

Giovanna Devetag, Hykel Hosni and Giacomo Sillari
**You better play 7: mutual versus common
knowledge of advice in a weak-link
experiment**

**Article (Accepted version)
(Refereed)**

Original citation:

Devetag, Giovanna and Hosni, Hykel and Sillari, Giacomo You better play 7: mutual versus common knowledge of advice in a weak-link experiment. [Synthese](#). ISSN 0039-7857 (In Press)

© 2012 [Springer](#)

This version available at: <http://eprints.lse.ac.uk/47260/>

Available in LSE Research Online: November 2012

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final manuscript accepted version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this version and the published version may remain. You are advised to consult the publisher's version if you wish to cite from it.

You Better Play 7: Mutual *versus* Common Knowledge of Advice in a Weak-link Experiment

Giovanna Devetag*

Hykel Hosni†

Giacomo Sillari‡

October 31, 2012

Abstract

This paper presents the results of an experiment on mutual *versus* common knowledge of advice in a two-player weak-link game with random matching. Our experimental subjects play in pairs for thirteen rounds. After a brief learning phase common to all treatments, we vary the knowledge levels associated with external advice given in the form of a suggestion to pick the strategy supporting the payoff-dominant equilibrium. Our results are somewhat surprising and can be summarized as follows: in all our treatments both the choice of the efficiency-inducing action and the percentage of efficient equilibrium play are higher with respect to the control treatment, revealing that even a condition as weak as mutual knowledge of level 1 is sufficient to significantly increase the salience of the efficient equilibrium with respect to the absence of advice. Furthermore, and contrary to our hypothesis, mutual knowledge of level 2 induces, under suitable conditions, successful coordination more frequently than common knowledge.

*University of Perugia, Via Pascoli 20, 06123 Perugia.

†Scuola Normale Superiore, Piazza dei Cavalieri 7, 56100 Pisa and LSE Centre for Philosophy of Natural and Social Science, London WC2A 2AE

‡Corresponding author. Scuola Normale Superiore, Piazza dei Cavalieri 7, 56100 Pisa.

Keywords: Coordination games; experimental philosophy; epistemic attitudes, weak-link game; conventions

JEL codes: D01, D83

1 Epistemic attitudes: The formal and the experimental

The purpose of this paper is to discuss—from the point of view of formal epistemology—some experimental results on the role of common knowledge as a coordination device in a weak-link coordination game.

Common knowledge has long been regarded as a fundamental component in the modeling of rational interaction. As such it features prominently, though informally, in Nash’s characterization of equilibrium. As game theory evolved towards capturing increasingly more refined solution concepts, common knowledge was there to stay as a fundamental modeling assumption. Harsanyi and Selten, to mention a particularly relevant example from the point of view of the present paper, put forward a view according to which the assumption of common knowledge of the players’ rationality is sufficient to ensure the joint selection of efficient equilibria in (non-pure) coordination games.

In a line of research which for a long time has been developed almost independently of the game-theoretic one, epistemic logicians have advanced, over the past two decades, rigorous characterizations of the epistemic attitudes of both individual and group knowledge and belief. The emergence of multi-agent epistemic logic thus enabled, around mid 1990s, the development of rigorous analyses of the epistemic conditions underlying game theoretic solution concepts.

The ensuing investigations on the connection between rational agents’ epistemic attitudes and their strategic behavior tended to fall into two seldom communicating categories, the formal and the experimental. This paper constitutes an attempt at fostering the cross-fertilization between those lines of research. It aims at so doing by taking full advantage of the mutual feedback that experimental economics and formal epistemology can give one another. Formal epistemology lends methodological clarity to experimental analysis, especially concerning modeling assumptions. In turn, experimental analysis gives empirical substance to the formal analysis of epistemic atti-

tudes of group knowledge. This results in a virtuous circle, which we expect to be a key feature of the emerging field of experimental formal epistemology. From this vantage point, our contribution speaks to the community of experimental economists, to game theorists interested in the refinement of epistemic conditions for solution concepts, and to the relatively younger community of logicians and philosophers of logic interested in empirical evidence on modes of reasoning¹)

The paper is organized as follows. Section 2 describes the weak-link coordination game used in our experiment and Section 2.1 discusses the relevant related literature. The experimental design is then illustrated in Section 3. Section 4 is devoted to the analysis of the experimental results regarding choice behavior (Section 4.2) and the subjects' beliefs (Section 4.1). Section 6 draws some general conclusions about our rather surprising findings and delineates what we envisage to be promising future lines of development of the research reported in this paper. The remainder of this Section aims at illustrating the motivation of this work by presenting an outline of how the recent research on epistemic attitudes has tended to fall under the “formal” and the “experimental” headings, with very little overlap.

1.1 The formal

Robert Aumann opens his classic paper on correlated equilibrium by reviewing the concept of Nash Equilibrium and asks

why should any player assume that the other players will play their components [of a Nash Equilibrium] and indeed why should they? [...] In a two-person game, for example, Player 1 would play his component only if he believes that Player 2 will play his; this in turn would be justified only by 2's belief that 1 will play his component; and so on. (Aumann, 1987)

The von Neumann-Nash analysis of game-theoretic equilibrium relies essentially on the concept of best-response which is normatively prescribed to a rational player only under the assumption of the players' *common knowledge* of each others' rationality (and, clearly, of the game). It is certainly beyond the scope of the present paper to discuss which solution concept is appropriate for a justified formal characterisation of rational interaction. Yet what

¹See van Benthem (2008) and Verbrugge (2009) for two influential position papers, and D'Agostino (2010) for an example of a logic for realistic agents.

Aumann’s remark clearly illustrates is the centrality of *epistemic attitudes* in the game-theoretic analysis of rational interaction.

In hindsight two things appear to be rather surprising. The first is that game theorists, and economists in general, have been rather happy with an intuitive (as opposed to mathematically rigorous) notion of “knowledge” for their modeling until [Aumann \(1976\)](#) put forward what have become known as *Aumann structures*. In an nutshell, the construction of the Aumann structure begins, in continuity with the (single-agent) decision theoretic tradition, by postulating the existence of a *state space*, which is assumed to provide a complete description of the “world”. To each individual agent (in a finite set, interpreted as “player”) is associated a specific partition of the state space. As in Savage’s model, events – the object of the agents’ knowledge– are subsets of the set of states. The operations of individual, mutual and common knowledge are defined by suitably constrained operators on the algebra generated by the partitions of the state space. Aumann’s formalization has effectively given rise the actively researched area of epistemic game theory².

This leads us to the second surprising fact (in hindsight), namely that it took a relatively long time for game theorists and formal epistemologists to realize that they were working on essentially similar problems. Indeed, philosophers have long been interested in epistemic attitudes as specific instances of the wider class of *intentional* attitudes. Put very crudely, epistemic attitudes capture the relation between an agent’s reasoning (and disposition to act) and the information they possess. The first systematic attempt at giving this relation a rigorous logical formalization appeared in ([Hintikka, 1962](#)). This work paved the way for the syntactic development of epistemic logic which was soon to be complemented by the introduction of relational semantics (or Kripke frames)³. By the mid-1990s epistemic logic was a standard analytical tool in artificial intelligence ([Meyer and van der Hoek, 1995](#))

Since Hintikka’s pioneering work, epistemic logicians have tended to focus on normative models of individual *knowing* and *believing*⁴. These notions were extended to their multi-agent counterpart in the context of artificial

²See [Brandenburger \(2008\)](#) for a terse overview of the field.

³See e.g. [Blackburn et al. \(2001\)](#) for a state-of-the-art introduction to Modal logic and [Blackburn et al. \(2007\)](#) for a comprehensive account of its development.

⁴Recent attention has been devoted to the logic of “being informed” and that of “being aware”. See [Allo \(2011\)](#) and [Halpern and Pucella \(2007\)](#), [Halpern and Rêgo \(2008\)](#) for recent overviews. On awareness, the interested reader can consult Burkhard Schipper’s bibliography [Schipper \(2012\)](#).

intelligence⁵. In an attempt to make this paper as self-contained as possible, we now briefly recall the central epistemic attitudes captured in multi-agent epistemic logic, namely individual, mutual and common knowledge⁶.

Let N be a set of agents, S be a set of states (or “possible worlds”) and θ, ϕ etc. be sentences of some propositional (modal) logic. The key intuition captured by epistemic logics is that when being in epistemic state, say, s , agent i may *access* other epistemic states, say t . This motivates the introduction of an *accessibility* relation $R_i \subseteq S^2$ (one for each agent $i \in N$). We say that an agent $i \in N$ in epistemic state $s \in S$ *knows* θ just if θ is true (in the classical logic sense) *in all epistemic states* which are accessible to i from s . Distinct formalizations of “knowledge” (and “belief”) arise from distinct constraints that we impose on the relation R . In the most widely studied (multi-agent) epistemic logic, known as $S5$, R is an equivalence relation, thus making the resulting logical characterization of knowledge effectively equivalent to that provided by Aumann’s structures.

This individual knowledge operator naturally extends to *groups* of agents as follows⁷. Suppose, for definitiveness, that i and j are the only members of group $G \subseteq N$. If both i and j know θ , we say that “ θ is *mutual knowledge* among group G ”. Since this latter can be expressed as a well-formed formula of epistemic logic, the construction can be iterated, so we can have that “ i and j know that θ is *mutual knowledge* among group G ”. This sentence thus expresses *second-level mutual knowledge*, because i knows that j knows that they both know θ , and so on. When the mutual knowledge operator is iterated infinitely many times, we say that θ is *common knowledge* between i and j , a concept introduced⁸ by Lewis (1969).

Over the past two decades or so, the formal analysis of (group) epistemic attitudes has led the momentous development of both formal epistemology (see, e.g. Hendricks, 2005) and epistemic game theory⁹. As to this latter, it is well worth recalling how the rigorous characterization of epistemic atti-

⁵Fagin et al. (1996) collects various works providing the first systematic formalization of the epistemic interaction of logical agents.

⁶Self-contained formal presentations can be found in Fagin et al. (1996) and van Ditmarsch et al. (2007).

⁷Cf. what Aumann dubs “syntactical structures” in Aumann (1999)

⁸For an overview on the notion of common knowledge, see Vanderschraaf and Sillari (2007).

⁹In turn, both areas contributed essentially to the development of the *social software* research programme (see, e.g. Parikh, 2002)

tudes has resulted in a much sharper analysis of the “epistemic conditions” yielding various kinds of solution concepts. A particularly interesting case in point, from the point of view of this paper, is the demotion of the common-knowledge condition for Nash equilibrium put forward by [Aumann and Brandenburger \(1995\)](#), where it is shown that decidedly weaker conditions (some involving mutual knowledge) suffice for the justification of the classical solution concept. The main result reported in this paper, namely that in the weak-link game of our experimental setting *mutual knowledge* leads to the selection of efficient equilibria more frequently than common knowledge, certainly adds experimental support to such a demotion of common knowledge in rational interaction.

