The Prescriptive Destiny of Predictive Attitudes: From Expectations to Norms via Conventions

Cristiano Castelfranchi (castel@ip.rm.cnr.it)
Institute of Cognitive Science and Technology – CNR
Viale Marx 15, Rome, 00137 ITALY

Francesca Giardini (giardini@media.unisi.it)
University of Siena, Cognitive Science Doctorate
Via Roma 47, Siena, 53100 ITALY

Emiliano Lorini (lorini@media.unisi.it)
University of Siena
Via Roma 47, Siena, 53100 ITALY

Luca Tummolini (tummoli@media.unisi.it)
University of Siena, Cognitive Science Doctorate
Via Roma 47, Siena, 53100 ITALY

Abstract

The goal of this paper is to show the normative component of a convention adopted by a population of cognitive agents. To address this aim we will defend two distinct thesis. The former is that even simple predictions developed to anticipate future state of affairs have an intrinsic tendency to evolve in full expectations and then in prescriptions. We consider this as a multilevel phenomenon occurring either at the individual psychological level or at the interpersonal one or, finally, at the collective macro social level. The latter thesis is that we consider this tendency as one of the possible paths of the spontaneous emergence of agents' commitments, of conventions and likely of real social norms: the tacit emergence of a prescriptive character and, then, of obligations and duties. We will examine the constitutive elements –both cognitive and relational – of this process of spontaneous transition from the predictive attitudes to the prescriptive ones and, on this basis, to real normative attitudes. Finally, we will discuss the inevitably normative component of conventions as traditionally described (Lewis 1969). We will argue that this fundamental process is notably left implicit or insufficiently explained.

Introduction

In everyday activities individuals are involved in social coordination problems. Such situations are commonly regulated by conventions. In Lewis (1969) classical definition, conventions are regularity of behavior where an agent *rationally* conforms "if and only if it is true that" the other agents conform too. In particular the belief that another agent will conform (what he defines as "expectation") provides the agent with a reason to conform too. Otherwise we propose that expectations are hybrid mental configurations whose components entail not only beliefs but also converging goals that those beliefs will be realized. This paper has the goal to show how such expectations in conventions tend to naturally evolve a prescription to conform to the regularity. We will argue that this is the way in which social norms can eventually emerge.

Consider the case of three students Mary, Julia and Barbara. They share an apartment and want to organize a party. Mary is teetotaller, Julia hates candies and Barbara's boyfriend is a DJ. They don't explicitly establish their own tasks, but they can coordinate simply referring to their tastes. So, Mary buys candies, Julia buys alcoholics and Barbara cares for the music. They have a great party, and the next time each of them will cover the same tasks relying on the expectations that the others will do the same. Since these expectations contain also a goal that the other will act in a particular way a mutual knowledge of these expectations is also an implicit request of compliance. We consider, as Lewis does, that a convention is already emerged. If next time Mary decides to buy alcoholics, instead of the expected candies, Barbara and Julia are entitled to tell her off for not conforming. They expected that she *ought* to conform and they have coordinated their actions to the expected conformity of Mary.

In this paper we will investigate how an implicit request can evolve in such a prescription to do what expected, where each agent will end up having a "right" to expect conformity and a "duty" to conform.

Cognitive Ontology

In order to argument our theses, we will specify the mental objects which we assume are present in the transition from mere beliefs to full prescriptions.

What a mere forecast and prediction are

We call *forecast* a belief about a future state of the world and we distinguish it from a simple *hypothesis*. The difference is in term of degree of certainty: an hypothesis involves the belief that future p is *possible,* whereas a forecast the belief that future p is *probable*. A forecast implies that the chance threshold has been exceeded (domain of probability). According to the agent's past experience or knowledge of physical or social rules and laws p *should* happen (in an epistemic sense). We have a *prediction* when the degree of certainty is close to 100 per cent.

