify a large group of materialistic players. We then ask participants to choose an action in the game and report their beliefs. The reported beliefs of materialistic players are positively correlated with their actions, i.e., they are more likely to choose an action the stronger is their belief that their opponent will also choose the same action. We view this pattern as evidence for the presence of self-similarity.

Keywords: self-similarity, strategic justification, false consensus, game theory

1 Introduction

This paper proposes that in strategic interactions a player’s beliefs about his opponent’s action are influenced by *self-similarity* considerations. Self-similarity means that a player who chooses some action X tends to believe, to a greater extent than a player who chooses a different action, that other players will also choose action X. This effect would contrast with the conventional game-theoretic approach, which postulates that a player’s beliefs provide *strategic justification* for his chosen action, in the sense that the chosen action is optimal for him with respect to his beliefs.

To demonstrate the existence of self-similarity we require a game in which the effect of self-similarity on beliefs differs from that of strategic justification. We look for a game with two actions, A and B, in which strategic justification requires that choosers of A believe to a lesser extent than choosers of B that their opponent is likely to choose A. One such game is the following two-player symmetric game:

	Opponent’s choice	
Player’s choice	A	B
A	30, 30	30, 70
B	70, 30	0, 0

(The two numbers in each cell of the table represent payoffs for the Player (left) and his Opponent (right).)

In this game, a player can guarantee himself a payoff of 30 by choosing the “safe” action A. Alternatively, he can take a risk and choose the action B in which case he obtains a payoff of 70 if his opponent chooses A and a payoff of 0 if his opponent also chooses B.

A “materialistic” player wishes to maximize his own payoff. He considers the payoff of his opponent only when comparing outcomes in which the player himself obtains the same payoff. Such a player should choose the safe action A if he believes that his opponent is sufficiently likely to choose the risky action B. On the other hand, if he believes that his opponent is sufficiently likely to choose A, then he should choose B. The exact threshold for switching between the two actions is idiosyncratic to the player. Thus, if the reported beliefs of materialistic players provide strategic justification for their actions, then the proportion of materialistic players who choose A should *decrease* with the probability that players assign to their opponent choosing A. If we find that the A-choice proportion of materialistic players *increases* with the probability they assign to their opponent choosing A, this result would constitute evidence in favor of the presence of self-similarity.

The first version of this paper was circulated in July 2014 under the title “They do what I do: Positive Correlation in Ex-Post Beliefs”. We thank Hadar Binsky for his assistance in analyzing the data and Eli Zvuluny for constructing and maintaining the web platform on which the research was conducted. We also thank Jon Baron, Karl Schlag, Andy Schotter, Jörg Spenkuch, and James Tremewan for their feedback. The first author acknowledges financial support from ERC grant 269143.

*Tel Aviv University and New York University. Email: rariel@post.tau.ac.il.

†Northwestern University. Email: y-salant@northwestern.edu.

However, not all players are materialistic. Some players also care about their opponent's payoff even at the expense of their own. We classify a player as "cooperative" if he views the outcome (A, A) as best and the outcome of (B, B) as worst. (An outcome (X, Y) summarizes the choices made by the player [X] and his opponent [Y]). In the game presented above, such a player is willing to forgo increasing his payoff by 40 without hurting his opponent (as revealed by his preference for (A, A) over (B, A)), and prefers not to sacrifice 30 in order to "punish" his opponent for choosing B (as revealed by his preference for (A, B) over (B, B)). A cooperative player's dominant strategy (at least in the weak sense) is to choose A. Therefore, his reported beliefs cannot conflict with strategic justification no matter what they are.

Participants in our main experiment were asked to play the game against an anonymous opponent who had been selected randomly from among several hundred individuals. They first received a description of the game and were then asked to report their preferences over its four possible outcomes. Participants were then randomly assigned to one of two treatments: Treatment 1, in which participants were asked to choose an action and then report their beliefs about their opponent's action and Treatment 2, in which they reported their beliefs and then chose an action.

We find that 74% of the participants are materialistic. In this group, the proportion of players who chose A *increases* significantly with the probability that they assign to their opponent choosing A. This is true in both treatments but the effect is stronger in Treatment 1, in which players report their beliefs after choosing an action. Thus, the reported beliefs of materialistic participants are consistent with self-similarity and cannot be explained by strategic justification.

A minority of 13% of the participants are cooperative. Almost all of them chose A. The beliefs of these participants are similar to those of the materialistic participants who chose A. However, unlike in the case of the materialistic participants, this does not constitute evidence in favor of self-similarity since every possible belief of a cooperative player is consistent with strategic justification for choosing A.