In addition to making rigorous the otherwise rather subtle distinction between mutual and common knowledge, the general framework of formal epistemology caters for another conceptual distinction playing a fundamental role in the present investigation, namely the distinction between normative and descriptive accounts of epistemic attitudes. Whilst the former prescribe how idealized agents *should* behave in order to qualify as rational, the latter aims at describing the behavior of actual agents (e.g. our experimental subjects) facing choice problems which involve reasoning about other agents’ epistemic attitudes.

It is commonly accepted that normative models of rationality play a crucial role in helping real decision makers correcting their mistakes. This role of normative models is grounded on the assumption that whilst their cognitive limitations may never allow them to fully attain the recommendations of the model, real decision makers should aim at the standard of rationality captured by the norm. Now, the formal work on multiagent epistemic logics which we have briefly surveyed in this Section undoubtedly suggests that common-knowledge provides the most favourable condition for coordination in strategic interaction. The experimental results discussed in this paper, however, tell another story.

1.2 The experimental

With the present contribution we relate to two connected strands of experimental and behavioral literature¹⁰, namely the extensive experimental litera-

¹⁰For a primer on the use of the experimental method in economics, see ([Friedman et al, 1994](#)) for a critical survey of experiments in game theory see ([Camerer, 2003](#)).

ture on coordination games and the philosophical and experimental literature on social conventions and social norms.

Among the early advocates of the empirical investigation of coordination games is Thomas Schelling, who, at the beginning of the 1960s pointed out how

One cannot, without empirical evidence, deduce what understandings can be perceived in a non-zero sum game of maneuver any more than one can prove, by purely formal deduction, that a particular joke is bound to be funny (Schelling, 1980, p.164)

The power of “focal points” to facilitate coordination, which Schelling himself tested in a series of informal experiments with his students, was first identified empirically by Mehta et al. (1994), who introduced the notion of “Schelling salience” to refer to features - such as action labels or others (see, e.g. Bacharach and Bernasconi, 1997; Cubitt and Sudgen, 2003) - able to confer unambiguous distinctiveness to one specific equilibrium. The power of “Schelling salience” to focus players’ expectations on some unique outcome and hence to solve coordination problems is most evident in so-called games of pure coordination, i.e., games in which all equilibria are payoff-equivalent. As a result, the essence of this kind of coordination problems consists only in avoiding a disequilibrium outcome. A different, but closely related class of coordination games contains games with multiple equilibria, some of which are more rewarding than others for all players. Typically, these games embody a tension between efficiency and security (or risk-dominance), as players choosing strategies that support payoff-superior equilibria incur in greater losses if their choice is not matched by all other players. As a consequence, two types of coordination failure may occur: a disequilibrium outcome and coordination on a suboptimal equilibrium.

Harsanyi and Selten (1988) formalized the tradeoff implicit in these games and stated that rational players, in the presence of common knowledge of rationality, should select the efficient (i.e., payoff-dominant) equilibrium. However, a plethora of experimental studies on the stag hunt and weak-link games (see, e.g. Cooper et al., 1990, 1992; Van Huyck et al., 1990) has convincingly demonstrated that Harsanyi and Selten’s prediction holds true only under very special conditions, while in all remaining cases coordination failure occurs almost invariably. Starting from these results, most of the subsequent experimental studies on coordination have been aimed at identifying the

efficiency-enhancing properties of several features of the strategic interaction at hand. Our study contributes to this strand of literature by identifying the *epistemic conditions* that, in a specific experimental setup highly conducive to coordination failure, render external interventions in the form of third party announcements effective in directing the players' choices towards superior and more risky equilibria.

The second strand of research to which our investigation is related includes the philosophical and experimental literature on social conventions and social norms pioneered by the game theoretically informed work of Lewis (1969) on conventions. David Lewis' account of social convention is based on coordination games. When players' interests are aligned (as in pure coordination games) and interactions are recurrent, players may succeed in coordinating their behavior by selecting a salient equilibrium. If they recurrently succeed in coordinating, the regular solution becomes salient by virtue of precedent, and, according to Lewis's definition, is a *social convention*. When there is a trade-off between security and efficiency (as in the kind of coordination game our experimental work focuses on), payoff considerations make the choice of salient equilibrium particularly relevant. Lewis's account requires that conventions be common knowledge and the importance of such an epistemic requirement, especially with reference to salience, has been a debated issue in the literature (see Section 2.1 below). The present contribution investigates the effectiveness of various levels of epistemic access to the salience-inducing coordination device (i.e. a suggestion to play the strategy supporting the efficient equilibrium), thus contributing to the discussion on the role of common knowledge in coordination and convention.

Lewis's work on social conventions has spun a vast literature in philosophy of social sciences. Particularly relevant are accounts of social norms based, to varying degrees, on Lewis's idea of convention as coordination (cf. for instance Ullman-Margalit, 1977; Sugden, 1986, 2004; Young, 1998; Vanderschraaf, 1995), culminating in Cristina Bicchieri's theory of social norms (Bicchieri, 2006). In Bicchieri's account, a social norm transforms a mixed-motive game into a coordination game: when a social norm is in place and agents hold the relevant expectations about others' preferences, norm-abidance behavior is the preferred course of action while the same course of action would have been dispreferred in the absence of a social norm. The expectations involved in Bicchieri's account of social norms are of two (kinds: (i) *empirical expectations* (first-order beliefs that others will conform) and (ii) *normative expectations* (second-order beliefs that other will expect one

to conform). Extensive experimental work by Bicchieri and colleagues (cf., e.g., [Bicchieri and Xiao, 2009](#); [Bicchieri and Chavez, 2010](#)) show that empirical expectations by themselves may not be sufficient to motivate compliance to social norms, and that the presence of aligned second-order, normative expectations, can trigger compliance. While the present experiment does not deal directly with social norms, it does look at how first and second-order expectations as elicited by different experimental conditions affect coordination outcomes, thus offering insights on the role and nature of coordinating expectations.

2 The game

Arguably the most prominent testbeds of the risk-efficiency tradeoff in the domain of coordination games are the stag hunt and minimum effort (or weak link) games. We limit ourselves to describe the latter¹¹. On-time aircraft departures are prototypical examples of coordination problems of the weak-link type (see, e.g., ([Knez and Simester, 2001](#)),) since the airplane cannot take-off before all operations (e.g., fueling, security checks, loading of luggage, boarding of passengers, etc.) have been completed. Other examples include relationships between different branches of a bank, the writing of a grant proposal involving several participants, an edited volume involving several authors, and many others ([Camerer and Knez, 1996, 1997](#)). Common to all aforementioned examples is the fact that output is determined by the agent exerting the lowest level of effort (the “weak link”), and any effort above the minimum is wasted.

A weak-link game is defined by N players who must simultaneously choose a natural number x in $\{1, 2, \dots, X\}$. The payoff function for the normal form game is defined as follows:

$$\pi(x_i) = a + b * \min(x_1, x_2, \dots, x_n) - c[x_i - \min(x_1, x_2, \dots, x_n)] \quad (1)$$

where x_i is the number chosen by player i , and a , b and c are positive parameters. x_i is intuitively interpreted as the “effort level chosen by player i ”. Provided that the payoff function is common knowledge, the game has X

¹¹In fact the minimum effort game is an extension of the stag hunt game to the case of n strategies with $n > 2$. See, e.g., [Camerer 2003](#), ch. 7.

strict Nash equilibria, corresponding to the X action combinations in which all players select the same effort level. Furthermore, the Nash equilibria can be Pareto-ranked, with the combination $\{x_1 = x_2 = \dots = x_n = X\}$ being the *efficient*, or payoff-dominant equilibrium. Viceversa, the $\{x_1 = x_2 = \dots = x_n = 1\}$ corresponds to the *secure* equilibrium ¹².

In our experiment we chose the following parameter values: $N = 2$, $X = 7$, $a = 6$, $b = c = 1$. For these parameters, the game presents seven Pareto-ranked Nash equilibria, with the combinations $\{x_1 = x_2 = 7\}$ and $\{x_1 = x_2 = 1\}$ representing the payoff-dominant equilibrium and the secure equilibrium, respectively. Each player is penalized the further her choice is from the minimum in the group. The resulting payoff matrix¹³ is shown in Table 1.

		Minumum number chosen						
		1	2	3	4	5	6	7
My choice	1	7	-	-	-	-	-	-
	2	6	8	-	-	-	-	-
	3	5	7	9	-	-	-	-
	4	4	6	8	10	-	-	-
	5	3	5	7	9	11	-	-
	6	2	4	6	8	10	12	-
	7	1	3	5	7	9	11	13

Table 1: Payoff of the weak-link game.

2.1 Related literature

The first comprehensive experiment on the weak link game was conducted by [Van Huyck et al. \(1990\)](#), varying group size, parameter values and matching protocol: unravelling to the lowest minimum over time always occurred with large groups (i.e., $N=9,14$) interacting repeatedly. Pairs did slightly better

¹²“Security” is an extension of Harsanyi and Selten’s risk-dominance principle to games with more than two strategies. A “secure” equilibrium is selected when all players choose in accordance with the maximin criterion.

