Promises and Threats with Conditionals and Disjunctions

Robert van Rooij and Michael Franke

Version of: January 25, 2010

Abstract

With a conditional "If you do ..., I'll do ..." we can make promises and threats. But with a disjunction "You do ..., or I'll do ..." we can only make threats, no promises. We suggest that this so because disjunctive promises would be a suboptimal strategic commitment in a game-like situation where the speaker is trying to influence the hearer's choice of action.

1 No Disjunctive Promise

In propositional logic, formulas $\neg A \rightarrow P$ and $A \vee P$ are equivalent, and our intuitions about their natural language counterparts also, for the most part, support this equivalence. However, as is often the case where two expressions are logically equivalent, the pragmatics of conversation poses restrictions on the acceptability of one that do not seem to apply to the other. In particular, in the case of conditionals and disjunctions an interesting pragmatic difference surfaces in the context of *inducements*, where the speaker is trying to influence the behavior of the hearer by conditional promises and threats: whereas the conditional statements in (??) and (??) can be a threat and a promise, used respectively to induce the hearer to hand over her wallet, the alleged disjunctive equivalents in (??) and (??) are *both* preferably read as threats.[1]

(1)   a. If you don't give me your wallet, I will punish you severely.
          $\neg A \rightarrow P$ threat

          b. If you give me your wallet, I will reward you splendidly.
          $A \rightarrow R$ promise

          c. You will give me your wallet or I will punish you severely.
          $A \vee P$ threat

          d. ? You will not give me your wallet or I will reward you splendidly.
          $\neg A \vee R$ threat

---

[1] Notice that in (??) letters P and R stand for punishment and reward respectively.

This is astonishing, in particular because if we expect (**??**) and (**??**) to be logically equivalent, we could also expect (**??**) to be a promise, just as (**??**) is. But, if felicitous at all, (**??**) rather reads as an inducement *not* to hand over the wallet, and as a threat that the hearer will be *punished* with a splendid reward if he does not comply. Since especially the possibility of a punishment by a "splendid reward" is awkward, the preferred threat-reading of (**??**) may therefore lead to verdicts of pragmatic infelicity.

The question that we would like to address in the following therefore is: what makes it so that a disjunction of the form $A \lor P$ is preferably read as a threat and not as a promise? The answer we give is that, from the point of view of a rational speaker who is concerned with the efficacy of her inducements, an alleged disjunctive promise, such as in (**??**), is a *suboptimal*, i.e., *irrational*, conversational move. Therefore, the preferred way of interpreting, i.e., rationalizing, a disjunctive inducement given by a rational speaker is as a threat, even if that goes against contextual assumptions of hearer-desirability.

2   Relation with Previous Accounts

The pragmatic puzzle we are trying to account for here may be familiar from discussions of so-called pseudo-imperatives (see **????**, for early contributions). Pseudo-imperatives (shortly, PIs) are mixed mood sentences where an imperative clause is followed by conjunction "and" or disjunction "or" and a declarative sentence. In other words, a PI is a sentence of the form:

(2)   a.   Do A, and X will happen/be the case/be done.
         b.   Do A, or X will happen/be the case/be done.

Interestingly, PIs behave quite similar to the sentences in (**??**):

(3)   a.   Close the window and I will kiss you.   threat
            Do A and R
         b.   Close the window and I will kill you.   promise
            Do A and P
         c.   Close the window or I will kill you.   threat
            Do A or P
         d.   ? Close the window or I will kiss you.   threat
            Do A or R

What is peculiar is that we can find examples of conjunctive PIs which read as conditional threats (**??**) and examples which read as conditional promises (**??**), depending on whether we assume that the hearer wants the declarative second conjunct to be realized. But, for disjunction, only threat-readings, so to speak, are possible: any second disjunct is either construed as hearer-undesirable or else the whole disjunctive PI seems pragmatically infelicitous.