Self-similarity is related to the false consensus phenomenon in non-strategic decision making, whereby individuals tend to perceive their own choices and judgments as being more common in the population than they actually are. The false-consensus effect was first studied by Ross, Greene and House (1977) and has been documented since then in a variety of settings and explained by various psychological theories. Marks and Miller (1987) survey this literature, and Dawes (1989) and Dawes and Mulford (1996) provide a critical discussion of the effect.

In the context of strategic decision making, Dawes, McTavish and Shaklee (1977) and Messé and Sivacek (1979) report that players in the Prisoner's Dilemma (PD) game tend to attribute their own action to other players. Dawes et al. (1977) conjecture that either strategic justification or self-similarity (or both) may explain this finding.

We conducted a separate experiment that illustrates the unsuitability of the PD game for demonstrating self-similarity. In the experiment, 184 participants (from the same pool as in the main experiment) played the following PD game:

	Opponent's choice	
Player's choice	C	D
C	50, 50	0, 60
D	60, 0	10, 10

and then reported their preferences over the game's four outcomes.

A large group of 36% of the participants reported cooperative preferences in which they rank (C, C) as the best outcome and (C, D) as the worst. Strategic justification requires that the proportion of these participants who choose C *increases* with the probability that they assign to their opponent choosing C. Thus, the fact that such participants assign a higher probability to their opponent playing C might be due to strategic justification rather than self-similarity.

Another large group of 39% of the participants reported materialistic preferences. These participants have a dominant strategy of choosing D, and indeed 83% of them do so. Their beliefs cannot conflict with strategic justification no matter what those beliefs are.

Therefore, the PD game is not appropriate for differentiating between self-similarity and strategic justification as causes for the positive correlation between a player's action and his belief that his opponent will choose the same action. In contrast, self-similarity operates in the opposite direction of strategic justification for materialistic players in our game, thus making it possible to test for the presence of self-similarity.

2 Method

The didactic site http://gametheory.tau.ac.il served as the platform for the experiment. Participants were recruited by email from a pool of current and former students dispersed among a large number of countries, almost all of whom had taken an undergraduate course in game theory. The participants had used the site previously during their studies and had agreed to participate in additional online survey experiments.

On the opening screen of the experiment, participants were told that they are about to play a game against an opponent who had been selected randomly from among several hundred people. They received a description of the game including the following payoff table:

	Your opponent's choice	
Your choice	A	B
A	30, 30	30, 70
B	70, 30	0, 0

It was promised that, from among all the pairs of players that would be playing the game, five would be chosen randomly to actually receive payment in US dollars according to the result of their game.

Participants were then asked to answer three questions:

The first concerned their preferences over the outcomes of the game: "Please rank the four possible results of the game according to your preferences. Assign 1 to the best result, assign 2 to the second-best result, and so on. If you are indifferent between two results, assign the same number to both of them." The four possible results of the game were presented to participants in the form of "You choose X; your opponent chooses Y." The order of presentation was randomized for each participant.

The second question asked the participants to choose action A or B.

The third question asked the participants to report their beliefs about the distribution of choices in the population of players: "What are your beliefs regarding the distribution of choices among the several hundred people who are also playing this game? I believe that ______ % will choose A and ______ % will choose B."

In answering these three questions, participants were randomly assigned to either:

Treatment 1: making a choice (question 2) before stating a belief (question 3)

Treatment 2: stating a belief (question 3) before making a choice (question 2).

3 Results and discussion

Of 604 respondents, we omitted 12 whose reported beliefs did not sum up to 99% or 100% and 45 who ranked the outcome (B, B) as weakly superior to the outcome (A, A). We were left with 547 participants.

3.1 Ranking of outcomes

There are 81 possible rankings of the four outcomes (allowing for indifferences), but the vast majority of the reported rankings belong to one of two groups: materialistic or cooperative.

About 74% of the participants are materialistic. They reported one of the three rankings in which the outcome (B, A) is ranked as the uniquely best outcome and the outcome (B, B) as the uniquely worst outcome. The three rankings differ in the comparison between (A, A) and (A, B). About 35% of the materialistic participants were indifferent between these two outcomes, 36% preferred (A, A) to (A, B), and 29% expressed a preference for (A, B) over (A, A).

Table 1: Reported beliefs in Treatment 1 (N=204, standard errors in parentheses).