¹³While in a two-person game the natural choice would that of depicting one’s own choice against the other player’s choice instead of own against minimum of the two, we opted to use the latter representation in order to make our results fully comparable with the previous literature.

but never converged when randomly rematched in each round: only when interacting repeatedly in fixed pairs were players able to coordinate on the payoff-dominant equilibrium. The results in [Van Huyck et al. \(1990\)](#) have been extensively replicated *ceteris paribus*. Taking these findings as point of departure, subsequent studies have investigated different ways to engineer coordination success in the lab: pre-play communication, inter-group competition, increased financial incentives, gradual group growth, and leadership are all mechanisms that have been shown to increase the likelihood of efficient coordination; others, such as increased feedback, have been tested with inconclusive results (see [Devetag and Ortmann \(2007\)](#) for a review of experiments on both the weak link and stag hunt coordination games). Of particular relevance for our study are those experiments that tested the effectiveness of advice, either from a third party or from players themselves. [Van Huyck et al. \(1992\)](#) tested the role of non-binding, publicly announced external assignments in two-person coordination games; their results indicate that assignments are not considered credible when they conflict with payoff-dominance or symmetry. It has to be pointed out, however, that in their games the two principles of payoff-dominance and risk-dominance do not conflict. [Chaudhuri et. al \(2009\)](#) tested intergenerational advice in the minimum effort game with large groups: they found out that the payoff-dominant equilibrium was selected only when advice not only was public and publicly shared (i.e., when all advice from previous generation members was made accessible to all new generation members, and this fact was common knowledge), but was also common knowledge in the sense of being read aloud. The authors call the first condition “almost common knowledge” as in [Rubinstein \(1989\)](#), and in their experiments it leads groups to select the Pareto-worst outcome. Only the common knowledge condition leads to Pareto efficiency. [Bangun et al. \(2006\)](#) replicate the design by [Van Huyck et al. \(1992\)](#) using a stag-hunt game embedding a clear conflict between payoff-dominance and risk-dominance. They show that, unlike the [Van Huyck et al. \(1992\)](#) experiment, an external assignment is sufficient to produce coordination on the payoff-dominant outcome, even when the assignment is “almost common knowledge”. Finally, [Chaudhuri and Paichayontvijit \(2010\)](#) test the power of an external assignment (in the form of a recommendation to choose the strategy consistent with the efficient equilibrium) in the weak link game with groups of five players: their 2x2 between-subject design tests fixed vs. random matching, and “common knowledge” vs. “almost common knowledge” of the assignment. A further treatment tests the power of an increased

bonus for coordination. As in the [Chaudhuri et. al \(2009\)](#) study, “common knowledge” refers to the recommendation being read aloud, while “almost common knowledge” implies that the recommendation is handed out to each subject to be read privately, and the fact that everybody is given the same recommendation is common knowledge. Their findings show behavior that is largely consistent with behavior in ([Chaudhuri et. al , 2009](#)): with random re-matching of groups in every round, only the increased financial incentive (publicly announced) can induce efficient coordination; neither “common knowledge” nor “almost common knowledge” of the recommendation are sufficient; fixed groups are generally able to reach successful coordination even in absence of external interventions. This last finding is at odds with the finding of [Chaudhuri et. al \(2009\)](#), in which players interacting in fixed groups failed to coordinate successfully even in the presence of a credible assignment; the difference is most likely due to the different content of the advice (in the study done by Chaudhuri and colleagues, many players advised their successors to actually select the secure equilibrium), and to the larger group size of the [Chaudhuri et. al \(2009\)](#) design, a parameter which has been shown to be crucial in determining the probability of efficient coordination in the weak-link game.

While in ([Chaudhuri and Paichayontvijit, 2010](#)) a distinction is made between *public* and *common* knowledge of advice (their conditions 1 and 2, respectively), it is clear from the discussion above that the two conditions are both common knowledge conditions. The former condition (in which subjects receive a sheet of paper stating the advice *and* that every other player receives the exact same advice and message) induces first-order knowledge of the advice (because of the text on the paper), second-order knowledge that everyone has knowledge of the advice (because the fact that everyone receives the same text is publicly known), third-order knowledge that everyone knows that everyone has knowledge of the advice (because everyone knows that everyone knows everyone has received the same text), and so on. Of course, subjects could be doubtful about whether others have read the message, paid attention to all of its contents, etc., hence the notion that this is an *almost* common knowledge condition.

None of the previous studies on the efficiency-enhancing properties of external advice has tested epistemic conditions weaker than common or at least almost common knowledge. Our paper fills this gap in the literature, examining, along with common knowledge, first- and second-order mutual knowledge of advice. First-order mutual knowledge is obtained by an ex-

perimental setup that induces knowledge of advice in each player, without offering subjects the opportunity to draw inferences beyond first-order, i.e. the experimenter does not give them enough information to infer that others have knowledge of advice. Second-order mutual knowledge is similarly obtained by a setup in which knowledge that one's opponent has knowledge of advice can be inferred, but knowing that the opponent knows that one knows that cannot.

Our experimental work is linked to a further strand of literature, namely that on game-theoretic accounts of social conventions. Cubitt and Sugden (in [Cubitt and Sugden \(2003\)](#)) have advanced a formal reconstruction of Lewis's account of social convention that stresses the importance of inductive reasoning as based on common knowledge of precedent. On the other hand, evolutionary accounts of social conventions (cf. [Binmore \(2008\)](#), [Skyrms \(1996\)](#)) deny the relevance of common knowledge for sustaining a social convention (cf. [Sillari \(2008\)](#) and [Rescorla \(2011\)](#) for overviews of the debate.) Our experimental study provides a contribution to this debate, as our experimental conditions allow us to observe the impact of different kinds of group knowledge on the formation of conventions. To be sure, our experiment does not deal with convention formation, as iterations of the interactive choice are numerous enough to speak of an emerging convention. However, our experiment does yield insights on a crucial mechanism of convention formation, i.e. that of salience building. In Lewis's account of convention as it is discussed in the recent debate, a specific kind of salience (precedent) is at work in perpetuating a convention. Instead, we focus on plausible mechanisms other than precedent for the emergence of salience, in particular we observe how subjects react to advice. Notice that the experimenters' advice does not add any factual knowledge to what subjects are supposed to (commonly) know, i.e. the structure of the game and payoffs. Rather than providing knowledge, the suggestion received by the subjects makes them aware of a relevant item of knowledge or, in other words, makes such an item salient to them. Assuming, with Lewis, that subjects coordinate by virtue of salience, our experiment then discriminates the effectiveness of salience-building mechanisms in (weak-link) coordination games.

3 Experimental Design

Ninety participants were recruited from Luiss Guido Carli in Rome. Participants were undergraduate students enrolled in the various programs of the University, economics, law and political science. All experiments were conducted at the CESARE experimental economics Lab at Luiss, by using a dedicated software. The payoff function was the one shown in Table 1 above.

Most sessions were conducted with 10 participants, with a few containing 8 participants. Each session consisted of a total of 13 rounds of play of the stage game. Participants were randomly assigned to computer terminals isolated from one another by glass separators. Initial instructions (equal for all treatments) explaining general rules were distributed in paper copy and were also read aloud by the experimenter. At this stage, subjects were instructed that the experiment was divided in two parts, and that further instructions would appear on everybody's monitor as the experiment progressed. Control questions were administered to assure that the rules of the game had been correctly understood by all. A paper copy of the payoff matrix was given to each subject, and the same matrix also appeared on each participant's screen when submitting their choice for the round. Subjects were told that they would be randomly paired with another anonymous participant from the room in each round. When everybody had entered their choice, each subject received private feedback regarding chosen number, minimum number chosen in the pair and payoff earned in the round. In order to enhance independence between rounds, the final payment was based on a randomly selected round, and this was explicitly stated in the instructions.

The first three rounds were equal for all treatments and were played with no external intervention, to increase comparability with previous experiments (Chaudhuri and Paichayontvijit, 2010) and to allow adequate learning of the payoff function. In fact, we wanted to avoid subjects following the experimenter's suggestion simply as a result of inexperience and being unaware of all its possible implications.

After the third round had been completed and before subjects could enter their choice for the fourth round, the experiment in all treatments was interrupted and subjects were asked to state the following two forecasts on a sheet of paper:

- the average number chosen in period 4 by all participants in the session
- the average forecast expressed by all participants in the session

Sincere belief elicitation was incentivized in the following way: for each forecast, subjects earned 1 ECU minus the difference between the stated average and the actual average, up to 0. After all participants had stated their beliefs, the experiment was resumed, and subjects could make their choices for round 4 and subsequent rounds.

The following four treatments were conducted. In the *control* treatment (2 sessions, $N = 16$), subjects were asked to state their beliefs at the end of the third round, and then asked to proceed with making their choices for round 4 and the remaining rounds, with no further instructions.

In the *mutual knowledge of level 1* treatment (3 sessions, $N = 30$), before subjects were asked to make their predictions and choice for round 4 and at the beginning of all subsequent rounds of play, the following statement appeared on every participant's monitor:

WE SUGGEST YOU TO PICK NUMBER 7
IN FACT, NOTE THAT, IF BOTH YOU AND THE OTHER PARTICIPANT PICK
7, YOU BOTH EARN THE MAXIMUM POSSIBLE PAYOFF, 13 ECUS

The fact that the same suggestion appeared on every participant's monitor was *not* common knowledge¹⁴. Subjects had to click on "continue" to move from the suggestion screen to the choice screen, to assure that the suggestion was not overlooked.

In the *mutual knowledge of level 2* treatment (3 sessions, $N = 26$), before subjects were asked to make their forecasts and choice for period 4, and at the beginning of all subsequent rounds of play, the same statement shown above appeared on every participant's monitor. In addition, the lower part of the screen reported the following:

CLICK ON "SEND" TO SEND THIS SUGGESTION TO THE
PARTICIPANT WITH WHOM YOU ARE PAIRED IN THIS ROUND.

Subjects had to click "send" to move to the following screen. The following screen reported the sentence shown below:

THE PARTICIPANT WITH WHOM YOU ARE PAIRED IN THIS
ROUND IS SENDING YOU THIS MESSAGE:

¹⁴This design feature marks an important difference in our setup with respect to that in (Chaudhuri and Paichayontvijit, 2010).

followed by the suggestion text.

The subject was asked to click “read” to send the other player a confirmation that the message had been read before moving to the next screen. The last screen reported:

THE PARTICIPANT WITH WHOM YOU ARE PAIRED NOTIFIES
YOU THAT THEY HAVE RECEIVED THE MESSAGE THAT YOU
SENT THEM.

Then the choice screen appeared. Neither the suggestion appearance nor the message exchange was common knowledge among participants¹⁵.

Finally, in the *common knowledge* treatment (2 sessions, $N = 18$), the suggestion text was read aloud by the experimenter before subjects were asked to state their forecast for period 4. The suggestion was read aloud at the beginning of round 4 only, and not at the beginning of every subsequent round. Figures 1 and 2 illustrate two samples of the computer interface used in the experiment.

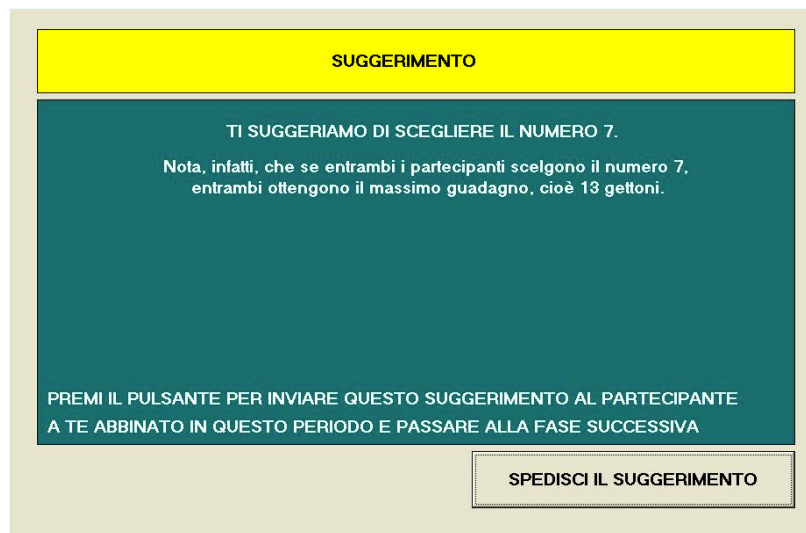


Figure 1: The suggestion screen

¹⁵Notice that while this protocol can legitimately generate knowledge of the suggestion up to level 2, it also legitimately generates a third-order *belief*, as each player can believe (but not know) that the other player has seen the acknowledgment sent by the former player.