What a real expectation is and how it can be positive or negative (hopecasts and fearcasts)

Expectations in our ontology are not indifferent hypothesis, forecasts or predictions. They imply a subjective concern in the realization of p. This is why one speaks of "positive" or "negative" expectations. In order to discriminate positive from negative expectations and weak (possibility) from strong (probability) expectations we introduce four terms (two of them do not exist in English): *hope, hope-cast, fear, fear-cast*. In the case of a *hope-cast* (positive strong expectation), I want p to be true; in the case of a *fear-cast* (negative strong expectation), I want it to be false. In both cases I believe that p is probable. In the case of a *hope* (positive weak expectation), I want p to be true; in the case of a *fear* (negative weak expectation), I want it to be false. In both cases I believe that p is possible.

The following table shows formulas for previous mental objects: simple hypothesis and forecasts and expectations (positive and negative)[1].

Table 1: Basic cognitive ontology

HYPOTHESIS$_i$ p $BEL_i\ (w_i\ (LATER\ p) \leq ½)$
FORECAST$_i$ p $BEL_i\ (w_i\ (LATER\ p) > ½)$
HOPE-CAST$_i$ p $BEL_i(w_i(LATER\ p)>½) \wedge GOAL_i\ LATER\ p$
FEAR-CAST$_i$ p $BEL_i\ (w_i\ (LATER\ p)>½) \wedge GOAL_i \neg LATER\ p$
HOPE$_i$ p $BEL_i(w_i(LATERp) \leq ½) \wedge GOAL_i\ LATER\ p$
FEAR$_i$ p $BEL_i(w_i(LATER\ p) \leq ½) \wedge GOAL_i \neg LATER\ p$

For instance, *Agent i has a hope-cast about p* means that *Agents i believes p will happen later (Agent i does not believe p is holding at the present) with a degree of certainty up to 50 per cent and Agent i wants p to hold later*[2]. We can define either expectations about actions (our actions or other agent's action) or about events of the world. For instance, *Agent x has a hope-cast about Agent y's action* when *Agents x believes (with a degree of certainty up to 50 per cent) Agent y will execute action a and Agent x wants it*[3]. We have specified expectations in term of goals and beliefs. However in this analysis we assume that the agent's conative component is an *intention* (Bratman, 1988).

The socio-cognitive essential nature of the "prescription"

A *prescription* is an illocutionary act based on a complex socio-cognitive structure (Searle, 1969). We consider however that it does not necessarily imply linguistic communication between the agents and can be supported by tacit behavioral communication between them.

Along the path from simple prediction to true prescription we distinguish four cognitive steps:

---

[1] We use an extension of Cohen and Levesque's multimodal logic for intentional states (1990) to probabilistic logic for knowledge and belief (Fagin & Halpern, 1994; Fagin, Halpern & Megiddo, 1990). We briefly summarize the semantics for *FORECAST p* formula in order to explicit the logic we are using.

In Cohen and Levesque's logic, models are defined as M = <Θ, P, E, Agt, T, B, G, Φ>, where Θ is a set of objects, P a set of agents, E a set of primitive events.

Agt∈ [E → P] specify agents in relation to events, T⊆ [***Z*** →E] is set of possible courses of events (worlds) specified as a functions from integers to elements in E, B⊆ T x P x ***Z*** x T is the accessibility relation for beliefs; G⊆ T x P x ***Z*** x T is the accessibility relation for goals, and Φ is the interpretation for predicates.

Given the semantics of Fagin and Halpern for uncertain knowledge, our models are defined as M = <Θ, P, E, Agt, T, B, G, Φ,**P**>.

