Evaluating conversational success: Weighted Message Exchange Games

Nicholas Asher and **Soumya Paul** *
IRIT, Université Paul Sabatier
118 Route de Narbonne
31062 Toulouse, France
nicholas.asher@irit.fr, soumya.paul@gmail.com

Abstract

We analyze evaluations of conversational success and how such evaluations relate to notions of discourse content and structure. To do so, we extend the framework of Message Exchange (ME) games by adding weights or scores to the players' moves and then accumulating these weights using discounting to evaluate a conversationalist's performance. We illustrate our analysis on a fragment of a recent political debate.

1 Introduction

As is by now well accepted, a discourse is more than an unstructured set of utterances; these utterances should, for example, be related to one another in a coherent fashion. But in general, not just any coherent arrangement of utterances will do. If one's goal is merely to avoid an awkward silence, then maintaining conversational coherence might suffice to achieve one's ends, but conversational goals are frequently more ambitious than this. Sometimes interlocutors converse to get to the truth of a matter; other times, a speaker says what she does to convince her interlocutor or a third party, an observer, to do something or to adopt a certain belief; in the latter case, the truth of what she says might be less important than its persuasiveness. One might win a political debate, for instance, even if the majority of the claims one asserts in that debate are false, as the 2016 series of debates between Republican candidates for the U.S. Presidency illustrates.

With (Grice, 1975), we hold that conversations are rational activities, and that agents act so as to maximize their conversational success. But in order for that to be possible, conversational agents, and observers, must be able to evaluate conversations for such success, and this requires moving beyond evaluations of discourse content in terms of truth or satisfaction. In particular, we want to know how the linguistic and discourse structure and content of a speaker's contributions affect that evaluation. In this paper, we propose a model of context-sensitive *evaluations of conversational success* and investigate how such evaluations relate to notions of discourse content and structure. In our view, a better understanding of conversational success will shed light on how agents structure their contributions and how these contributions affect the overall shape and content of the conversation.

Conversational success need not be shared by all members of a conversation; speakers can have different and even opposed conversational goals. We thus develop our model of conversational success using the framework of Message Exchange (ME) Games (Asher et al., 2016), in which a conversation is understood as a sequential, extended game that does not require interlocutors to share interests or goals. To avoid troublesome backwards induction results that predict that no conversation takes place in cases of opposed interests (Crawford and Sobel, 1982; Asher et al., 2016), an ME game conceives of a conversation without a commonly known, set end and thus models conversations as infinitary games. Such games are evaluated by a *Jury*. Intuitively, a Jury is any entity or a group of entities that evaluates a conversation and decides the winner. For example, in a courtroom situation, the Jury is the courtroom Jury itself whereas in a political debate, the Jury is the audience of the debate which maybe the entire citizenry of a country. A Jury can even be one of the participants of the conversation itself. Thus, a Jury for a particular conversation setup depends entirely on the context. But given such a setup, it is always clear who or what constitutes the Jury. We formalize the Jury here as a *weight function* or *scoring function* over the sequence of

*The authors thank ERC grant 269427 for research support.

Appears in *JerSem: The 20th Workshop on the Semantics and Pragmatics of Dialogue,* edited by Julie Hunter, Mandy Simons and Matthew Stone. New Brunswick, NJ USA, July 16-18, 2016, pages 55-64.

conversational moves. To accumulate the individual weights to obtain a global score of a conversation for the players, we will use techniques of discounting (Shapley, 1953).

To motivate these decisions in our analysis ofconversational success, consider a recent example from the U.S. Republican primary debates (February 6, 2016) where things go dramatically wrong for a candidate Marco Rubio (R), the junior US senator from Florida. The crucial episode can be viewed at (Christie-Rubio-debate, 2016), and the transcript at (Christie-Rubio-transcript, 2016).

We describe the relevant part of that conversation below where the numbers correspond to blocks of sequential discourse moves making up a coherent unit. In terms of the linguistic theory SDRT, these blocks correspond to *complex discourse units* or CDUs (Asher and Lascarides, 2003). A CDU is a structure consisting of elementary discourse units (typically clauses) that are linked together by discourse relations and, crucially, that bear together some rhetorical relation to another discourse unit. For example, the block (3) below in an SDRT analysis would yield a CDU consisting of several EDUs; the first sentence (a) yields an EDU that is elaborated on by the EDU derived from (b), with the (c) and (d) elaborating on (b). The division of the conversation into CDUs and their numbering will help us in carrying out a detailed analysis in Section 4.