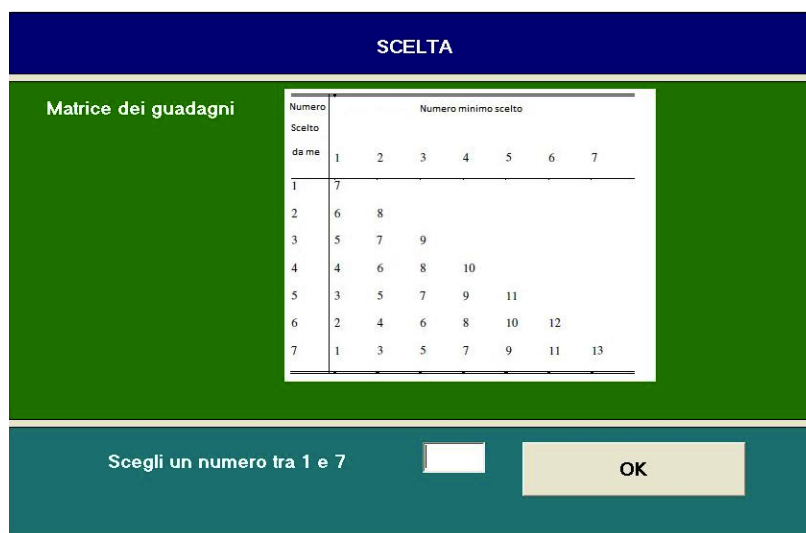


Figure 2: The choice screen

Before moving on to stating our hypotheses and analyzing the data, we would like to address the issue of what in fact is the *content* of the proposition on which the present experiment hinges. The suggestion that the experimenters send to the subjects contains two elements: one is the suggestion itself (“we suggest you pick number 7”) and the other is a factual explanation of the rationale behind the suggestion (“note that if both you and the other participant pick 7, you both earn the maximum possible payoff, 13 ECU’s’.) As the message read by subjects contains a factual explanation of the rationale behind the suggestion, it is reasonable to believe that subjects are led to an *indicative* reading of the suggestion, so that the content of, e.g., the proposition commonly known in the third treatments is “experimenters suggest to pick 7,” or perhaps “experimenters suggest that I pick 7.”¹⁶

To facilitate the exposition of the results, we will refer to our four treatments according to the following stipulation:

- *control* - C = control
- *mutual knowledge* - MK = mutual knowledge of level 1

¹⁶The indication of a rationale for the suggestion is thus important to avoid that subjects understand the suggestion in imperative terms, as for instance in “experimenters suggest that I should pick 7.” We are grateful to an anonymous referee for prompting us to make this clarification.

- *message* - *Message* = mutual knowledge of level 2
- *common knowledge* - *CK* = common knowledge.

The control treatment is necessary to provide us with a benchmark for the frequency of coordination, as well as to control for learning and restart effects. The other three treatments allow us to test the effect that adding one level of knowledge has on the probability of coordination on the efficient equilibrium.

The choice to use the payoff function of (Van Huyck et al., 1990) is motivated by the fact that this payoff function has already been tested for a group size of two players with random matching. Results have shown high variability and no convergence over time in this condition. To our knowledge, this is the worst coordination outcome that has been reported in a minimum effort game with group size of two. As a high level of coordination failure in our baseline treatment is a necessary condition for our experimental treatments to be of interest, we chose to implement the same payoff function used by Van Huyck and co-authors.

Experiments lasted forty minutes on average, including instruction time. Average total earnings were equal to 13.5 euro, including the 2 euro show up fee.

We formulated the following hypotheses:

- coordination failure is the modal outcome in the control treatment
- both the share of choices equal to 7 and the frequency of efficient coordination outcomes increase as more knowledge levels are added

Hence, we expected to observe widespread coordination failure in the control treatment, and higher frequencies of coordination success as moving from the *control* to the *common knowledge* treatment.

4 Results

4.1 Analysis of choice behavior

Figure 3 reports the average choices, whilst average minima and relative frequency of choices equal to 7 over time in the four treatments are reported in Figure 4 and Figure 5, respectively. All values are pooled across sessions.

The following facts are worth noticing: first, both minima and 7 choices in the control treatment indicate widespread coordination failure with no detectable trend over time. Minima stabilize around 4 and choices of 7 around 40 per cent. Hence, the control treatment by and large replicates previous findings obtained with the same group size and matching protocol in (Van Huyck et al., 1990). Secondly, all experimental treatments reveal an improvement in coordination with respect to the control, as hypothesized.

However, somewhat surprisingly and contrary to our hypothesis, the increase is not monotonic in knowledge levels. More specifically, the highest values are attained in the message treatment, while the values for the mutual knowledge and common knowledge treatments lie between the two extremes of “control” and “message”.




Figure 3: Average choices over time, all sessions pooled

Figure 6 reports the frequency of the choice of 7 in rounds 4-13 separately for each treatment, and in rounds 1-3 for all treatments combined. Figure 7 reports the same information, respectively, for round 4 only and for round 1 only.

Figure 6 clearly shows that, overall, the message treatment was by far the most effective in fostering choice of the strategy supporting the efficient equilibrium. The second finding worth noticing is that mutual knowledge of




Figure 4: Average minima over time, all sessions pooled




Figure 5: Relative frequency of choices equal to 7 over time, all sessions pooled

level 1 is able to induce the choice of 7 in a proportion comparable to that attained in the common knowledge treatment.

By comparing the same data in round 4 only, as reported in Figure 7, we can measure the effect of different levels of knowledge excluding the potentially confounding factor given by message repetition. Moreover, the salience

of the publicly read announcement in the CK treatment is supposed to be highest in round 4 as compared to subsequent rounds. Hence, the CK treatment is given its best chance to produce efficient coordination immediately after the public announcement. Data from round 4 suggest that the message treatment is the most effective even without the aid of message repetition; in fact, choices of 7 amount to 80 per cent of all choices in that treatment. The CK treatment is the second best, with 60 per cent of 7 choices. MK, however, does almost as well as the CK treatment, whereas the percentage in the control treatment does not differ from that observed in the first three rounds of play.




Figure 6: Frequency of choices equal to 7 in rounds 4-13, disaggregated by treatment, and in rounds 1-3, all treatments pooled

The relative ranking of the four different knowledge levels on goodness of coordination is further supported by data on equilibrium play. Figure 8 reports the percentage of play consistent with the efficient equilibrium (i.e., both players in the pair picking 7) in the four treatments in rounds 4-13. In the message treatment more than half of play (52 per cent) resulted in coordination on the efficient equilibrium, whereas data for the CK treatment and for the MK treatment amount to 28.8 per cent and 27.3 per cent, respectively. Hence, again, mutual knowledge of level 1 is comparable to common knowledge in its efficiency-enhancing effects, and mutual knowledge of level

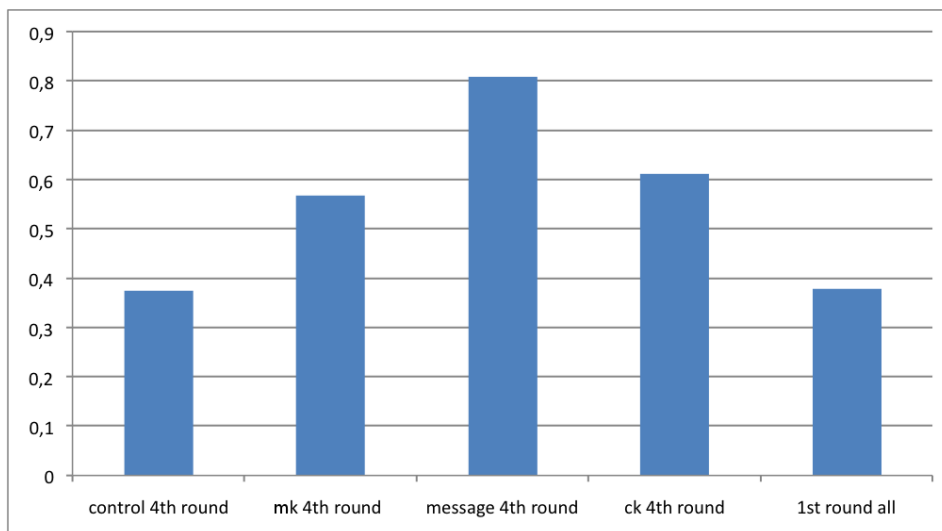


Figure 7: Frequency of choices equal to 7 in round 4, disaggregated by treatment, and in round 1, all treatments pooled.

2 fares substantially better than common knowledge. In accordance with our hypothesis and with previous results, the control treatment results in the highest inefficiency, with only 13.8 per cent of efficient equilibrium play. Consistently, data on the standard deviation of choices in the four treatments in round 4-13 (Figure 9) reveal that the message treatment reduces strategic uncertainty the most with respect to the control, while the MK and CK treatment values are indistinguishable.

Pairwise comparisons between choices in the first round in the different treatments reveal that there are no statistically significant differences in choice distributions in period 1 (Mann-Whitney U test; control-mutual knowledge $p = .65$; control-message $p = .499$; control-common knowledge $p = .95$; mutual knowledge-message $p = .83$; mutual knowledge-common knowledge $p = .51$; message-common knowledge $p = .44$. All p -values are two-tailed).

Table 2 reports results of the Fisher's exact test on the differences in the share of 7 choices in rounds 4-13, and Table 3 reports results of the same test on the frequency of efficient equilibrium play. All differences are highly significant *except* that between MK and CK.