So,

M,σ,v,n $FORECAST_i\ p$ that is

M,σ,v,n $BEL_i\ (w_i\ (LATER\ p) > ½)$ iff for all σ* such that $< \sigma,n> B[v(x)]\ \sigma^*$,

M, σ*,v,n* $(w_i\ (LATER\ p) > ½)$, that is

$\mu_{i,\ (\sigma^*,n)}\ (T_{i,\ (\sigma^*,n)}\ (LATER\ p)) \geq ½$

In the present extension **P** is a probability assignment such that, for each agent i ∈ {1,……,n} and a time point n in a course of events *σ*, assigns a probability space $P(i, (\sigma^*,n)) = (T_{i,\ (\sigma^*,n)}, \chi_{i,\ (\sigma^*,n)}, \mu_{i,\ (\sigma^*,n)})$ with $T_{i,\ (\sigma^*,n)} \subseteq T$ an arbitrary set of T (sample space), $\chi_{i,\ (\sigma^*,n)}$ a σ-algebra on measurable subsets of $T_{i,\ (\sigma^*,n)}$ and $\mu_{i,\ (\sigma^*,n)}$ a probability measure on elements of $\chi_{i,\ (\sigma^*,n)}$. $\mu_{i,\ (\sigma^*,n)}\ (T_{i,\ (\sigma^*,n)}\ (LATER\ p))$ gives the probability on the measurable subset of $T_{i,\ (\sigma^*,n)}$ where (LATER p) holds.

[2] We assume 50 per cent as the chance threshold. For simplicity we are assuming a probability space composed by two events (p and ¬p). If we consider n different events, the chance threshold would be 1/n.

[3] The formula is $HOPE\text{-}CAST_x\ y\ a =$
$BEL_x\ (w_x(LATER\ (DONE\ y\ a))>½) \wedge A\text{-}GOAL_x\ DONE\ y\ a$.
We could also define fear-cast about agent y's action (even for hope and fear) as Agent x's belief that Agent y will do action a and the goal that Agent y will not do that action.

(I) EXPECTATION: the generation of a Goal relative to Agent y's action a; i.e. the "intention that" y Does a (Grosz & Kraus, 1996); (this – combined with the belief - produces a full "expectation");
(II) INFLUENCING GOAL: the idea that it might be useful inducing the other to Do a, thus the goal of influencing him;
(III) REQUEST: the goal of influencing him through the (tacit or explicit) communication of such a goal, by exploiting his adhesion to my request;
(IV) Full PRESCRIPTION: I want your adhesion to my goal on the basis of an *obligation* of yours.

(IV) implies (III) that implies (II) that implies (I). Each step adds some additional mental and relational ingredient and defines a more restricted set.

In a sense, (II) and (III) are weak forms of "prescription", just subjective, merely in my mind, since I want that you do something and I'm acting in order to influence you to do so. Expectations about intentional agents[4] contain true (tacit) imperatives (III) and arrive to full prescriptions (IV). In both there is a goal that Agent *y* does a, based on a goal that Agent *y* intends to do a, based on a goal that Agent *y* comes to know that *x* has the former goal. To have real prescriptions or imperatives *y* must be a cognitive, intentional agent; thus the aim of the imperative or prescription is in fact an intention of the other to do so (to perform a)[5]. The prescription presupposes that *y* does not already and certainly intend to do a and will do a autonomously. Agent *y* will probably be induced to intend to do so also by the fact of believing that *x* has the goal that *y* does a; i.e. for goal-adoption and more specifically for "adhesion" to the request/prescription. This means that x has the goal that *y* knows about his goal. So, x can use some form of tacit or explicit communication to *y* to achieve the goal, or he can have some reasons for assuming that *y* already knows or will understand that *x* has such an expectation on her behavior.

However, full prescriptions are only those also based on the idea that there is an obligation impinging on you, and binding you to do so *because* you have an obligation. Either this obligation is created by communication (orders), or it is just instantiated since a norm is already there (solicitations and reminders), or even is elicited by an agent's social commitment – a (tacit) promise – and the obligation is created by the agent himself while committing to do so (contract proposal).

This means, in our framework (Conte & Castelfranchi, 1995), that there is a Normative Belief[6].