Several issues are worth the linguist's attention here. Firstly, it needs to be explained how a conjunction "and" can obtain a kind of conditional reading in

the first place, especially one in which the illocutionary force typically associated with an imperative clause cancels out. Secondly, it needs to be explained how a disjunctive PI cannot function as a promise.[2] It is the latter problem that this paper deals with, including but not restricted to PIs. In our exposition, we will focus on the contrast between conditionals and disjunctions of declarative sentences of the form in (??). This carries over to PIs if we may assume that (i) conjunctive PIs have conditional readings, and that (ii) imperatives have some descriptive content that refers to a hearer action or a hypothetical state of affairs. Our aim, then, is to be as linguistically sober as possible: treating clauses as denoting propositions and remaining as conservative as possible in our analysis of conditionals, negation and disjunction, we would like to explore to what extent rationales of influencing others' behavior by promises and threats alone can explain the contrast between conditionals and disjunctions.

3   Promises & Threats as Strategic Commitments

An account of conditionals and disjunctions as possible promises and threats requires us to establish sufficient transparency of our view of the basic notions of promise and threat. In line with both linguistic and game-theoretic analyses, we take these to be special kinds of speech acts with which the speaker commits herself to a particular course of events.[3] Following the seminal work of ?, we will explore how such commitment can be used *strategically* in order to influence the hearer's behavior, and what difference it may make to strategic commitments whether they are of conditional or disjunctive form. Towards this end, we will first look briefly at commitment-based analyses of speech-acts in sections ?? and ??. After that we will zoom in on the rationality requirements of strategic commitments in section ??.

3.1   Speech-Acts and Speaker Commitment

In linguistic pragmatics, ?, for instance, proposes to analyze speech acts in general as functions that change a context in such a way that it alters the speaker's commitments. An assertion that φ is true, for example, would come out as an update function on the context of utterance such that after the utterance (and its uptake) the speaker is committed, at least if prompted and within certain limits, to defend the truth of φ (c.f. ?983:Asserting1,). Under this view, promises and threats can be regarded as speaker commitments as well. For instance, ? (?, p.69) writes: "A promise that φ is a function that changes a context in which the speaker is not committed to bringing φ about into one in

---

[2]Attempts to solve the first problem have been made by, for instance, ?, ?, ?, ? and ?. Explanations of the second problem feature of disjunctive PIs have been attempted, for instance, by ?, ?, ? and ?.

[3]Whether it is in her power to actually bring the event to which she commits herself about, or whether it is only a doxastic commitment is not crucial for our concern.

which he is so committed."[4]

What is a threat, and what is the difference between a threat and a promise? Psychologically speaking, one can feel threatened by another person, just because one fears that this other person might be harmful. Notice that according to this conception, one can feel threatened by someone without this person having done anything. However, in this paper we will adopt a more operational notion. A threat is a commitment by one person intended to change another's person future behavior. But, how, then, does a threat differ from a promise? Also a promise involves a commitment of the speaker often with the intention to change the hearer's future actions. However, the major difference is that whereas in case of a threat, the commitment has negative consequences for the other, with a promise these consequences are positive.[5]

Still, there is a further difference between promises and threats. According to **?**, unlike promises, "no obligation is involved in threatening" (p.193).[6] That means that although by uttering a threat the speaker commits herself to carrying out a sanction if needed, she is strangely enough no *obliged* to do so in the same way that she would be in case of a promise. Clearly, if, say, Jones threatens to punish Smith in case he giggles, then, after giggling, Smith will not *insist* on his punishment, and, more importantly, also cannot lay claim to a social obligation that he be punished. The case of a promise, obviously, is essentially different in this respect. This is an interesting point to which we will return later: from the speaker's point of view, a threat is *cheaper in expectation* than a promise, all else being equal, because, as **?** put it, "threatening is not as institutionally dependent as promising" (p.193).

3.2 Conditionality of a Commitment

In order to be effective as an inducement, promises and threats should not be made unconditionally: intuitively, the sentences in (**??**) do not commit the speaker to severe punishment or splendid rewards *come what may*, but only in case the hearer did or did not hand over her wallet; a promise of a splendid reward come-what-may will not convince anyone to hand over his wallet, and neither will the threat (if you can call it such in the first place) that one is to be severely punished independent of one's own choices. For our present purposes, it is actually a mood point to ponder whether such inducements are essentially conditional commitments or commitments to conditionals.[7] What ultimately counts for the efficacy of an inducement by promises and threats is that the

---

[4]There are further felicity conditions that one would have to consult if a flawless conceptual characterization of promises and threats qua speech act was at stake (the locus classicus is Searle's (**?**)'s exemplary analysis of a promise). This is not crucial though for any of our present concerns.