Figure 8: Percentage of equilibrium play in rounds 4-13

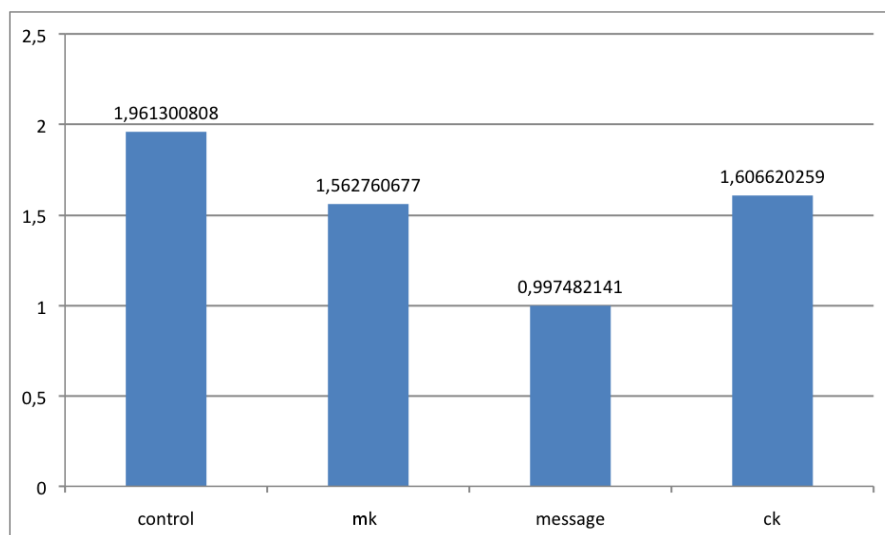


Figure 9: Standard deviation of choices in rounds 4-13, disaggregated by treatment

4.2 Analysis of beliefs

Figure 10 shows the distribution of first order beliefs (i.e., subjects' forecasts about the average of all participants' choices in round 4) by treatment. It

	Control	Mutual knowledge	Message	Common knowledge
Control		$p = .0061$	$p = .0001$	$p = .0047$
Mutual knowledge			$p = .0001$	$p = .7057$
Message				$p = .0008$
Common knowledge				

Table 2: Fisher’s exact test on the frequency of choices 7 in the different treatments in rounds 4-13 (all p -values are two-tailed).

	Control	Mutual knowledge	Message	Common knowledge
Control		$p = .0010$	$p = .0001$	$p = .0009$
Mutual knowledge			$p = .0001$	$p = .7528$
Message				$p = .0001$
Common knowledge				

Table 3: Fisher’s exact test on the frequency of equilibrium play in the different treatments in rounds 4-13 (all p -values are two-tailed).

is noteworthy that beliefs behavior is highly consistent with choice behavior, in that a shift of the distribution in the direction of higher numbers can be detected when moving from the control treatment to the MK, CK and message treatments. In the message treatment, roughly 65 per cent of forecasts are concentrated on the value 7.

Table 4 reports some descriptive statistics on first order, second order beliefs and choices in round 4 by treatment¹⁷.

High consistency between beliefs and choices is revealed by looking at the column of mean values. Data on beliefs reveal important insights: first, they support the hypothesis that the suggestion (and its related knowledge levels) rendered subjects on average more optimistic, and greater optimism in turn induced the choice of the risky action significantly more often with respect to the absence of an external suggestion¹⁸. In other words, an experimenter

¹⁷One subject in the control treatment stated a first order belief equal to 10, one subject in the message treatment stated a second order belief equal to 42, and one subject in the common knowledge treatment stated a second order belief equal to 8. We eliminated these outliers from the computation of the average values and of the subsequent correlation indexes. The results of the correlation do not change if we include the outliers. We did not observe any anomalous behavior in these subjects at the choice level, hence they may have not understood the prediction task correctly.

¹⁸Notice also that, since the elicited beliefs are beliefs in the average choice of action of all other participants in the session, choosing an action higher than one’s expressed prediction denotes optimist risk-seeking behavior.

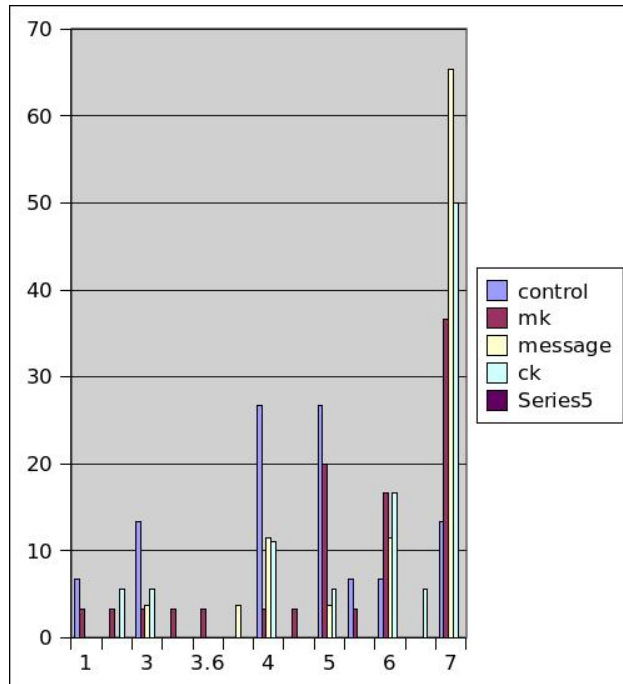


Figure 10: Distribution of beliefs by treatment

	N	Minimum	Maximum	Mean	Std. Deviation
1st order belief CONTROL	15	1	7	4,57	1,568
2nd order belief CONTROL	16	1	7	4,31	1,740
Choice in round 4 CONTROL	16	1	7	5,13	1,962
1st order belief MK	30	1	7	5,47	1,630
2nd order belief MK	30	1	7	5,58	1,602
Choice in round MK	30	1	7	5,73	1,999
1st order belief MESSAGE	26	3	7	6,18	1,311
2nd order belief MESSAGE	25	4	7	6,37	1,026
Choice in round MESSAGE	26	3	7	6,58	1,027
1st order belief CK	18	2	7	5,88	1,597
2nd order belief CK	17	4	7	6,00	1,225
Choice in round CK	18	2	7	5,83	1,724

Table 4: Descriptive statistics on first order, second order beliefs and choices in round 4 by treatment

demand effect ([Zizzo, 2010](#)) as an explanation of behavior in our experimental treatments can be safely excluded. Rather, the most plausible explanation of

our subjects' behavior involves increased optimism and trust. Moreover, data on beliefs strongly point at the fact that subjects in the MK and message treatments believed that other participants in the session received the same type of information that they themselves received.

The positive effect of *any* knowledge level can be detected by looking at minimum values in column 3. Unlike the control treatment, where a minimum of 1 is observed both at the belief and at the choice level, in all experimental treatments the external announcement is able to eliminate the secure action (i.e., the choice of 1) from both participants choices and forecasts. Finally, Table 5 reports the correlation coefficients between first order beliefs, second order beliefs and choices in round 4 separately for each treatment. As the table shows, all relevant correlation coefficients are positive and highly significant.

		1st order pred. C	2nd order pred. C	4th round choice C
1st order pred. C	Corr. Coeff.	1,00	0,86	0,61
	Sig.		0,00	0,02
	N	15,00	15,00	15,00
2nd order pred. C	Corr. Coeff.	0,86	1,00	0,61
	Sig.	0,00		0,01
	N	15,00	16,00	16,00
		1st order pred. MK	2nd order pred. MK	4th round choice MK
1st order pred. MK	Corr. Coeff.	1,00	0,80	0,38
	Sig.	.	0,00	0,04
	N	30,00	30,00	30,00
2nd order pred. MK	Corr. Coeff.	0,80	1,00	0,59
	Sig.	0,00	.	0,00
	N	30,00	30,00	30,00
		1st order pred. M	2nd order pred. M	4th round choice M
1st order pred. M	Corr. Coeff.	1,00	0,92	0,69
	Sig.		0,00	0,00
	N	26,00	25,00	26,00
2nd order pred. M	Corr. Coeff.	0,92	1,00	0,51
	Sig.	0,00	.	0,01
	N	25,00	25,00	25,00
		1st order pred. CK	2nd order pred. CK	4th round choice CK
1st order pred. CK	Corr. Coeff.	1,00	0,74	0,76
	Sig.	.	0,00	0,00
	N	18,00	17,00	18,00
2nd order pred. CK	Corr. Coeff.	0,74	1,00	0,72
	Sig.	0,00		0,00
	N	17,00	17,00	17,00

Table 5: Spearman's rho correlation coefficients (with relative p -values, two-tailed) between subjects' 1st order, 2nd order predictions, and 4th round choices, divided by treatment. The relevant significance values are reported in bold face.

5 Refined *message* treatment

The rationale for the next treatment is to assess the potential impact of the change in propositional content caused by the linguistic form of the *message* treatment¹⁹. More specifically, the act of sending a message to the other player, combined with the ambiguous message content (“*We* suggest you to pick 7...”) might have created normative pressure that is absent in the other two conditions, deriving from a possible interpretation of the suggestion as coming from the other player as well as from the experimenter. To address this issue, we designed and conducted a further “level-2” treatment (MK2 henceforth) in which any possible source of normative pressure was eliminated. In this further treatment, the message exchange phase was left unaltered, but the whole procedure was rephrased so as to make it explicit that the suggestion was only from the experimenter, and even the minimal interaction between players was eliminated.

Hence, the subject’s monitor showed the experimenter’s suggestion, as in the previous treatment. The lower part of the screen reported now the following sentence:

“CLICK HERE TO HAVE THE EXPERIMENTER’S SUGGESTION
VISUALIZED ON THE SCREEN OF THE PARTICIPANT WITH
WHOM YOU ARE PAIRED IN THIS ROUND”

Subjects had to click to move to the following screen, which reported the sentence shown below:

THE PARTICIPANT WITH WHOM YOU ARE PAIRED IN THIS
ROUND IS FORWARDING YOU THIS SUGGESTION FROM THE
EXPERIMENTER

followed by the suggestion text.

The subject was asked to click “read” to send the other player a confirmation. The last screen reported:

WE CONFIRM THAT THE EXPERIMENTER’S SUGGESTION HAS
BEEN VISUALIZED ON THE SCREEN OF THE PARTICIPANT
WITH WHOM YOU ARE PAIRED

¹⁹We most warmly thank one anonymous referee for pointing out this.

Then the choice screen appeared. As before, neither the suggestion appearance nor the message exchange were common knowledge among participants.

Three sessions of the new treatment were conducted (two sessions with 8 players and one session with 10 players) for a total of 18 subjects. Figure 11 shows average choices over time for the five treatments combined, Figure 13 reports average minima and, finally, figures 14 report the frequency of “7” choices and 15 that of equilibrium play.




Figure 11: Average choices over time for the five treatments combined

Pairwise comparisons between choice distributions in the first round reveal, as before, no statistically significant differences (Mann-Whitney U test; control-MK2 $p = .83$; mutual knowledge-MK2 $p = .73$; message-MK2 $p = .58$; common knowledge-MK2 $p = .67$.) Instead, the share of observed “7” choices in MK2 in rounds 4-13 is significantly different from both the control ($p = .006$) and the message treatment ($p = .0001$) by a Fisher exact test. However, it is not different from the mutual knowledge ($p = .9$) and from the common knowledge treatment ($p = .77$). Restricting the analysis to the share of “7” choices in round 4 only (i.e., immediately after the announcement is publicly made in the common knowledge treatment), the only significant difference found is the one between MK2 and the control treatment ($p = .05$, Fisher’s exact test, one-tailed).