Fielding a question about his experience to be president given that he is a very junior US senator, R initially responds with (1) a summary of his record in the Senate, (2) a short argument that experience isn't sufficient for being President and then concludes (3) by drawing a comparison between himself and Obama, who, like R, had only one term of political experience at the national level before running for President:

(3) *"(a) And let's dispel once and for all with this fiction that Barack Obama doesn't know what he's doing. (b) He knows exactly what he's doing. (c) Barack Obama is undertaking a systematic effort to change this country, (d) to make America more like the rest of the world."*

(3) is a coherent move when R utters it. The question to which R is responding carries with it an implicit argument against him. The major premise of that argument is that no one who has had only one term of legislative experience could be a President who *"knows what he's doing."* R argues that Obama was very effective, thus challenging this premise.

The floor then goes to Governor Christie (C) of New Jersey, who takes issue with R's response and attacks his record in the Senate (4) and picks up the comparison to Obama (5). R responds by attacking C's record as governor (6), which is a natural move. But then something strange happens: in (7), R goes back and repeats (3) almost verbatim:

(7) *"But I would add this. Let's dispel with this fiction that Barack Obama doesn't know what he's doing. He knows exactly what he's doing. He is trying to change this country. He wants America to become more like the rest of the world."*

C then characterizes R's response in an extremely damaging way:

(8) *"That's what Washington, D.C. Does. The drive-by shot at the beginning with incorrect and incomplete information and then the memorized 25-second speech that is exactly what his advisers gave him."*

The debate continues with R again attacking C's record (9). Had R stuck to this strategy, he might have recovered from his faux pas repetition; but instead, he goes back and repeats in block (10) the material in (3) and (7) without any attempt to respond to C's characterization of the repetition in (7). In block (11) C once again points out the *"memorized text"* to R's detriment. The effect of this repetition and his failure to counter C's negative characterization of it was disastrous for R as pundits claimed and subsequent polls confirmed; C's characterization gave a label for R's *"robotic performance,"* and the video in (Christie-Rubio-debate, 2016) went viral.

While prior work on a conversationalist's success or 'power status' has focused on superficial features like the number of turns the speaker has, the length of time she has spoken, or word bigrams (Prabhakaran et al., 2012; Prabhakaran et al., 2013), examples like the Rubio gaffe show that a dialogue participant's success in meeting her conversational objectives depends upon the individual moves that she makes *in the particular dialogue context*. When pundits and the public evaluated the debate performance of the candidates, they justified their evaluations by making reference to particular moves in the debate, including R's 'robotic' repetitions. Had R simply given (3) in his response to the moderator's leading question, the response would have been fine. But same message (e.g. (7) and (10)) in a different context (e.g. following (4) and (5) and then (8)) gets a very different and bad score. Further, R's 'robotic' response affects the evaluation of the rest of the conversation, penalizing his subsequent performance.

To model evaluations of conversational success, we need to answer three questions: (a) how do we characterize the context upon which the evaluation is based? (b) in virtue of what does one give such an evaluation? (c) how does the evaluation proceed? Given our characterization of Rubio's performance, evaluators are sensitive to the exact words used, to the conversational string, but they also evaluate whether a particular discourse move or sequence of moves performs a coherent rhetorical role, like answering a question, amplifying on a response to a question, rebutting a prior attack move by another participant, and so on.

With respect to question b, evaluators exploit criteria like responsiveness and coherence, taking, e.g., an attack on an agent i to which i has no coherent rebuttal to contribute to a negative evaluation of a response given by i. Evaluation of conversational success also depends, however, on what is needed to persuade the evaluator that an agent has been successful. This may depend upon the agent's own global goals like defending a particular position, but it may also depend upon the evaluator's preconception of what a successful conversation for i would be.