[5]This does not exclude that we sometimes use the terms in a slightly misleading way, as in "If you 'lend' me your wallet, I promise you I won't hurt you."

[6]This point is corroborated by empirical data reported by **?**.

[7]**?**, for instance, argues that under commitment-based analyses of assertion conditional assertions and assertions of conditionals come down to essentially the same thing.

hearer will get to understand that a certain reward or punishment is *conditional upon* his choice of action.

From this point of view, the presumption of rationality of strategic commitments explains the fact that conditional promises and threats readily receive biconditional, so-called *conditional-perfection* readings.[8] A conditional threat also implicates a promise, and the other way around. The threat that you will be punished if you do not hand over your wallet is effective only to the extent that (it is commonly understood that) the speaker promises that you will *not* be punished if you *do* hand over the wallet. An analogous argument applies to inducements by promises, of course. Hence, under the assumption that the speaker is rational, i.e., concerned with the maximal efficacy of her statements also on the level of what is implicated, we should enrich a conditional promise or threat to a biconditional reading. Obviously, the same rationale applies to exclusive readings of disjunctions: also a threat such as (**??**) is read exhaustively; and so should, all else being equal, a promise like (**??**), if it was felicitous after all. We will come back to this point also later in our analysis in section **??**.

3.3 Strategic Commitments

The view that conditional promises and threats are *strategic commitments* of the speaker with the aim of influencing the hearer's choice of action in a game-like situation, is a centerpiece in the analysis of human interaction that Thomas Schelling became famous for (**?**).[9] In sequential games it is normally the agent who acts first who has the advantage. The idea of making a strategic commitment is to seize the initiative, even if your action follows that of the other. On **?**'s account, a *commitment* of the speaker is modelled as a pruning of the game tree: by a binding public announcement the speaker ostensibly excludes some of her action alternatives in a sequential game (cf. **?**). Look at the simple game in figure **??**: the opening move is the hearer's choice of A or $\neg A$ and subsequently the speaker performs some action X or Y. The commitment "if A, then X", for example, is modelled as pruning (at least) the branch w_2 from the game (and possibly also, if conditional perfection is taken into account, the branch w_3).

A *strategic commitment* is a commitment in this sense with the intent to influence the hearer's behavior. Obviously, by the prospect of a splendid reward, for instance, the speaker can persuade the hearer into performing an action that the speaker likes, but the hearer dislikes. Such a strategic commitment is subject to a number of obvious *rationality* constraints, in particular whether the commitments are *credible*, *beneficial* and *efficacious* (see **?**, for a more thorough formal account).

---

[8] For theoretical assessments of reading "if" as "if and only if" see, for instance, **?**, **?**, **?** and references therein. Empirical research supporting the strong availability of an "if and only if"-reading in conditional promises and threats is reported by **?** and **?**.

[9] It is here also that promises and threats differ from assurances and warnings: they might have the same effect, but only the former two involve commitments.
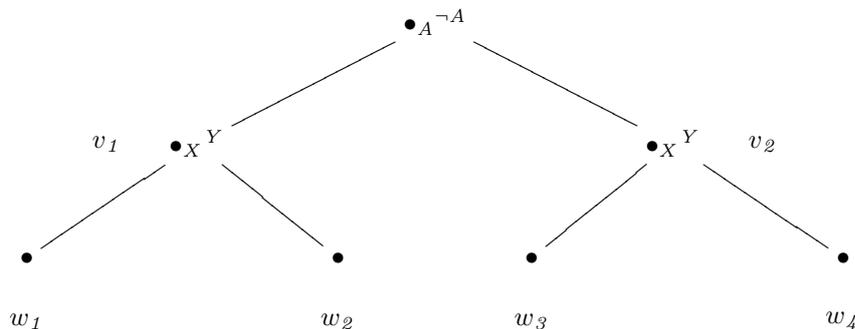
Figure 1: Sequential Game with Potential for Commitment

Credibility. The rationality constraint that received most attention in the economic literature is that a threat or promise should be *credible*. To punish or reward someone else can be *costly*. This is obvious if you promise to give some money, but the (future) consequences of punishing (e.g. killing) somebody if she doesn't perform the desired action can be very costly as well. As already observed by Schelling (1960, p. 177), from this it immediately follows that threats costs more when they fail, while promises cost more when they succeed. But given that effective threats and promises involve only conditional commitments, a major issue arises as to whether the threat or promise was *credible* (see **?**).