Finally, the last comparison concerns the percentage of equilibrium play: all differences are highly significant by the same Fisher exact test (control-




Figure 12: Average minima

MK2 $p = .000$; mutual knowledge-MK2 $p = .02$; message-MK2 $p = .000$ *except* that between MK2 and the common knowledge treatment ($p = .3$).

6 Conclusions and Further Work

To summarize, our results show that the “control” and “message” treatments envelope all others along the most important dimensions (average, minima, percentage of choices of 7, percentage of equilibrium play.) Thus, our experiment reveals two major findings: first, while group knowledge of advice in weak-link games does increase efficient play and coordination, levels of iteration of knowledge of advice do not matter with regards to the amount of the increase (in particular common knowledge does not significantly improve over other kinds of group knowledge); second, second-order mutual knowledge of advice has significantly different effects, as a fostering coordination device, depending on the specific mode with which subjects form expectations about the behavior of their fellow players.

Intuition, sometimes backed by formal results as in the analysis of the so-called *coordinated attack problem* (Fagin et al., 1996), has long been suggesting that the rational solution of coordination games required common-knowledge of the players’ rationality and choices.²⁰ Our results to the effect

²⁰See however Parikh (2005) for an extensive discussion on the far reaching consequence of finite, indeed *small*, level knowledge and Hosni and Paris (2005) for an analysis of




Figure 13: Relative frequency of “7”

that common knowledge of advice does no better than mutual knowledge as a salience-generating device appears to be quite surprising and meaningful nonetheless. For instance, in the debate on whether conventions need to be common knowledge²¹ weighs in favor of the position that they need not.

Our attempt at explaining this unexpected result is best framed in terms of the correlated unexpected finding, namely that second-order mutual knowledge of advice may lead to the selection of efficient equilibria more frequently than any other epistemic condition we investigated.

Why is our second-order mutual knowledge “message” condition so effective at promoting efficient coordination? To identify a plausible explanation, it is convenient to focus on the differences between the second-order mutual knowledge “message” condition and all other conditions in our experiment. To induce the second-order expectation, subjects are instructed to exchange messages in order to continue the experiment. Such a direct, first-person communication exchange is only present in the second-order “message” condition, therefore we surmise that it is an important factor in the explanation of our results. Indeed, computer-mediated message exchange is known to give rise to expectations of mutual commitment between players²², both in

coordination which dispenses with common knowledge altogether.

²¹See for instance [Binmore \(2008\)](#) and [Sillari \(2008\)](#).

²²For studies on the normative impact of (even very limited) computer-mediated communication in mixed-motive games see [Bicchieri and Lev-On \(2007\)](#) and [Bicchieri, Chavez and Lev-On \(2010\)](#), among others.

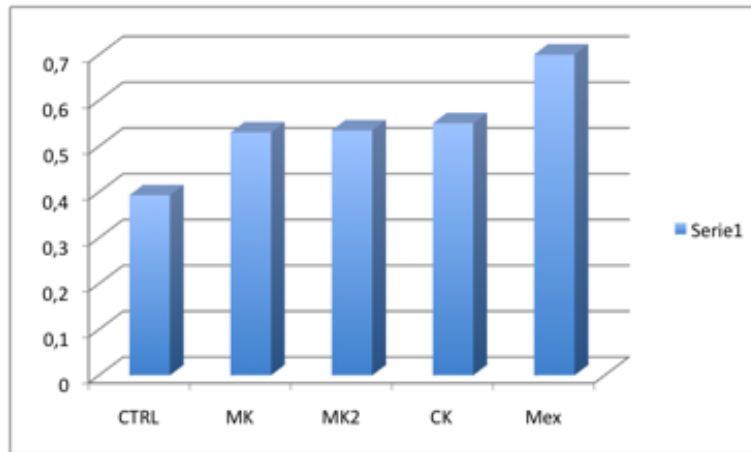


Figure 14: Frequency of “7”

the sense that receiving the message makes the receiver believe that the sender will do her part to honor the advice, and in the sense that sending the message makes the sender believe that the receiver will expect that the sender ought to conform to the advice she sent. In short, first-person message exchange may elicit a norm of promise-keeping.

Thus, not only it is relevant what kind of beliefs players hold, but also how they arrived at holding such beliefs. In particular, beliefs formed through the embryonic communication used in our second-order “message” condition seem to foster compliance with the content of the exchanged message. Recall how in Bicchieri’s theory of social norms second-order expectations play a crucial role as *normative* expectations. While first-order, empirical expectations may not suffice to motivate compliance with the social norm, first- and second-order expectations, when aligned, may succeed in bringing about norm compliance. Our study is not concerned with social norms, however the coordination success observed in our second-order mutual knowledge condition, paired with the observation of high correlation between first-, second-order beliefs, and choices, highlights the relevance of second-order expectations for motivating action. In light of this observation, our hypothesis that subjects’ second-order beliefs have normative content calls for further research.

Our result that second-order mutual knowledge works better than first-order in fostering efficient equilibrium choices may lead to the generalization

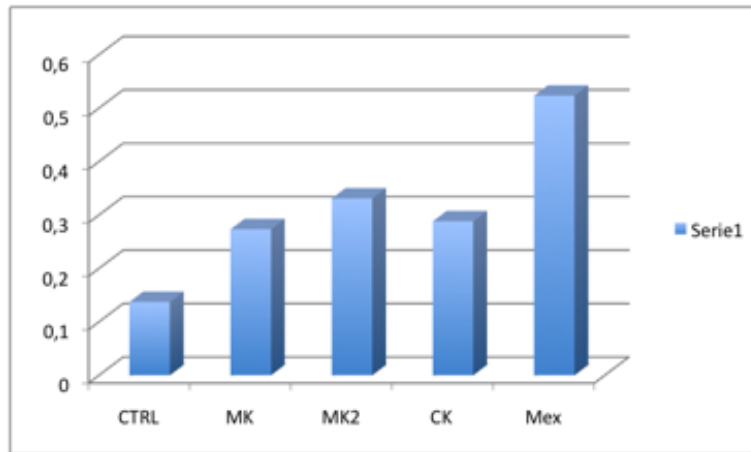


Figure 15: Frequency of equilibrium plays

that there is a monotonic relationship between players' abidance to advice and levels of mutual knowledge of advice. If that were the case, common knowledge of advice would result in highest level of efficient coordination. However, that is not what is observed in our common knowledge condition, lending support to the hypothesis formulated above that second-order expectations carry more weight because of their normative character. Our results, however, suggest not only that common knowledge of advice is not necessary to achieve efficient coordination, but also that it is not sufficient. This result bears on the literature on the relation between convention and common knowledge, showing that common knowledge is not the most effective device for creating salience and fostering efficient coordination, at least in the short run.

We would like to close this paper by insisting on our finding that coordination is more likely under mutual knowledge than common knowledge of advice. As we noted before the surprise might be accounted for by assuming a psychological relevance of the suggestion. If true this might suggest that belief formation mechanism weights more than the sole epistemic condition (common knowledge) towards the selection of the efficient equilibrium. This interpretation seems to be posing a challenge for bayesian epistemology, insofar as the latter focusses only on the information revealed by the agent's (degrees of) belief. Yet the challenge is only apparent, for the surprising performance of mutual knowledge in the message treatment might well depend

on the subjects' "updating" on the expectation of mutual commitment. This emerging notion of belief update, however, clearly deserves further investigation.

Acknowledgments

Financial support from the Italian Ministry of Education, University and Research (MIUR, project PRIN 2008) is gratefully acknowledged. We thank Daniela Di Cagno, Massimo Egidi, Marco Tecilla, Mariagrazia Palmieri and Dearn Sulonjaku for their help in developing the experimental software and in implementing and conducting the experiments. We thank Cristina Bicchieri, Ryan Muldoon, the participants of the Pittsburgh-Tilburg Workshop on *Formal Epistemology Meets Experimental Philosophy* and particularly two anonymous referees for their helpful comments and suggestions. The usual disclaimer applies.

References

- P. Allo. The logic of being informed revisited and revised. *Philosophical Studies*, 153(3):417–434, February 2011.
- R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.
- R. J. Aumann and Adam Brandenburger. Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63(5):1161–1180 1995.
- R. J. Aumann. Correlated Equilibrium as an Expression of Bayesian Rationality. *The Econometric Society*, 55(1):1–18, 1987.
- R. J. Aumann. Interactive Epistemology I: Knowledge *International Journal of Game Theory*, 28:263-300, 1999.
- M. Bacharach and M. Bernasconi, The Variable Frame Theory of Focal Points: An Experimental Study, *Games and Economic Behavior*, 19(1), 1-45, 1997.
- L. Bangun, A. Chaudhuri, P. Prak and C. Zhou, Common and Almost Common Knowledge of Credible Assignments in a Coordination Game. *Economics Bulletin* 3, 1-10, 2006.

- C. Bicchieri. *The Grammar of Society*. Cambridge University Press, 2006.
- Cristina Bicchieri and Alex Chavez. Behaving as expected: Public information and fairness norms *Journal of Behavioral Decision Making* 23, 161-178, 2010
- Cristina Bicchieri and Erte Xiao. Do the right thing: But only if others do so *Journal of Behavioral Decision Making* 22, 191-208, 2009
- Ken Binmore. Do Conventions Need to Be Common Knowledge? *Topoi*, 27:17-27, 2008.
- P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, June 2001.
- P. Blackburn, J. Van Benthem, and F. Wolter. *Handbook of modal logic*. Elsevier, 2007.
- R. J Aumann and Adam Brandenburger. Epistemic Game Theory: An Overview. in *The New Palgrave Dictionary of Economics*, 2nd edition, S. Durlauf and L. Blume (eds.), Palgrave Macmillan, 2008.
- C. Camerer, *Behavioral Game Theory. Experiments in Strategic Interaction*. Princeton University Press, 2003
- C. Camerer and M. Knez, Coordination, Organizational Boundaries and Fads in Business Practice, *Industrial and Corporate Change*, 5, pp. 89-112, 1996.
- C. Camerer and M. Knez, Coordination in Organizations: A Game-theoretic Perspective, in Z. Shapira (Ed.), *Organizational decision making*, Cambridge: Cambridge Series on Judgment and Decision Making, pp. 158-188, 1997.
- A. Chaudhuri, Ananish, B. Sopher, and A. Schotter. Talking Ourselves to Efficiency: Coordination in Inter-Generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice *Economic Journal* 119, N. 534, 91-122, 2009.
- A. Chaudhuri, and T. Paichayontvijit, Recommended play and performance bonuses in the minimum effort coordination game *Experimental Economics* 13, 346-363, 2010.