Finally, to answer question c, a global evaluation of Player i's contributions depends on the contributions she makes on each of her turns and how they are related to the discourse context. The evaluation of i's performance in the conversation should be a function of the evaluation she receives on each turn. We examine a normalized, additive function that assigns to each turn for every debater i a score in $\{0, 1, \ldots, d\}$ where d is a positive integer. However, a bad evaluation on one turn like that of Rubio's (or 1988 Vice-Presidential candidate Quayle's famous gaffe (Asher and Paul, 2013)) colors the evaluation of further turns, and several bad evaluations can doom the entire conversation by heavily 'discounting' the value of future moves.

The rest of the paper is organized as follows. Section 2.1 introduces *weighted ME games*—that is, ME games with *weights* or *scores* for each move of a play. The weights are accumulated over the entire play by the method of discounting. Section 2.2 extensively discusses a discounting factor to account for the penalties that the speakers incur from making disastrous discourse moves. As we show in Section 3, the discounting factor entails the existence of ϵ-Nash equilibria for weighted ME games, meaning that a notion of optimal rational play exists for our games. Section 4 applies our notion of weighted ME games to the Rubio/Christie exchange, while Section 5 considers related work. Section 6 concludes the paper.

2 The model

In this section we introduce Weighted Message Exchange games and formulate a discounting mechanism to accumulate the weights of the moves along a play.

2.1 ME and WME games

Definition 1 (ME game (Asher et al., 2016)) *A Message Exchange game (ME game) is a tuple* $\mathcal{G} = ((V_0 \cup V_1)^\omega, Win_0, Win_1)$ *with* $Win_0, Win_1 \subseteq (V_0 \cup V_1)^\omega$.

V_0 and V_1 are called the *vocabularies* of players 0 and 1 respectively. The intuitive idea behind an ME game is that a conversation proceeds in turns where in each turn one of the players 'speaks' or plays a string of letters from her own vocabulary. However, the player does not speak any garbled sequence of strings but sentences or sets of sentences that 'make sense'. We capture by setting V_0 and V_1 to be SDRSs (Asher and Lascarides, 2003). See (Asher et al., 2016) for a detailed discussion on this topic and the motivation behind the formal setting of ME games.

Formally the ME game $\mathcal{G}$ is played as follows. Player 0 starts the game by playing a non-empty sequence in V_0^+. The turn then moves to Player 1 who plays a non-empty sequence from V_1^+. The turn then goes back to Player 0 and so on. The game generates a play p_n after n (≥ 0) turns, where by convention, $p_0 = \epsilon$ (the empty move). A play can potentially go on forever generating an infinite play p_ω, or more simply p. Plays are segmented into *rounds*—a move by Player 0 followed by a move by Player 1. A finite play of an ME game is (also) called a history, and is denoted by h. Let Z be the set of all such histories, $Z \subseteq (V_0 \cup V_1)^*$, where $\epsilon \in Z$ is the empty history and where a history of the form $(V_0 \cup V_1)^+ V_0^+$ is a 0-history and one of the form $(V_0 \cup V_1)^+ V_1^+$ is a 1-history. We denote the set of 0-histories (1-histories) by Z_0 (Z_1). Thus $Z = Z_0 \cup Z_1$. For $h \in Z$, $\mathsf{turns}(h)$ denotes the total number of turns (by either player) in h.

We are interested in an extension of ME games where a *Jury* assigns a non-negative integer *weight* or *score* to every move by each player. The Jury then accumulates these weights in a way it deems suitable to compute the global score of the play for each player. In what follows, unless otherwise mentioned, i will range over the set of players, here $\{0, 1\}$. Thus, Player $(1 - i)$ denotes Player i's opponent.

Let $\mathbb{Z}$ be the set of all integers and $\mathbb{Z}_+$ be the set of non-negative integers. For any $n \in \mathbb{Z}_+$ let $[n] = [0, n-1] \cap \mathbb{Z}_+ = \{0, 1, \ldots, n-1\}$. A weight function is a function $w : (Z_0 \times V_1^+ \cup Z_1 \times V_0^+) \to \mathbb{Z} \times \mathbb{Z}$. Intuitively, given a history $h \in Z$, w assigns a tuple of integers $(a_0, a_1) = w(h, x)$ to the next legal move x of the play h. Note that the weight function, w depends on the current history of the game in that, given two different histories $h_1, h_2 \in Z$, it might be the case that $w(h_1, x) \neq w(h_2, x)$ for the same continuing move x. For notational simplicity, in what follows, given a play $p = x_0 x_1 \ldots$ of $\mathcal{G}$, we shall denote by $w_i^n(p)$, the weight assigned by w to Player i in the nth turn of p ($n \geq 1$). That is, if $w(p_{n-1}, x_n) = (a_0, a_1)$, then $w_0^n(p) = a_0$ and $w_1^n(p) = a_1$