Above, we have suggested to model commitments as pruning of a game tree. But the idea that a formerly possible action is entirely excluded from a player's choice set just by public announcement that she will not choose to play so, is unrealistically strong. In reality, the option of playing Y after A, even after having said "if you do A, I will do X" remains. The problem is that as long as the option Y is still available in principle, it *might* still be chosen, and, more strongly even, it might even be rational to choose it.

Suppose that you made the conditional threat or promise to do X, if the hearer performed A. That means that you are committed to do X *after* the hearer performed A. But, irrespective of anger or gratitude, if you like doing Y more than doing X, i.e., if you prefer outcome w_2 over outcome w_1 in figure **??**, then it is (only) rational to deviate from your previous commitment. In case of a threat, harm is already done, and in case of a promise, you already have what you desired. Why would you stick to your commitment?

There is an obvious reason why you should carry out your commitment in these circumstances after all, if there is a good chance that you will be engaged with the other person in similar circumstances again in the future. Carrying out the commitment strengthens your *reputation*, while not carrying it out only destroys it. But if reputation cannot be brought into the picture, the only way in which threats and promises like in (**??**) can be credible is when they are (seen to be) *costless*. And this also makes intuitive sense. Take for instance the threat in (**??**). If the threatener is seen to be desperate or irrational enough,

one cannot rule out that he won't kill you if you don't give your wallet: perhaps he simply doesn't care about the possible consequences if things don't go as he desires. And indeed, to make your threat look credible, it is not unwise to act (as if you are) irresponsible.

It is clear that a lot can (and has been) theorized about credibility of inducements by conditional promises and threats from a game-theoretic point of view. For the purposes of this paper, however, it is sufficient to simply assume that promises and threats *are* credible: the involved obligations are binding. More precisely even, our analysis proceeds from the assumption that the speaker believes that her statements will be believed; it is inessential whether this belief is actually correct.

Benefit & Efficacy. Since we are concerned mostly with the speaker's perspective, we will concentrate, instead, on the other rationality constraints of promises and threats: that strategic commitments should be *beneficial* for the speaker (in expectation), and that they should be *efficacious* (in expectation).

Take a conditional promise like "if you do A, I will reward you with R" with which the speaker commits herself to a reward R after the hearer has performed A. Naturally, the *benefit* for the speaker of having A performed should exceed the speaker's detrimental cost of paying the reward R in order to count as a *rational* inducement. (This also entails that the speaker should not, of course, promise to reward performance of an action that the speaker would prefer *not* to have performed.) It is also clear that a strategic promise is only *efficacious* to the extent that the promised reward R is a sufficient incentive for the hearer to perform an otherwise dispreferred action A. So, while for the speaker the danger of having to pay the cost of the reward R must not exceed the expected gain of the action A (benefit), for the hearer the gain of obtaining R must exceed the loss in performing A (efficacy). Thus, "if you do A, I will reward you with R" is a rational and effective promise just in case both the speaker and the hearer prefer $A \wedge R$ above $\neg A \wedge \neg R$.

Similar considerations apply to conditional threats. Uttering a threat "if you do $\neg A$, I will punish you with P" commits the speaker to punishing the hearer if he does not perform A. This is *beneficial* only if the chance that the hearer performs A and the associated benefit of that for the speaker exceeds the possible cost that P might have if the hearer does not perform A. In order to be *efficacious*, the danger of being punished by P when A is not performed must outweigh the loss of performing A. So, while for the speaker the danger of having to punish the hearer by P must not exceed the expected gain of having A performed (benefit), for the hearer the danger of being punished by P must exceed the loss of performing A (efficacy). Thus, "if you do $\neg A$, I will punish you with P" is a rational and effective threat just in case both the speaker and the hearer prefer $A \wedge \neg P$ above $\neg A \wedge P$.