- R. Cooper, D. De Jong, R. Forsythe, and T. Ross, Selection Criteria in Coordination Games: Some Experimental Results *American Economic Review*, 80, 218-233, 1990.
- R. Cooper, D. De Jong, R. Forsythe, and T. Ross, Communication in Coordination Games. *Quarterly Journal of Economics*, 107, 739-771, 1992.
- R.P. Cubitt and R. Sugden, Common knowledge, salience and convention: a reconstruction of David Lewis's game theory, *Economics and Philosophy* 19, 175-210, 2003.
- M. D'Agostino. Tractable depth-bounded logics and the problem of logical omniscience. In H Hosni and F Montagna, editors, *Probability, Uncertainty and Rationality*, pages 245–275. Edizioni della Normale, 2010.
- C. Bicchieri and A. Lev-On. Computer-mediated communication and cooperation in social dilemmas. *Philosophy, Politics and Economics*, 6, 139-168, 2007
- C. Bicchieri, A. Chavez and A. Lev-On. The medium or the message? Communication relevance and richness in trust games. *Synthese*, 176, 125-147, 2010.
- M. D'Agostino and L. Floridi. The enduring scandal of deduction. *Synthese*, 167(2), 2009.
- G. Devetag and A. Ortmann, When and Why? A Critical Survey on Coordination Failure in the Laboratory, *Experimental Economics*, 10, 331-344, 2007
- R. Fagin, J.Y. Halpern, M. Yoram, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1996.
- Friedman, Daniel, Sunder and Shyam *Experimental Methods: A Primer for Economists*, Cambridge University Press, Cambridge. 1994
- J.Y. Halpern and Leandro Rêgo. Interactive unawareness revisited. *Games and Economic Behavior*, 62(1), 232-262, 2008
- J.Y. Halpern and R. Pucella. Dealing with logical omniscience. In *Proceedings of the 11th conference on Theoretical Aspects of Rationality and Knowledge*. ACM, 2007.

- J. Harsanyi, and R. Selten. *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: The MIT Press.
- Vincent F. Hendricks. *Mainstream and Formal Epistemology*. Cambridge University Press, Cambridge, 2005.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- Hykel Hosni and Jeff Paris. Rationality As Conformity. *Synthese*, 144(2): 249–285, 2005.
- Knez and Simester Firm-wide incentives and mutual monitoring at Continental Airlines, *Journal of Labor Economics*, 19, pp. 743-772, 2001.
- D.K. Lewis. *Convention: A philosophical study*. Harvard University Press, 1969.
- J. Metha, C. Starmer, and R. Sugden. The Nature of Salience: An Experimental Investigation of Pure Coordination Games. *American Economic Review*, 84, pp. 658-673, 1994
- J.-J.Ch Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press, 1995.
- R. Parikh. Social Software. *Syn2*, 132:187–21, 2002.
- R. Parikh. Logical omniscience and common knowledge: WHAT do we know and what do WE. *on Theoretical aspects of rationality and knowledge*, pages 62–77, 2005.
- Michael Rescorla Convention *The Stanford Encyclopedia of Philosophy (Spring 2011 Edition)* Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2011/entries/convention/>> Last checked: January 2012
- A. Rubinstein. The Electronic Mail Game: Strategic Behavior under “Almost Common Knowledge”, *American Economic Review*, 79, 385-391, 1989
- T. Schelling. *The Strategy of Conflict* (second edition), Harvard University Press, Cambridge, MA, 1980

- Burkhard Schipper. Unawareness Bibliography, Burkhard Schipper (ed.), URL = <<http://www.econ.ucdavis.edu/faculty/schipper/unaw.htm>> Last checked: January 2012
- Brian Skyrms. *The Evolution of the Social Contract*, Cambridge University Press, 1996.
- Giacomo Sillari. Common Knowledge and Convention *Topoi*, 27:29-39, 2008.
- Robert Sugden *The Economics of Right, Co-operation and Welfare* (second edition) Palgrave, Macmillan, 2004 (1986)
- H. van Ditmarsch, W. van Der Hoek, and B. Kooi. *Dynamic epistemic logic*. Springer, 2007.
- Johan van Benthem. Logic and reasoning: Do the facts matter? *Studia Logica*, 88:67–84, 2008.
- J.B. Van Huyck, R.O. Beil, A. Gillette and R.C. Battalio. Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure. *The American Economic Review*, 80, 234-248, 1990.
- J.B. Van Huyck, A. Gillette and R.C. Battalio. Credible Assignments in Coordination Games. *Games and Economic Behavior*, 4, 606-626, 1992.
- Edna Ullman-Margalit. *The Emergence of Norms* Clarendon Press, 1977
- Peter Vanderschraaf. Convention As Correlated Equilibrium *Erkenntnis*, 42:65-87, 1995
- Peter Vanderschraaf and Giacomo Sillari. Common Knowledge *Stanford Encyclopedia of Philosophy (Spring 2009 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/entries/common-knowledge/>> Last checked: January 2012
- Rinike Verbrugge. Logic and social cognition: The facts matter, and so do computational models. *Journal of Philosophical Logic*, 38:649–680, 2009.
- Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* Princeton University Press, 1998
- H. P. Young. The Evolution of Conventions, *Econometrica*, 61:1, 57-84, 1993

D. Zizzo. Experimenter Demand Effect in Economic Experiments. *Experimental Economics* 13:75-98, 2010

A Experimental Instructions

A.1 General Instructions (common to all treatments)

A.1.1 Introduction

Welcome! You are about to participate in an experiment on decision making funded by the Ministry of Education, University and Research. Instructions are simple, and if you follow them correctly and make appropriate decisions you will be able to earn an appreciable sum of money that will be paid to you privately and in cash. All choices that you make will be stored and processed anonymously. All your earnings in the experiment will be expressed in Experimental Currency Units (ECUs). 1 ECU is worth 1 euro.

A.1.2 The choice task

The present experimental session is composed of 10 participants and will last a total of 13 periods, divided into 2 parts: the first part will last for 3 periods, and the remaining part will last for 10 periods. The choice task will be the same for all the 13 periods. You are now listening to the initial instructions, to be followed by further written instructions that will appear on your monitor. *At the beginning of each period the software will pair each of you with another participant from the room picked randomly.* You and the other participant will both have to choose - independently from one another and without any possibility to communicate - one and only one number from 1 to 7 (included). *The euros you earn will depend on the number you have chosen and on the minimum number that has been chosen in the pair* (that is, the lower number between the one chosen by you and the one chosen by the participant with whom you are paired). Each of you has been given a paper copy of the double-entry table that will help you compute your earnings. *The table is identical for all participants*, and the same table will appear on your monitor every time you will have to make your choice for the round. Please have a look at the table now. The rows of the table report your possible choices in every period, i.e., the numbers from 1 to 7. The columns report the *minimum* number that has been chosen in the pair.

The table cells report the ECUs that you earn in correspondence of each possible outcome. The ECUs you earn are in the cell at the intersection between the row corresponding to the number you have chosen and the column corresponding to the minimum number chosen. For example, suppose you choose number 4 and the other participant chooses number 3: the minimum number in the pair is 3: therefore, your earnings are in the cell at the intersection between row 4 and column 3 (that is 8 ECUs).

Now, please pay attention to your monitor. The central part of the screen (green background) reports your earnings table, which is identical to the paper copy that you have. You will have to make your choice by typing the number chosen in the appropriate text box at the center of the blue square and by clicking on 'OK'. When all participants have chosen, your monitor will show you the period outcome: the number you have chosen, the minimum number chosen and the ECUs you have earned. At any time you will be able to double-check that your earnings have been computed correctly by using the copy of the earnings table that has been given to you. By clicking on 'Continue' you will move to the next period. Please remember that at the beginning of each period the software will match you with a participant picked randomly; therefore, in general you will be paired with a different participant in each period, although it will happen that you will be paired with someone more than once. In any case you will never be told the identity of the participants with whom you are paired.

At the end of the third period the experiment will be interrupted and the monitor will show you further instructions for you to follow. At some point, **BEFORE YOU MAKE YOUR CHOICE FOR PERIOD 4**, the experimenter will ask you to take the blank sheet of paper that is on your desk and write the following 2 predictions:

- The *average number* that will be chosen by all participants in period 4, i.e. the arithmetic mean of all choices in period 4
- The *average prediction* of participants, i.e., the average of all participants' forecasts regarding choices in period 4

This information will be asked only at the beginning of period 4, and you will be paid in the following way: for each prediction, you will earn 1 euro minus the arithmetic difference between your predicted value and the realized value, for a maximum of 2 euro in the case in which both your predictions are

100 per cent correct. The true realized values will be computed and publicly revealed at the end of the experiment.

A.1.3 Your earnings

Your earnings will be computed as follows: at the end of the experiment, the software will randomly pick one among the 13 periods, and each of you will be paid according to the ECUs he/she has earned in that period. The final screenshot will inform you of which period has been extracted and of how many ECUs you have earned. This earnings will be summed to the earnings from your predictions and to the 2 euro show up fee. Each one of you will be paid privately in cash. After receiving your payment, you will be free to leave the lab.

We now ask you to respond to all the questions reported in the anonymous questionnaire that has just been handed to you, and to give the questionnaire back to the experimenter. In case of incorrect answers, the relevant part of the instructions will be repeated. During the experiment you are not allowed to communicate in any way. If any forms of verbal or non-verbal communication are detected, the session will be immediately interrupted and nobody will be paid.

If you have questions, please raise your hands now.

A.2 Control questions (common to all treatments)

The present questionnaire is anonymous and only serves the purpose of making sure that all participants have understood the rules of the experiment before it starts.

- **QUESTION 1.** Imagine you have chosen number 1 in a period, and imagine that the participant with whom you are paired has chosen number 3. According to the table, how many ECUs have you earned? How many ECUs has he/she earned?
- **QUESTION 2.** Imagine you have chosen number 3 in a period, and imagine that the participant with whom you are paired has chosen number 2. According to the table, how many ECUs have you earned? How many ECUs has he/she earned?

- QUESTION 3. Your earnings will be determined by the ECUs that you have earned in one of the 13 periods that will be picked randomly by the software: TRUE or FALSE (mark with an X)
- QUESTION 4. You will always be paired with the same participant for all 13 periods: TRUE or FALSE (mark with an X)

A.3 Treatment-specific instructions appearing on each subject's monitor at the start of round 4

Note that these instructions were not read aloud.