We argue that expectations frequently not only entail goals (that is why they can be "positive" or "negative"), but are much more cogent, since they entail also influencing goals and even true prescriptions (although tacit). Especially in social conventions, the evolution from mere beliefs to prescriptions is unavoidable, and is based on communication, on mutual social commitments, and on the following emergent circular process: *a consolidated convention, thanks to the prescriptive nature of the expectations, becomes a social norm, and on such a basis expectations on the others become "prescriptions".*

The psychological tendency: From predictions to goals and to prescription

We examine here several independent reasons why a simple belief about a future action or event tends to become a full "expectation" (i.e. to be joined by a converging goal), and how that expectation tends to imply a normative component. Those mental mechanisms from the individual to social situations are: uncertainty and need for prediction; acquaintance with the expectation, reliance and delegation and expectations based on commitments and obligations.

Uncertainty and need for prediction

The first process we want to focus on has been investigated in psychological literature. We are referring to *predictability*, that is the cognitive component of *self-efficacy* (Bandura, 1982): the need to anticipate future events, and the consequent need to find such anticipation validated by facts. This need for prediction[7] is functional in order to avoid anxiety, disorientation and distress. Cooper and Fazio (1984) have experimentally proved that people act in order to find their forecasts (predictions) validated by facts, and that feel distressed by invalidation. In many cases, people have the tendency to behave in accordance with their predictions (Sherman, 1980). We believe that, when an individual x forecasts (predicts) an individual y's action, he needs to validate his prediction. He has the tendency to behave *as if* he wanted the predicted action's execution.

[4] Even predictions about unanimated agents (weather; sun; traffic) become hope/fear-casts (concerned expectations): i.e. a goal joins the belief.

[5] When applied to a cognitive agent, the goal that agent y does a presupposes the goal that y intends to do the action (Castelfranchi, 1998).

[6] A formula for Normative Belief is:

$$(\text{N-BEL}_x\ y\ a)=\text{BEL}_x(\text{OUGHT}(\text{Does}\ y\ a))$$

We have introduced the predicate OUGHT in order to fully express that there is some sort of obligation on y to perform a given action. For the purposes of this paper we will assume obligation as a primitive which defines the set of worlds in which the action follows from obligations

[7] Miceli and Castelfranchi (2002) consider the need for prediction as a *metagoal* of the mind in the sense that it is a regulatory principle concerning one's mental functioning (vs a regulatory state explicitly represented).

Acquaintance with the expected scenario

As we have already noticed, when a positive expectation is invalidated people suffer distress. That type of distress is not simple disappointment (a form of distress for a failure in goal achievement "mixed" with surprise for the invalidated prediction). Rather, an invalidated hope-cast looks like an ill-treatment. A common reaction to invalidated positive strong expectations is anger (Averill, 1982; Burgoon, 1993), coupled with a sense of injustice.

The stronger the hope-cast –that is, the more certain its implied prediction and the more important its implied goal- the stronger is the sense of injustice. Therefore an invalidated hope-cast is often a *violated* one. Since the belief that p will happen is well grounded and there is the goal that it does happen, p turns in something *bound* to happen. Some norm or prescription is already implied at this pure psychological level. People have the tendency to add a normative component to their positive strong expectation. When we predict that a future event will happen and we have the goal that happens, an implicit norm that it *ought to happen* is there[8].

Reliance and Delegation

Delegation (i.e. an agent relying on another agent's action) is a typical process for turning a simple prediction or forecast into an expectation. We consider here weak delegation in a situation of positive interference, by adopting the perspective of the single Agent x. Positive interference is a social precognitive notion[9], and weak delegation is the unilateral decision of an agent to delegate (part of the plan for reaching a certain state of the world), without any agreement *expressed* by the other agent. In weak delegation Agent x is aware of Agent y's positive interference[10]. We add something more: Agent x is almost sure Agent y will do action a (since he believes Agent y wants to execute that action). This condition extends the condition of *positive interference awareness* and is necessary in order to guarantee delegation. The condition is: *Agent x forecasts that if Agent y will do action a then p will be reached and Agent x forecasts Agent y will execute that action (since he believes Agent y has it as a goal) and Agent x has p as a achievement goal*[11]. In such a case, it is likely that Agent x relies on Agent y for action execution: since Agent x's degree of belief is high, he delegates part of his plan to Agent y in order to get p. Agent x prefers to delegate than executing the action by himself. Since Agent x has built a plan that includes Agent y's action (a sort of multi-agent plan) from now on x is *committed* on Agent y's action. In other words, Agent x has an "intention that" Agent y will execute the delegated sub-plan. Agent x expects that Agent y will execute the action, in his mind action a *should* be executed by Agent y.