An analogue story holds for disjunctive threats of the form "You do A or I will punish you with P." In this case, the speaker commits himself to punish the hearer who does not perform A with P. This is *beneficial* just in case the

speaker expects that the benefit of action A taking place exceeds the cost of performing P in case the hearer abstains from A. The threat is *efficacious* in case the hearer fears P more than the harm he expects from performing A. Thus, "You will do A or I will punish you with P" is a rational and effective threat just in case both the speaker and the hearer prefer $A \wedge \neg P$ above $\neg A \wedge P$.

All these arguments seem completely straightforward and intuitive. The problem is that a seemingly equally straightforward argument can be given for why, and when, disjunctive promises of the form "You do $\neg A$ or I will reward you with R" make sense. What would be wrong with a disjunctive promise, if both speaker and hearer would prefer $A \wedge R$ to $\neg A \wedge \neg R$? As far as our argument goes so far, there is nothing that would suggest that disjunctive promises are any different from a conditional promise. The disjunctive promise would be *beneficial* just in case the speaker's loss of A not being performed exceeds the cost of giving the reward R, and it would be *efficacious* iff the reward R would be bigger for the hearer than the cost of performing A. The problem is that a disjunctive promise of the form "You do $\neg A$ or I will reward you with R" has exactly the same preference structure as the conditional promise "If you do A, I will reward you with R." Still, the latter is acceptable, but the former is not (as a promise). The problem is *why*? What is the difference between a conditional and a conjunctive promise that makes the former a reasonable strategic commitment and thereby a feasible move in dialogue, but not the latter?

4   Disjunctive Promises are Risky Inducements

In response to this question, we suggest to scrutinize more carefully the speaker's expected utility by taking into account that the speaker cannot know for sure what preferences and beliefs the hearer has. Doing so, we propose that disjunctive promises like (??) are suboptimal inducements because they are the only inducements in (??) that contain a *risk* of inefficiency in the light of the speaker's uncertainty that cannot be compensated.

More concretely, the main idea, to be spelled out in this section, is that, even where semantically equivalent, there is a minor linguistic difference between a conditional $\neg A \rightarrow X$ and a disjunction $A \vee X$: the former mentions the possibility "$\neg A$", while the latter mentions the possibility "A." We will assume that mentioning raises salience, however slightly, and that mentioning will therefore increase the probability with which the speaker expects the mentioned possibility to be realized. This then affects the speaker's *expected utility* of choosing a given inducement. Clearly, for (??) and (??), mentioning a speaker-desirable option only increases the speaker's expected utility and is unproblematic. Alternatives (??) and (??), on the other hand, mention a speaker-undesirable option which slightly decreases expected utility. This puts the efficacy of these statements at risk, given that the speaker cannot be certain about the hearer's actual preferences and beliefs. However, a conditional threat like (??) can compensate for this risk by committing to a stronger punishment, which is cheap in expecta-

tion, as explained in section **??**. Committing to a stronger reward is *not* cheap in expectation, and therefore a disjunctive promise cannot compensate for the risk of inefficacy.

In order to spell out this idea, we first have to enlarge on the assumption that mentioning an alternative slightly raises the speaker's expectation of realization (sections **??** and **??**). Secondly, we have to spell out the structure of the speaker's uncertainty (section **??**) and how this all affects the expected utility of inducements (**??**) .

4.1   Mentioning in Disjunctions and Conditionals

The standard view of the effect of an assertion that φ (and its acceptance) is that it eliminates all possibilities from the common ground where φ is not true (**?**). But even before an assertion is accepted, the fact that the assertion was made and understood had already another effect on the common ground: the possibility that ϕ is or might become true is brought to the (joint) attention of the participants of the conversation, in particular, to the hearer (c.f. **??**). For simple assertions with the content "It is raining" this extra effect is negligible because it is a side-effect of the acceptance of the assertion anyway. For more complex assertions, however, this extra effect does not fall out as a consequence of the assertion by itself. To see that this is so also for disjunctions and conditionals look at the simple question-answer pairs in (**??**).

(4) Who (of John and Mary) came to the party?
   a. John.
   b. John or John and Mary.
   c. John, if not also Mary.

Although all three answers are semantically equivalent,[10] they certainly do not convey the same idea. Whereas (**??**) implicates that Mary did *not* come, answers (**??**) and (**??**) both implicate that it is also possible that Mary might have come together with John. The intuitive reason why is because otherwise the speaker would not have *mentioned* this latter possibility, despite the semantic equivalence with the answer in (**??**) (cf. **??**).