Action in game:	A (N=129)	B (N=75)
Average belief in A	60.4% (1.7)	46.3% (2.5)
Average belief in B	39.6%	53.7%
Median belief in A	65%	45%

About 13% of the participants are cooperative. They reported one of the 8 rankings in which the outcome (A, A) is one of the best outcomes, the outcome (B, B) is one of the worst outcomes, and (A, A) is strictly preferred to (B, B). The three most popular rankings in this group were:

$(A, A) \succ (B, A) \succ (A, B) \succ (B, B)$ as reported by 44% of the group,

$(A, A) \succ (A, B) \succ (B, A) \succ (B, B)$ as reported by 20%, and

$(A, A) \succ (A, B) \sim (B, A) \succ (B, B)$ as reported by 17%.

(The symbol $\succ$ denotes strict preference and the symbol $\sim$ denotes indifference.) Cooperative participants have a (weakly) dominant strategy of choosing the action A.

The remaining 13% of the participants, who are neither materialistic nor cooperative, reported 16 different rankings. For example, 21% of them expressed a preference for maximizing the sum of payoffs in the game, i.e., they prefer the outcomes (A, B) and (B, A) to the outcome (A, A), which in turn is preferred to (B, B). Given the large number of rationales for the rankings in this group, we do not report any additional results about their choices or beliefs.

We now turn to analyzing the responses of the materialistic and cooperative participants to the remaining questions by treatment.

3.2 Materialistic participants

3.2.1 Treatment 1 (choice, then belief)

About 63% of the participants in this treatment chose action A and 37% chose action B. Table 1 summarizes the participants' reported beliefs according to their action. Clearly, choosers of A tend to believe that other players will also choose A to a greater extent than choosers of B.

Figure 1 provides a finer description of the data by means of two Cumulative Distribution Functions (CDFs) of the likelihood assigned to A, one for choosers of A and the other for choosers of B. For each group, we plot for every

Table 2: Proportion choosing A for various reported beliefs on A: actual and logistic regression values in %.

Intervals of belief on A	0–19	20–39	40–59	60–79	80–99	100
Number of observations in interval	9	34	53	79	27	2
Actual proportion choosing A	11	44	62	70	85	100
Regression prediction at mid-point of interval	28	44	60	75	85	89

Figure 1: CDFs of likelihood assigned to A in Treatment 1.

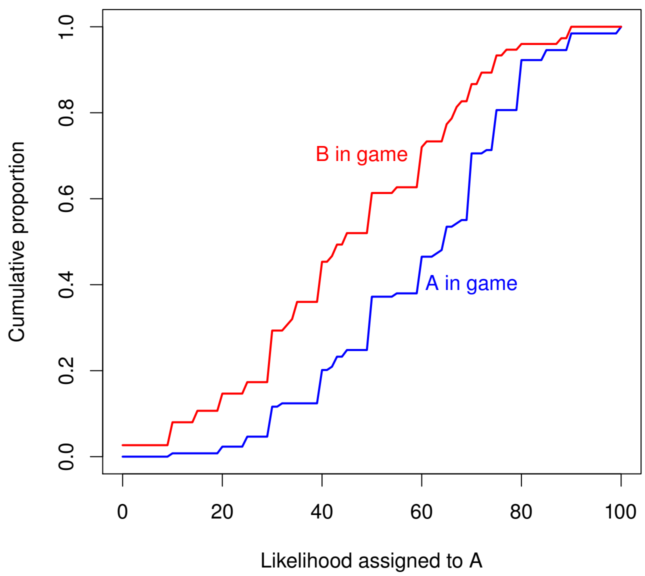

$0 \leq x \leq 100$ the proportion of participants who assigned a belief weakly less than $x\%$ to A.

The CDFs differ significantly (the Kolmogorov–Smirnov test statistic for the equality of the CDFs is 0.28 ($p = 0.001$)). Furthermore, the CDF for the choosers of A first-order stochastically dominates the CDF for the choosers of B. The first-order stochastic dominance reflects the strong tendency of choosers of A to assign higher likelihoods to others choosing A than choosers of B.

Table 2 presents the data from a different perspective, by reporting the proportion of participants choosing A as a function of their beliefs. The table also includes the predictions of a logistic regression of "choosing A" on the "belief that others choose A". (The regression predicts that $\text{logit(probability of choosing A)} = -1.26 + 3.36\text{(belief in A)}$, $OR = 28.86$, $p < 0.001$, where the variable "belief in A" is in the [0, 1] interval.) Clearly, the proportion of participants who choose A increases with the belief on A.

To summarize, two equivalent patterns emerge in Treatment 1. First, materialistic players tend to believe that others behave similarly to them (as shown in Table 1 and Figure 1). Second, the tendency to choose A increases with the likelihood that players assign to others choosing A (Table 2).

Table 3: Reported beliefs in Treatment 2 (N=204).