Expectations based on commitments and obligations

In mutual interference situations, each agent adapts his plan relying on his expectation about the others. Since every agent is changing his autonomous course of action, to achieve coordination they need to be aware of each other coordinating intentions. Having the positive expectation, the agent has generated the goal that the other do the delegated action. We have previously considered this cognitive step as an *influencing goal*: the goal that the other has a certain goal. In mutual interference the agent believes that also the other agents need to act relying on their expectations. So the agent will have also the *communicative goal* that the others know about his expectation because then they will have a reason to do what expected. The practical action of coordination can be either a form of implicit communication (my action is deliberately intended to generate also a belief about my expectation by means of the recognition of this intention) or simply a sign of the underlying expectations (observing a regularity in behavior in others reveal their mutual expectations). A regularity in coordination can spread the belief about the knowledge of expectation. The agents will act relying also on this belief while having the corresponding goal. They have a proper *expectation*, in our sense, about the knowledge of each other expectations. This process leads beyond the mutual knowledge of expectations as assumed in Lewis to mutual expectation about expectations.

Agent x can also come to assume that Agent y has agreed to his delegating decision. This is a weak sense of agreement, it is a tacit agreement which x counts upon for further decisions. In such a situation positive expectations become *entitled expectations* because mutual knowledge of positive expectations are considered by the agents as tacit *requests*. Because Agent x expects Agent y to know his expectation, if y would not like to adopt x's goal she should have rejected the request. Even a tacit agreement can create a social commitment for agent y. With *social commitment*

---

[8] The process is even more complex. In fact the sense of injustice involved in violated hope-casts is also linked to a sense of loss. See Miceli and Castelfranchi (2002) for a discussion.

[9] We have positive interference of Agent y's action a with Agent y's goal p when *Agent y's action a favors Agent x in achieving a goal p when p is a consequence of action a*, i.e., $(FAVOUR_x\ y\ a\ p) = ((DONE\ y\ a) \rightarrow p) \wedge (A\text{-}GOAL_x\ p)$.

[10] Agent x is aware of that positive interference when *forecasts that if Agent y will do action a then p will be reached and Agent x has p as an achievement goal,* i.e., $(AWARE\text{-}FAVOUR_x\ y\ a\ p) = FORECAST_x\ ((DONE\ y\ a) \rightarrow q) \wedge (A\text{-}GOAL_x\ q)$.

[11] The formalization is $FORECAST_x\ ((DONE\ y\ a) \rightarrow q) \wedge FORECAST_x(DONE\ y\ a) \wedge BEL_x\ (A\text{-}GOAL_y\ (DONE\ y\ a)) \wedge (A\text{-}GOAL_x\ q)$.

(Castelfranchi 1998) we intend to refer to a form of Goal Adoption where Agent y is committed to Agent x to do a if Agent x is interested in a. Both the agents know that y intends to do the action whose result p is a goal of Agent x (as a consequence of the mutual knowledge of expectations) and that, as for a, x is *entitled* to expect y doing a (x has a sort of "right" on y) and hence he wants it. When even the entitled expectations are mutually expected than this socio-cognitive structure creates for both the agents an (interpersonal) *obligation* to do the expected action. Having mutual entitled expectations the agents are not only disappointed if somebody violate them, but can feel a stronger sense of injustice. Mutually attributing a social commitment the agents feel entitled and *prescribe* the expected action.