Definition 2 (WME game) *A weighted ME game (WME game) is a tuple $\mathcal{G} = ((V_0 \cup V_1)^\omega, w)$ where w is a weight function.*

In Section 3, We will formally define a *Jury* who assigns weights to the moves of the game in a play p and accumulates them in a way it deems suitable to have a global evaluation of p for both the players. One of the standard methods for performing such an accumulation is 'discounting' (Shapley, 1953). In discounting, along a play p, the immediate moves are assigned high values and the moves further and further into the future are assigned lower and lower values. This is achieved by multiplying the weight of every subsequent move by a factor λ, which is usually fixed to be a constant between 0 and 1. However, in our case, to capture the context dependence of evaluations, we shall set λ to be a function of the history h, $\lambda : Z \to (0, 1)$.

Before fixing λ, we define first the discounted weight of a play and a discounted WME game.

Definition 3 (Discounted-payoff) *Let p be a play of $\mathcal{G}$ and let λ be a discounting function. Then the discounted-payoff of p for Player i is given by*

$$w_i^D(p) = \sum_{n \geq 1} \lambda(p_{n-1})^{n-1} w_i^n(p)$$

Definition 4 (Discounted WME game) *Let w be a weight function and λ be a discounting function. A discounted WME game with discount λ is a tuple $\mathcal{G}_D[\lambda] = ((V_0 \cup V_1)^\omega, w)$ such that for every play p, Player i receives a payoff of $w_i^D(p)$.*

When λ is clear from the context, we shall simply write $\mathcal{G}_D$ instead of $\mathcal{G}_D[\lambda]$. A (pure) strategy σ_i for Player i is defined in the standard way, $\sigma_i : Z_{1-i} \to V_i^+$. A play $p = x_0 x_1 x_2 \ldots$ conforms to a strategy σ_i of Player i if she always plays according to σ_i in p, that is, for every $j > 0$, $j - 1 = i \pmod 2$ implies $x_j = \sigma_i(p_{j-1})$. We denote by $p_{(\sigma_0, \sigma_1)}$ the unique play conforming to the tuple of strategies (σ_0, σ_1).

Definition 5 (Best-response / Nash-equilibrium) *A strategy σ_i of Player i is a best-response to a strategy σ_{1-i} of Player $(1 - i)$ if for every other strategy σ'_i of Player i, we have*

$$w_i^D(p_{(\sigma_i, \sigma_{1-i})}) \geq w_i^D(p_{(\sigma'_i, \sigma_{1-i})})$$

Given $\epsilon > 0$, σ_i is an ϵ-best-response to σ_{1-i} if for every other strategy σ'_i of Player i, we have

$$w_i^D(p_{(\sigma_i, \sigma_{1-i})}) \geq w_i^D(p_{(\sigma'_i, \sigma_{1-i})}) - \epsilon$$

A tuple of strategies (σ_0, σ_1) is a Nash equilibrium (resp. ϵ-Nash equilibrium) if σ_0 and σ_1 are mutual best-responses (resp. ϵ-best-responses).

We can also define natural notions of a *win*, *winning-strategy* etc. as follows, for both zero sum and non-zero sum games.

Definition 6 (Winning and winning strategy) *Let $\mathcal{G}_D[\lambda] = ((V_0 \cup V_1)^\omega, w)$ be a discounted WME game. Then (i)* **Zero-sum:** *Player i wins a play p of $\mathcal{G}_D[\lambda]$ if $w_i^D(p) \geq w_{1-i}^D(p)$. Player $(1 - i)$ wins p otherwise. (ii)* **Non-zero sum:** *Fix constants $\nu_i \in \mathbb{R}$ called 'thresholds'. Then Player i wins a play p if $w_i^D(p) \geq \nu_i$. (iii) A strategy σ_i is winning for Player i if she wins all plays p conforming to σ_i.*

2.2 The discounting factor

We now fix the exact form of the discounting factor λ to suit evaluations of conversational success. We assume that w is both integral and bounded, that is, the range of w is $[d]$ for some constant $d \in \mathbb{Z}_+$. A move with a weight of '0' is a 'failure' or a 'disastrous move' and heavily penalizes a player's future play. Also a move that gets weight 'd' is a 'brilliant move'; if such a move follows a disastrous move then it is a 'recovery move'.