Action in game:	A (N=115)	B (N=89)
Average belief in A	58.7% (2.1)	51.0% (2.5)
Average belief in B	41.3%	49.0%
Median belief in A	66%	50%

These patterns are consistent with self-similarity. In contrast, they cannot be explained using only strategic justification. To see this, let p denote the probability that a player assigns to his opponent choosing A. A player compares the sure payoff of 30, which he obtains by choosing A, with the lottery $L(p) = (p)[70] + (1 - p)[0]$, which he obtains by choosing B. The belief p provides strategic justification for the player's action if the player prefers the lottery $L(p)$ to the sure payoff of 30. Assume that a player's beliefs are independent of his preferences over lotteries. In that case, expected utility predicts that if two players have the same preferences and the one with the higher p chooses A, then the one with the lower p also chooses A. This "single crossing property" holds for any theory of choice under uncertainty that has a measure representation (see Segal (1993)) and, in particular, for any rank-dependent model.

Thus, strategic justification predicts that the proportion of players who choose A decreases as p increases. Table 2 and the logistic regression, in contrast, indicate that the opposite is true: the larger the weight players assign to their opponent choosing A, the more likely they are to choose A. Thus, strategic justification cannot explain the patterns in the data.

3.2.2 Treatment 2 (belief, then choice)

Participants in this treatment were asked to report their beliefs before choosing an action. Their responses demonstrate self-similarity but to a lesser extent than in Treatment 1.

Table 3 summarizes the participants' reported beliefs according to their action. Choosers of A tend to believe that other players will choose A to a greater extent than choosers of B although the gap is smaller than in Treatment 1.

Table 4: Proportion choosing A for various reported beliefs on A: actual and logistic regression values in %.

Intervals of belief on A	0–19	20–39	40–59	60–79	80–99	100
Number of observations in interval	12	39	36	85	30	2
Actual proportion choosing A	50	36	58	62	63	100
Regression prediction at mid-point of interval	40	47	55	62	68	71

Figure 2: CDFs of likelihood assigned to A in Treatment 2.

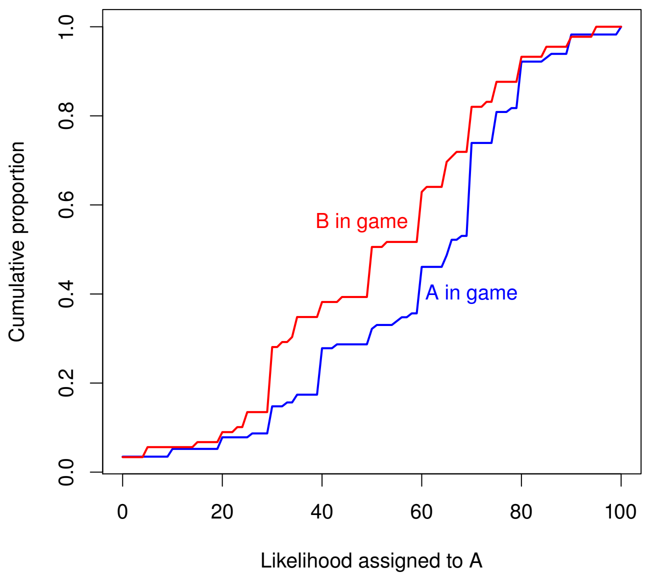

A similar pattern arises in Figure 2, which depicts the two CDFs of the likelihood assigned to A, one by choosers of A and the other by choosers of B. The CDFs differ significantly (the Kolmogorov–Smirnov test statistic for the equality of the CDFs is 0.21 ($p = 0.020$)), and the CDF for the choosers of A first-order stochastically dominates the CDF for the choosers of B. Again, the effect is smaller than in Treatment 1.

Finally, Table 4 reports the proportion of participants choosing A as a function of their beliefs and the predictions of the logistic regression. (The regression predicts that logit(probability of choosing A) $= -0.55 + 1.47$(belief in A), $OR = 4.33$, $p = 0.020$.) Again, the proportion of participants who choose A increases with the belief on A although the increase is less pronounced than in Treatment 1.

3.3 Cooperative participants

Recall that cooperative participants are those who rank the outcome (A, A) as one of the best outcomes and the outcome (B, B) as one of the worst. A participant with such a ranking should find the action A to be (at least weakly) dominant to action B and thus will choose A. Table 5 indicates that this is indeed the case: 68 of the 72 cooperative participants chose A.

For cooperative participants, action A dominates action B, so any belief that they express is consistent with strategic justification. The average and median reported beliefs of cooperative participants who chose A are about 60% in both treatments, which is remarkably similar to those of materialistic participants.