In general, from just eliminating those possibilities in the common ground where $\varphi \vee \psi$ or $\varphi \rightarrow \psi$ are true, the extra effect of bringing φ- and ψ-possibilities to the attention does not follow. In fact, it has been argued that a very important purpose (among others) of a disjunctive claim is to bring its disjuncts to the attention (c.f. **???**). Similarly, the antecedents of (indicative) conditionals $\varphi \rightarrow \psi$ are normally associated with a speaker presupposition that φ be possible (e.g. **?**). Taken together, whatever the concrete mechanisms at work, it is fair to say that, on top of their semantic meaning, it is important to the way that disjunctions and conditionals are processed in discourse that these constructions *mention*, or are *about*, certain states of affairs.

---

[10]This holds for logical disjunction and material implication, but also for other standard analyses of the conditional, such as (variably) strict implication.

4.2 Mentioning, Salience, and Priming

But now suppose that a conditional or a disjunction such as in (**??**) is uttered where the truth of the antecedent or first disjunct is under the control of the hearer. What effect does it have to mention an action under hearer control in a game-like setting where the speaker wants to induce a certain action in the hearer?

First of all, if the hearer has not been aware of it at all, then just mentioning an action will inevitably make him aware of it. To make the hearer aware of a possibility that the speaker does *not* want to be realized might therefore just be a very dumb move in conversation because it would put the wrong ideas into the hearer's head (**??**). In other words, if the hearer is (possibly) unaware of some action A that the speaker does not want to have performed, then it is, intuitively speaking, a deficient inducement strategy to mention it in the first place *unless* it is strongly and credibly discredited, such as by a threat of sanction or similar.

But what if an action is mentioned that the hearer is (most likely) already aware of? Even then, it may seem *prima facie* suboptimal to bring such a possibility to the hearer's attention in case the speaker doesn't want this possibility to be(come) true. Mere mentioning makes an option salient, and to increase the salience of a choice option for the hearer simply means that, from the speaker's point of view, the probability with which this option is chosen increases, even if only very slightly. Think of marketing and advertisement: you want your product name to be ubiquitous, you want it to be the first thing that comes to mind when consumers make a decision (c.f. **?**). But also in more on-the-spot decision making: the salience of choice options matters especially when these have a strong desire-raising component (compare the discussion of "incentive salience" by **?** and references therein).

Taken together, this seems to suggest that when trying to manipulate hearer behavior the speaker should not mention and thereby raise the salience of speaker-undesirable, yet hearer-desirable options without at the same time discrediting those. In other words, raising the salience of a choice option A stands a mild chance of priming the hearer into choosing A (if not actual, then at least in speaker expectation). In strategic inducements this increase in the probability that the hearer performs hearer-undesirable A leads to a decrease in expected utility. This decrease, intuitively speaking, must be compensated by highlighting negative aspects or consequences of A, or otherwise the strategic inducement might rather lead to performance of A rather than to abstinence from it. This is possible for threats that mention speaker-undesirables, but not for promises. It is this intuitive argument that the remainder of this section spells out in more detail.

4.3 Towards a More Realistic Game Model

We would like to compare in particular the benefit and efficacy of conditional and disjunctive promises and threats in the light of a sufficiently realistic game model that incorporates the speaker's natural uncertainty about the concrete
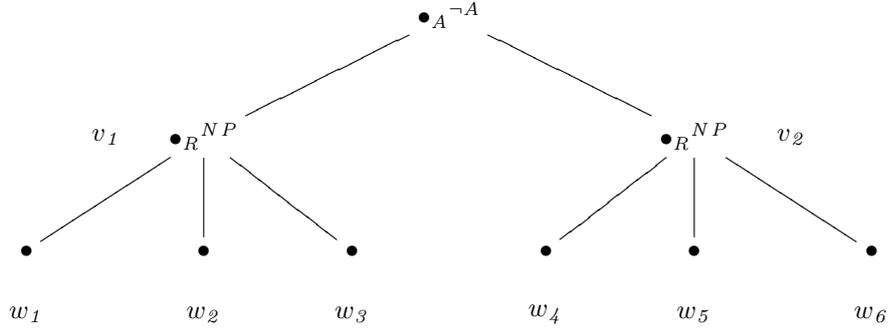
Figure 2: Game Model for Punishing and Rewarding Hearer Actions

preferences and beliefs of the hearer. Since performance of an action A is at stake, together with a potential reward or punishment thereafter, the minimal sequential game model we should consider is (something like) the one given in figure ??$laborate fig : SituationModel_E laborate the hearer's choice is between A and \neg A,$ ▪ and subsequently the speaker may choose to either reward the hearer (R), punish him (P) or stay neutral and abstain from both (N).