For any history h, the function λ consists of two terms

$$\lambda(h) = \lambda_1 \lambda_2^{\frac{\mathsf{rec}_i(h)}{\mathsf{turns}(h)-1}}$$

The first is the global discounting which weighs initial moves more than later ones. This reflects the intuition: "get your best licks in first" - the player who does better initially often has an upper hand throughout the course of the debate. The second term is the 'punishing factor' that heavily discounts disastrous moves of a player. It 'kicks in' after the first disastrous move made by the player and gets worse if she keeps making such moves. A player may also recover from a disastrous move by making a number of brilliant moves, after which the punishing factor disappears, but might kick in again in the future. $\mathsf{rec}_i(h)$ is thus the 'recovery index' of Player i at history h and is computed using Algorithm 1 [note that the denominator of $(\mathsf{turns}(h) - 1)$ occurs in the index of λ_2 so that the number of turns does not affect it like it does for the global discounting λ_1].

Algorithm 1: REC$_i(h)$

data:h; **result:**$\mathsf{rec}_i(h)$
let `rec_i` = 0; `good` = 0
for `j=1` **to** $\mathsf{turns}(h)$ **do**
  if $w_i^j(h) = 0$ **then** `rec_i++`
  if `rec_i=0` **then** `good=0`
  if `rec_i > 0` **then**
    if $w_i^j(h) = d$ **then** `good++`
  if `good=c` **then** `rec_i--; good=0`
return `rec_i`

Intuitively, Algorithm 1 starts accumulating the number of disastrous moves occurred. If Player i plays 'c' recovery moves after having played one or more disastrous move, the accumulated count of the disastrous moves decreases by 1. If i has fully recovered, it stops keeping track of the brilliant moves. The process repeats when i plays a disastrous move again.

3 Finite satisfiability and the Jury

We can now formalize the notion of the Jury. The Jury fixes the weights of the moves of the Players and also the parameters of the discounting function λ. That is, it fixes λ_1, λ_2 and c. Thus

Definition 7 (Jury) *The Jury for a discounted WME game $\mathcal{G}_D$ is a tuple $\mathcal{J} = (w, \lambda_1, \lambda_2, c)$ where w is a weight function.*

Although the game $\mathcal{G}_D$ can potentially go on forever, the Jury has to decide the winner after a finite number of turns. We can compute a bound on the number of turns after which the Jury can confidently decide the winner of the game. This is facilitated by the discounting of the weights and also the fact that w is integral and bounded. We have

Proposition 1 *Fix a discounted WME game $\mathcal{G}_D$ with a Jury $\mathcal{J} = (w, \lambda_1, \lambda_2)$ such that the range of w is $[d]$. Then given $\epsilon > 0$ we have for Player i and any play p of $\mathcal{G}_D$*

$$w_i^D(p) \leq \sum_{j=1}^{n_\epsilon} \lambda(p_{j-1})^{j-1} w_i^j(p) + \epsilon$$

where $n_\epsilon \leq \frac{\ln[\frac{\epsilon}{d}(1-\lambda_1)]}{\ln \lambda_1} - 1$.

Proof Suppose Player i does not play any disastrous move after n_ϵ turns. The maximum payoff she can gain after n_ϵ turns is $\lambda_1^{n_\epsilon+1} \frac{1}{1-\lambda_1} d$. Setting

$$\lambda_1^{n_\epsilon+1} \frac{1}{1-\lambda_1} d \leq \epsilon$$

we have $n_\epsilon \leq \frac{\ln[\frac{\epsilon}{d}(1-\lambda_1)]}{\ln \lambda_1} - 1$. ∎

Thus, if the Jury stops the game after n_ϵ turns, they can be sure no player would have gained more than ϵ, had the game been allowed to continue forever. Note that this result is fully general, but that values for n_ϵ will very much depend on the values set for λ_1 and λ_2.