4 Final comments

We conducted an experiment to test for the presence of self-similarity in a game with the following three features. First, a large majority of players in this game rank its outcomes in a way that makes it possible to differentiate self-similarity from strategic justification. Second, each action is chosen by a significant proportion of players, thus enabling the comparison of beliefs across actions. Third, the game is simple and essentially context-free.

The main finding is that choosers of an action tend to believe that other players will choose the same action to a greater extent than choosers of the other action. This evidence for self-similarity extends the findings in the psychology literature on the false consensus phenomenon in individual decision making to the domain of strategic interactions. Researchers in experimental game theory should be aware of the self-similarity effect when interpreting the reported beliefs of participants in experiments.

4.1 The Chicken game

Prior to the above experiment, we conducted another in which participants were asked to choose an action in the following game of Chicken and then report their beliefs (but we did not ask them about their rankings over outcomes):

	Opponent's choice	
Player's choice	Dove	Hawk
Dove	30, 30	20, 70
Hawk	70, 20	0, 0

The results, which are summarized in Table 6, are similar to those in Table 1: Players tend to report beliefs that gravi-

Table 5: Summary statistics of cooperative participants by treatment.

	Treatment 1 (N=40)	Treatment 2 (N=32)
Proportion choosing A	93%	97%
Choosers of A: Average belief in A	60.6% (3.1)	59.7% (3.7)
Choosers of A: Median belief in A	60%	60%

Table 6: Reported beliefs in Chicken Game (N=244).

Action in game:	Dove (N=162)	Hawk (N=82)
Average belief in Dove	59.9% (1.7)	43.3% (2.6)
Average belief in Hawk	40.1%	56.7%
Median belief in Dove	67%	50%

tate toward their own action. However, these results are not conclusive evidence for the presence of self-similarity, since we do not know the participants' preferences over the outcomes of this game. If the pattern is Table 6 exists among materialistic participants, then it supports self-similarity, but if it exists only among cooperative players (who rank (Dove, Dove) as the best outcome and (Hawk, Hawk) as the worst), then it does not, since the preference ranking of cooperative players is such that Dove is a weakly dominant action.

In a separate experiment, we asked 263 participants to rank the outcomes in the above Chicken game. About 69% of them are materialistic and about 15% are cooperative. However, without knowing the link between their preferences over outcomes, their actions, and their reported beliefs, we cannot distinguish self-similarity from strategic justification. The experimental design in the current paper takes care of this problem by first asking participants to report their preferences over outcomes and then analyzing the correlation between actions and beliefs for the group of materialistic participants.

4.2 Ex-ante versus ex-post beliefs

It is a common practice among experimental game theorists to refrain from asking participants about their beliefs prior to choosing their action. This is because the belief elicitation question might encourage more strategic thinking and less instinctive thinking (see Costa-Gomes & Weizsäcker (2008) for a discussion), thus influencing participants' choices.

We asked players about their beliefs before (Treatment 2) and after (Treatment 1) choosing an action. The results are similar but not identical. First, asking players to report their beliefs prior to choosing an action reduces the proportion of choosers of A by 5 percentage points, from 68% to 63%. Second, the self-similarity effect is weaker when participants are asked about their beliefs ex-ante.

We speculate that the reason for these differences is that the safe action A is somewhat more intuitive than the risky action B. Asking players to report their beliefs prior to making a choice encourages them to think about their opponent more carefully and thus to be more strategic. One possible line of strategic reasoning is the following: "My instinct is to choose A. But my opponent is likely to do the same, and thus it is better for me to choose B". (Of course, this reasoning does not go to the next step of asking whether the opponent uses the same reasoning.) This line of reasoning, which is more likely to emerge when participants are asked to report their beliefs prior to choosing an action, weakens the self-similarity effect.

References

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies*, *75*(3), 729–762.

Dawes, R. M. (1989). Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology*, *25*(1), 1–17.

Dawes, R. M. & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes*, *65*(3), 201–211.

Dawes, R. M., McTavish, J., & Shaklee, H. (1977). Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and Social Psychology*, *35*(1), 1–11.

Marks, G., & Miller, N. (1987). Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin*, *102*(1), 72–90.

Messé, L. & Sivacek, J. M. (1979). Predictions of others' responses in a mixed-motive game: self-justification or false consensus? *Journal of Personality and Social Psychology*, *37*(4), 602–607.

Ross, L., Greene, D., & House, P. (1977). The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, *13*(3), 279–301.

Segal, U. (1993). The measure representation: A correction. *Journal of Risk and Uncertainty*, *6*(1), 99–107.