To be entirely precise, there should really be a set of rewards $\{R_1, R_2, \dots\}$ (and similar for punishments). The hearer would prefer some of the rewards over others and, in the simplest case, the more the hearer desires a reward, the more costly it is for the speaker to give. We would then have to ponder whether this set is reasonably infinite, and which properties the hearer's and speaker's preference-ordering have and how these are related to each other. We will sidesteps these details here and simply consider some reward R that is speaker-costly but hearer-desirable, and some punishment P that is also speaker-costly and hearer-undesirable.

More specifically, we will assume that the speaker's preferences are qualitatively of the following form:

$$w_2 > w_3, w_1 > w_5 > w_6, w_4.$$

This captures the intuition that the speaker mostly cares about whether action A is performed. Subordinate to her preference for A, she would prefer to remain neutral over punishing and rewarding. In contrast to that, we should assume that the hearer prefers $\neg A$ over A, but that the reward R and the punishment P that we consider are potentially efficacious, so that they outweigh the hearer's preference about A. More precisely, the hearer's preferences are then qualitatively given as:

$$w_4 > w_1 > w_5 > w_2 > w_6 > w_3.$$

In other words, the hearer (is assumed by the speaker) to value most the reward, and prefers a neutral outcome over a punishment. Subordinate to these preferences is his preference of performing $\neg A$ over performing A.

11

It is not necessary, but also not desirable to specify the hearer's preferences any further than that. This is because a speaker will never be able to know for sure how exactly the hearer will value a promise or a reward. Our modelling here adopts the speaker's perspective and takes her natural uncertainty into account. In other words, the model assumes that the speaker believes the hearer's preferences are qualitatively as specified above, but that the speaker does not know for certain how strongly, for example, w_4 is preferred over w_1.

Similar remarks then also apply to the speaker's beliefs about the hearer's beliefs. Here it is most natural to suppose that the speaker believes that the hearer expects, all else being equal, a neutral outcome. We could even go as far as saying that the hearer might not even be *aware* of possible punishments and threats and that it is only when pointed out to him that he accommodates these possibilities into his decision-making. To keep matters simple here, we will refrain from representing such a sequential game with possibly unaware players (c.f **??**). For the present purpose it suffices to assume that the speaker believes that the hearer's beliefs are qualitatively as follows:

$$w_2, w_5 \gg w_1, w_3, w_4, w_5.$$

The idea is that the speaker again does not know precisely which probabilistic beliefs the hearer holds, but she does believe that, barring any speaker commitment, the hearer considers action N *substantially* more likely than either reward or punishment.

4.4 Risk of Strategic Inducements

Given the speaker's natural uncertainty about the hearer's precise preferences and beliefs, it turns out that disjunctive promises are risky, and therefore suboptimal in expectation, in a sense that threats and conditional promises are not. To see what is at stake, we need to compare the statements in (**??**) one by one as committing strategic inducements against the background of the speaker's uncertainty as described in the previous section. Let us look at threats and promises in turn and let us ask what update effects these statements would have on the hearer and how this affects the speaker's assessment of her expected utility of uttering these statements.

As for their semantic update effect, both the conditional threat $\neg A \rightarrow P$ in (**??**) and the disjunctive threat $A \vee P$ in (**??**) are semantically equivalent and denote, if taken as binding, the set $\{w_1, w_2, w_3, w_6\}$. If we also take conditional perfection, respectively exclusive readings of disjunctions, into account the impact of these threats is an update that leaves only outcomes $\{w_2, w_6\}$.