From conditional and mutual expectations (conventions) to norms

Once defined the psychological and social path leading from predictions to expectations and then to prescriptions, we want to compare our statements to traditional approaches to convention. A convention is considered as a regularity in behavior evolved by two or more people to solve recurrent coordination problems.

Thus a convention can be considered as a solution to a special case of positive interference. When there is a social coordination problem agents *reciprocally* interfere in their autonomous actions. Conventions deliver high degree of beliefs (forecasts or predictions) that the other agent will perform the needed action. This is why very likely we will have delegation in convention. Since all the agents build multi-agent plans involving each other actions, in convention deeply rest a structure of "intentions that" and hence of positive expectations (hope-casts). Positive expectations on the other agent's action imply the prescription to do as expected.

We consider ill defined the classical Lewis' definition of convention (1969), because it lacks this link between forecasts (the weak sense in which we interpret Lewis' expectations) and true prescriptions. However Lewis' proposal shares some relevant elements with our argumentation: a)It focuses on a regularity in the behavior of members of a population; b)It is the solution for a problem of reciprocal coordination, i.e. an agent following a convention can achieve a better coordination among agents and avoid interferences; c)It stresses the importance of expectations in order to drive the choices and the behavior of other agents; d)It focuses on the importance of conditional preferences and shows the relevance of shared knowledge, expectations and behavioral conformity.

However conventions are not simply beliefs about the behavior of another agent and his intentions to conform (since the others do the same), they also entail the goal that the others (continue to) conform. It is to some extent true that the higher a given behavioral regularity, the more likely it will be prescribed, so people assume that violations of expectations, in our sense, are disapproved and conformity is not simply expected, but also prescribed (Conte & Castelfranchi, 1999).

It seems that, for Lewis, the fact the other will conform is either only a belief justifying the choice or it is the aim I want to persecute by conforming. What is missed is the agent's goal (not simply the belief) that the other will conform, i.e. the goal about the other's mind. He seems to reduce expectations to justified forecasts merely based on past experiences.

Alterman and Garland (2000) share our criticisms about the lack of an adequate set of cognitive assumptions. They consider too static Lewis' notion of convention as a recurrent situation S, because what is a regularity S is negotiated among the participants as a part of their social interaction. They consider another problem with Lewis' definition the notion of convention as based on a fixed regularity R. The participants can share expectations about the structure of a conventional activity, but the actual structure of the conventional behavior on a given occasion will be uniquely determined on each occasion.

Mark (in press) proposes an analysis of conventions in organizational context, where they represent a solution in order to solve coordination problems among distributed groups attending to the same task. Conventions spread from mutual knowledge and expectations in the group, therefore the existence of a convention indicates that all in the group have common knowledge that this convention leads to solving a coordination problem. From these premises follows that as a consequence of beginning to conform to the convention, group members develop expectations, which we call positive expectations (hopecast), that others in the group will conform. This leads to commitments which sustain the convention.

Bicchieri (1990) wants to show why people should conform to social norms. She defines a social norm as function of individual choices and, ultimately, of individual preferences and beliefs, but she doesn't recognize the importance of expectations in prescribing a conforming behavior. Bicchieri uses expectations in order to explain people conformity behavior, but she sustains that the emergence of social norms is a matter of learning in a small group, rather than of the development of prescriptions.

Beyond any criticism, our final aim is to integrate Lewis' definition in order to stress the prescriptive nature of expectations and the emergence, on such basis, of obligations.

In our framework what happens is that:

- I conform to R because you expect it, in this sense I adopt your goal. Conformity can also stem from my adhesion to your request: I follow a prescription also because I know that it is your

desire so I expect (and demand) that you will reciprocate;
- My conformity is conditional to your choice to conform because "people do not want others in the same conditions as their own to sustain lower costs, benefits being equal"[12]

If everyone expects and prescribes conformity when the agent doesn't conform he is violating a norm and not simply making an unexpected action. The expectation that others will conform in the future is a reason to continue to conform: if I conform, then others will.