A.3.1 Control treatment

Before proceeding to make your choices for the remaining 10 periods, the experimenter will ask you to write down your predictions for period 4 on the blank sheet of paper that you can find on your desk. After all participants have written their predictions, the experiment will restart. The matching rules and the rules determining your earnings in each period remain invariant

- PREDICTION 1: the *average number* that will be chosen by all participants in period 4, i.e. the arithmetic mean of all choices in period 4
- PREDICTION 2: the *average prediction* of participants, i.e., the average of all participants' forecast regarding choices in period 4

Now please write your predictions.

A.3.2 Mutual knowledge treatment

Before proceeding to make your choices for the remaining 10 periods, the experimenter will ask you to write down your predictions for period 4 on the blank sheet of paper that you can find on your desk. After all participants have written their predictions, the experiment will restart. The matching rules and the rules determining your earnings in each period remain invariant

SUGGESTION

WE SUGGEST YOU TO CHOOSE NUMBER 7. Note, in fact, that if both participants choose number 7, they both earn their maximum earnings, that is 13 ECUs

Now please write your predictions:

- PREDICTION 1: the *average number* that will be chosen by all participants in period 4, i.e. the arithmetic mean of all choices in period 4
- PREDICTION 2: the *average prediction* of participants, i.e., the average of all participants' forecast regarding choices in period 4

A.3.3 Message treatment

Before proceeding to make your choices for the remaining 10 periods, follow the instructions that will appear on your monitor. The software will create the random pairings among participants. Then you will be shown a message and you will be asked to send the message to the participant with whom you have been paired for period 4.

BEFORE you make your choice the experimenter will ask you to write down your predictions for period 4 on the blank sheet of paper that you can find on your desk. After all participants have written their predictions, the experiment will restart. The matching rules and the rules determining your earnings in each period remain invariant. Now please follow the instructions on your monitor and then wait for the experimenter to tell you to write down your predictions. es

Reviewer 1

Thanks for the kind words. We have answered the reviewer's remarks punctually below:

1a first of all: one typo/notational remark: p. 13 ($x_1 = \dots = x_n = 7$).

In the experiment you have $n = 2$, thus it would suffice to write ($x = 1 = x_2 = 7$)

answer We have rewritten the paragraph taking into account the reviewer's suggestion.

1b Table 1 (Payoff table) In a two player game a natural idea would be to depict own choice against opponent's choice rather than own choice against minimum of the two. Is there a special reason to take the latter?

answer We have decided to keep our table in order to make results comparable to the previous literature, and explained this in a footnote.

1c Orders of Knowledge: I have the impression the Message-Message condition does not create second, but third order information. First order is the information itself, second order is sending a message about the information to the other player and third order is sending the confirmation back. Even when starting to count the levels at a different position that seems to be definitely two levels above the information given in the mutual knowledge condition, not only 1.

answer It seems as though the referee is talking about levels of belief rather than levels of knowledge. In the "Message-Message" condition we see that

1. Statement appears on my screen \rightarrow I know "suggestion"
2. I click "send" \rightarrow I know "suggestion" (knowledge level does not increase here, because I can't be sure you have received my message)
3. I receive the acknowledgment \rightarrow I know that you know "suggestion"

While, as the referee seem to be suggesting, it is likely that at (2) I will believe that you know the suggestion, what we are interested here is levels of “hard” information. In this sense, it is only when I receive confirmation at (3) about you having read the suggestion that second order knowledge is attained. As the epistemic analysis is a relatively new tool in the literature, we prefer to keep our focus on knowledge only, leaving the study of belief and other attitudes to future research.

We have added a footnote in the text addressing this issue.

- 1d Table 4:** It would be interesting to have information on how often the minimum/maximum are played. The way the table looks like the information is contingent and the conclusions drawn from it seem contingent - especially given the different group sizes for the different groups.

answer While we agree that the referee is raising a relevant point, we decided not to delve deeper in the analysis of beliefs, as the main focus of our contribution is the analysis of choices induced by the suggestion.

- 1e Another remark to table 4:** A remarkable fact is that in three of the four groups the actual choice averages higher than the average of both, the first and the 2nd order belief. Given that playing above the expected average seems to be risky strategy (resp. staying slightly below sounds like a safe strategy), this seems irrational... or at least it seems as if there is no direct connection between expectation and the own move.

answer We say that “rendered subjects on average more optimistic, and greater optimism in turn induced the choice of the risky action significantly more often”. We have expanded on this remarking that an average choice larger than the average prediction also indicates optimistic, risk-seeking behavior.

- 1f p.6.** ‘...emerge as a constraints that we are justified to impose on the relation R . This sentence sounds a bit mixed up. The differences in the formalization ARE the different relations (and their properties), but these differences are justified by the concepts of knowledge/belief itself, not by the relation.

answer We have rephrased the sentence to avoid potential ambiguities.

Reviewer 2

The second reviewer raises important criticisms that persuaded us of the necessity of including a further experimental condition in order to properly address them. The main point is that our “Message” treatment is not homogeneous with our “Mutual Knowledge” and “Common Knowledge” treatments, in that its propositional content (“I suggest you to play 7”) is different from the propositional content of the other treatments (“the experimenters suggest you to play 7.”) We agreed on the reviewer’s criticism and added a new treatment homogenous with “Mutual Knowledge” and “Common Knowledge.” Punctual answers below:

2a The first suggestion is merely cosmetic: it is that section 1 be shortened. The general remarks about the relationship between formal epistemology and experimental analysis are unnecessary, since they don’t really reveal anything about the bearing of this particular experiment on the formal epistemology literature. In addition, the authors need not spend so much time spelling out the whole formal framework (accessibility relation, etc.), given that neither the experiment itself nor the discussion of it makes use of it. But this is merely a cosmetic point: I didn’t find anything factually or philosophically objectionable about the content of this section.

answer Our experimental setting suggests that the mode of communication may be more relevant than the actual epistemic content of the communicative act. This flies in the face of the qualitative difference between the finitary operation of mutual knowledge and the infinitary one of common knowledge. The purpose of having a rather detailed (albeit still very sketchy!) analysis of the formalisation of epistemic logic is fully functional to marking this contrast. We have added a paragraph in the conclusions emphasising this fact.

2b I have several worries concerning about the content of the proposition that is mutual-1/mutual-2/common knowledge. The proposition which the subject sees is this: ”We suggest that you pick 7.” There are

two ways to read this proposition: as an imperative and as an indicative. This suggests two different types of mutual knowledge that might be generated. On the imperative reading, level 1 would be "player 1 knows that he should pick 7 and player 2 knows that she should pick 7," and level 2 would be "player 2 knows that player 1 knows player 2 should pick 7 and player 1 knows that player 2 knows that player 1 should pick 7." On the indicative reading, level 1 would be "player 1 knows that he is told to pick 7 and player 2 knows that she is told to pick 7" and level 2 would be "player 2 knows that player 1 is told to pick 7 and player 1 knows that player 2 is told to pick 7." (Or even: "player 1 knows that if both pick 7 then they will get the highest payoff" etc.) The authors ought to explicitly state what they are trying to elicit levels of knowledge of. I take it to be the indicative reading, but it should be explicit. (This is more or less a cosmetic suggestion.)

answer We have added a paragraph (end of page 12) clarifying what is the propositional content elicited in the experimental treatments.

2c Levels of knowledge of either of these propositions would be interesting to examine. However, it appears to me that what is actually elicited in the mutual-knowledge-level-2 condition is "player 2 knows that player 1 is telling her to pick 7," since, as far as I can tell from the experimental design, the message that each subject is sending to the other player is "We suggest you to pick number 7" (not, for example, "The experimenter has suggested to me that I pick 7") and it is called "the suggestion" (not, for example, "information about the suggestion").

There are two major substantive problems with this. The first is that there is normative pressure to choose 7 in the mutual-2 condition that there is not in the mutual-1 condition. This arises from two different sources: first, that the other player is telling you to choose 7, and second, that you are telling the other player to choose 7 (and so one might think you are incurring an obligation to do what is best for the other player, conditional on her choosing 7). The experiment is billed as examining a purely epistemic phenomenon, and so the very plausible existence of normative pressure does undermine this claim. The authors do examine this possibility in the conclusion, and cite it as an area for further research: however, it seems to me to be the obvious explanation

for the data, and therefore an experimental design flaw. In other words, the authors say that they've found a surprising result, and suggest a factor that might be the cause and should be explored further: but given that the factor is present as a result of the experimental design, it seems that we don't yet have a "surprising" result. (If the authors want to keep the experiment as is, they ought to further explain how the finding moves the debate forward.)

answer We agree with the referee's assessment and we have changed the phrasing of our conclusions accordingly.

2d The second, and more devastating, problem is that given what is actually elicited in the mutual-2 condition, the proposition in which the players have common knowledge in the "common knowledge" condition different than the proposition in which the players have mutual knowledge in the mutual-2 condition. Therefore, the experiment hasn't shown that there is a difference between mutual and common knowledge of the same proposition. I suggest that the authors either reword the mutual-knowledge-level-2 condition or rewording the common-knowledge condition so that they differ only in knowledge structure, not in propositional content.

answer This is the remark that prompted us to run the new treatment session. We have modified our "message" treatment so that it is now crystal-clear that subjects are merely forwarding the *experimenter's* suggestion and not sending their own. The rephrasing of the experimental setup in our view makes the new treatment semantically homogeneous to the ones based on mutual and common knowledge. The results of this new session are described and analyzed in section 5, and are taken into account in the Conclusions of our paper.

2e The following is a substantive point, but smaller. I would like the authors to say something about how knowledge of "advice" is different from knowledge of facts. In the original Lewis discussion of salience, the idea was that convention is established when two potential coordinators have (levels of) knowledge that a particular choice is salient. This knowledge then exerts normative pressure because it produces mutual

expectations, and if one expects the other to choose etc., then if one does not choose in accordance, one fails to do what is best for oneself. But the content of the knowledge is not itself supposed to be normative in character. Indeed, if the proposition one had (levels of) knowledge of was that a particular outcome is normative (or recommended or advised) rather than salient, then this undermines the claim that the game gives rise to a convention in Lewis's sense. The Chaudhuri studies et al cited by the authors use the idea of advice, too, so I suppose it is not unprecedented in this literature, but I think it needs to be defended, at least briefly.

answer We have added a paragraph at the end of section 2 addressing this remark. The main idea is that “knowledge of advice” as it is studied in our paper and in the strand of literature we mention (Chaudhuri, etc.) is in fact relevant to Lewis-conventions. In Lewis' account, (common) knowledge of precedent induces in the agents a belief that they will pick the action made salient by its being a successful precedent. In our experimental setup, knowledge of advice induces in the agents a belief that they will pick the action made salient by being the recommended one. The symmetry reveals that, at bottom level, we are looking at the effectiveness of advice as a salience-building mechanism.