Still, in line with our reasoning above, there should be a small difference between the conditional and the disjunctive threat. Whereas the conditional threat slightly increases the probability of w_6 (in the expectation of the speaker), the disjunctive threat slightly increases the probability of w_2. This is because the conditional mentions "$\neg A$" and so the speaker will assume a slight increase in the chance that the hearer will play this option. For the disjunctive threat rather

the hearer choice A is given a slightly higher probability. That means that mentioning the speaker-desirable action A in the disjunctive threat actually has a slight increasing effect on the expected utility of that statement, as compared to the conditional threat that mentions "$\neg A$." However, this slightly detrimental effect of mentioning the speaker-undesirable action is relatively harmless, because the speaker believes that the hearer prefers w_2 over w_6 *and* the speaker believes that the hearer considers w_2 much more likely than w_6. It is therefore not likely that the conditional threat would not be efficacious despite the fact that it might slightly increase the chance of performance of $\neg A$.

Moreover, and more importantly, the speaker can compensate for the risk of a conditional threat by choosing a sufficiently stronger punishment. As noted in section **??**, this need not decrease the speaker's expected utility, because threats are (relatively) cheap in expectation in that the (possibly costly) punishment is not (as) socially binding as in the case of a promise.

This is different for disjunctive promises. Again, the conditional promise $A \rightarrow R$ in (**??**) and the disjunctive promise $\neg A \vee R$ in (**??**) are semantically equivalent. Their update effect is to eliminate outcomes w_2 and w_3, and additionally, if perfection and exclusive readings are taken into account, restrict the options under consideration to $\{w_1, w_5\}$. Once more, we also attest a difference from mentioning different alternatives: whereas the conditional promise slightly increases the speaker's expected utility because it mentions the desirable option A and thus increases the probability that w_1 is realized, the disjunctive promise slightly decreases the expected utility by mentioning the undesirable option $\neg A$ and thereby increasing the probability of w_5. However, unlike with threats, this latter decrease is more *risky* from the point of view of an uncertain speaker: the problem is that although w_1 is assumed more hearer-desirable than w_5, the latter is naturally assumed substantially more likely. If the speaker is uncertain about the extent to which the hearer prefers w_1 over w_5, mentioning the undesirable option puts the efficacy of the inducement at risk.

But could the speaker not compensate this risk, as she could with conditional threats, by promising a higher reward, so as to make sure that w_1 is sufficiently preferred over w_5? She probably could, but not necessarily without sacrificing even more on expected utility. Promises are costly when efficacious and more institutionally dependent than threats. To the extent that the speaker would like to invest on the promise to compensate risk, the statement's expected utility decreases, because any stronger reward would only be more costly and thus further decrease the speaker's expected utility. This is then the main differences between conditional threats and disjunctive promises: although both mention a speaker-undesirable option, which is risky under uncertainty, threats, but not promises, can be "pumped up cheaply", so to speak, to compensate for the risk.

Taken together, we argue that disjunctive promises are suboptimal, because they emphasize the wrong alternative and cannot compensate for any negative effects that this might have. It is the combination of natural speaker uncertainty, priming by mentioning and the asymmetry of when punishments and rewards are speaker-costly that explains why disjunctive promises are a deficient inducement strategy.

5 Conclusion

In this paper we explained promises and threats as strategic commitments of the speaker. We discussed the constraints that such commitments have to obey to be rational. We saw that from the speaker's point of view, a threat is cheaper in expectation than a promise. The main idea presented in this paper was that disjunctions like (??) preferably get a threat-reading because their use as a disjunctive promise is a risky, and therefore suboptimal strategic inducement. This explanation involved the idea that mere mentioning a possibility raises its salience. This helps to explain the difference in acceptability between the conditional promise $A \to R$ and the disjunctive promise $\neg A \vee R$. The idea that a threat is cheaper in expectation than a promise, on the other hand, explains why the disjunctive promise is also more risky than the conditional threat $\neg A \to P$.

This explanatory strategy certainly raises a number of concerns. Perhaps the most pressing is the question whether the alleged suboptimality of a disjunctive promise is something that is checked on-the-spot, every time anew a speaker would like to influence a hearer. We emphatically do not subscribe to this obviously nonsensical view. Rather we suggest here that the suboptimality of disjunctive promises is a force that informs language organization, not *ad hoc* choice of formulation. Certain locutions and grammatical constructions, and not others, are conveniently and conventionally used for certain discourse functions, and not others. It is at this level of functional organization that, we suggest, evolutionary pressures have weeded out disjunctions as a viable vehicle of making promises.