Conclusion

Our aim has been to show that a social norm (tacit unless someone tells someone else off for not conforming) can stem from a recurrent situation where an expected (forecast) action becomes a prescribed one, thanks to the prescriptive force of hopecasts/expectations. First, we have analyzed the cognitive ontology of forecasts, predictions and prescriptions, we have defined their role in social coordination and finally we have shown the path leading from conventions, as defined by Lewis, to social norms. Our hypothesis is that mutual expectations, based on shared knowledge, have a prescriptive nature and not simply a predictive one.

This prescriptive nature of expectations can make a social norm out of a consolidated convention. A convention here is considered as an objective social phenomenon, which does not need to be represented as a convention to exist. A convention is there when a situation of mutual interference between agents develop a pattern of mutual expectations. Differently a norm must be explicitly represented in the agents' mind to exist (Conte and Castelfranchi 1995). When the convention emerges cognitively it becomes a social norm: a prescription issued by the whole community which acts, above personal requests and individual desires, as a sort of impersonal authority. Such a social norm is not perceived as a personal request based on private interests, but as a request of an abstract entity with power exerted for general interests. In this perspective, deviance from a convention/social norm is not simply a disappointment of private interests but a violation of general ones and authority. A shared regularity of action create a normative reality: it *ought* to be done like this because everybody does like this.

Acknowledgments

This research has been partially supported by the MIUR Project "Agenti software e commercio elettronico".

[12] see Conte & Castelfranchi (1999) for a discussion.

References

Alterman, R. & Garland, A. (2000). Convention in joint activity. *Cognitive Science*, 25(4), 611-657.

Averill, J. (1982). *Anger and aggression: An essay on emotion*. New York: Springer.

Bandura, A., (1982). Self-efficacy mechanism in human agency. *American Psychologist*, 37, 122-147.

Bicchieri, C. (1990). Norms of Cooperation, *Ethics*, 100, 838-861.

Bratman, M. E. (1988). *Intentions, plans, and practical reason*, Cambridge, MA: Harvard University Press.

Burgoon, J.K. (1993). Interpersonal expectations, expectancy, expectancy violations, and emotional communication. *Journal of Language and Social Psychology*, 12, 30-48.

Castelfranchi, C. (1998). Individual Social Action. *Artificial Intelligence*, 103, 157-182.

Cohen, P. R. & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, 213-261.

Conte, R. & Castelfranchi, C. (1999). From conventions to prescriptions. Towards an integrated view of norms, *Artificial Intelligence and Law*, 7, pp. 323-340

Conte, R. & Castelfranchi, C. (1995). *Cognitive and Social Action*, London: UCL Press.

Cooper, J. & Fazio, R.H. (1984). A new look at dissonance theory. In L. Berkovitz (Ed.), *Advances in experimental social psychology*, Vol. 17, (pp. 229-266). San Diego, CA: Academic Press.

Fagin, R. & Halpern, J.Y. (1994). Reasoning about Knowledge and Probability. *Journal of the ACM*, 41(2), 340-367.

Fagin, R., Halpern, J.Y. & Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation*, 87(1/2), 78-128.

Grosz, B.J. & Kraus, S. (1996). Collaborative Plans for Complex Group Action. *Artificial Intelligence*, 86, 269-357.

Lewis, D. K., (1969). *Convention: a philosophical study*, Cambridge: Harvard University Press

Mark, G., (in press). Conventions and Commitments in Distributed Groups. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*.

Miceli, M. & Castelfranchi, C. (2002). The Mind and the Future. The (Negative) Power of Expectations. *Theory & Psychology*, 12(3), 335-366.

Searle, J. R. (1969). *Speech Acts*. Cambridge: Cambridge University Press.

Sherman, S.J. (1980). On the self-erasing nature of errors in predictions. *Journal of Personality and Social Psychology*, 39, 211-221.