Remark Note that it is crucial to assume that the players are unaware of the parameters of the Jury, w, λ_1, λ_2 and c. Otherwise, they can compute n_ϵ on their own. The game then becomes equivalent to a finite extensive form game with a set end, which is against the view on modeling strategic conversations defended in (Asher et al., 2016) that we have adopted. Thus, although the Jury takes a decision on the outcome of the game after a finite number of turns, the players do not know when that decision takes place. Thus, the game still appears to the players as potentially unbounded.

From Proposition 1, it also follows that ϵ-Nash equilibria always exist in our discounted WME games in pure strategies. However, since our space of strategies is uncountably infinite, the existence of Nash equilibria is a delicate matter (see for e.g. (Levy, 2013)) and we intend to explore it further in future work.

Corollary 1 *Given $\epsilon > 0$, a discounted WME game always has an ϵ-Nash equilibrium.*

Proof Consider the 'finite' discounted WME game for n_ϵ turns where n_ϵ is given by Proposition 1. Define the relation $\sim$ on plays of n_ϵ turns as: for two plays p and p', $p \sim p'$ iff for all $j : 1 \leq j \leq n_\epsilon$, $w_i^j(p) = w_i^j(p')$ and $w_{1-i}^j(p) = w_{1-i}^j(p')$. Clearly, $\sim$ is an equivalence relation. Also, since w is integral and bounded, there are only a finitely many possibilities for the weights of each Player i along any play p, and thus $\sim$ has finitely many equivalence classes. Thus there is a finite number of discounted payoffs possible (one for each equivalence class of $\sim$) after n_ϵ turns. A backward induction procedure on the equivalence classes of $\sim$ gives an ϵ-Nash equilibrium tuple of strategies $([\sigma_0], [\sigma_1])$ on these classes. Indeed, since by Proposition 1, no player can gain more than ϵ by deviating from it. Lifting $[\sigma_0]$ and $[\sigma_1]$ to corresponding representative elements of functions over actual histories gives us a required ϵ-Nash equilibrium (σ_0, σ_1). ∎

4 Applications

In Section 2, we developed weighting functions with two discounting parameters, λ_1 and λ_2 and a recovery constant c. λ_1 discounts future moves in the standard way agreeing with our intuition that good moves carry more value if played earlier than later. λ_2 is particular to WME games, that derives from agents' bad moves a penalty that adversely affects their score. c represents the number of brilliant moves required by a player to recover from a single disastrous move. These parameters are decided by the Jury. In this section we examine an WME game evaluation of our example dialogue, framed by the question as to whether Rubio has the experience to be president to be a dialogue on its own. The exchange is rather lengthy from the perspective of giving a complete discourse structure in which each clause is linked to other clauses via one or more rhetorical relations; this particular part of the political debate has over 200 clauses or elementary discourse unit (EDU). However, SDRT groups EDUs into more complex units or CDUs, small discourse graphs on their own that also have rhetorical links to other discourse units (Asher and Lascarides, 2003). As coherence is assured amongst the EDUs within the blocks, we will look only at the organization of CDUs and their relation to the whole dialogue, for it is there where the Jury has an important effect.

Our example is a fragment of a zero sum WME game. Let us denote the actual debate that unfolded between Rubio (R) and Christie (C), which is a play of the above game, as p_{RC}. Rubio's goal is to provide a convincing answer to the moderator's (M) question: to convince the public that he has the experience to be President. The goal of the antagonist, here C, is to destroy that answer, and C is very effective in doing that. Let us see how.

To do so, we will examine the role of the CDU blocks of the debate, which we've numbered in the introduction as (1)-(11), in the context of the Jury which is here the audience in the debate. For the sake of concreteness, we will take a particular integer scale and discount values for the weighting scheme; we feel that the scheme is defensible, though we acknowledge that there are many weighting schemes to choose from and we are unsure at this point exactly how to determine optimal weighting schemes or even whether such exits. We will also leave the tie between the details of the discourse structure and the weighting scheme relatively programmatic for now, as we have not fully figured out at present all the parameters of variation in this relation. Based on the Jury's evaluations and its applause reactions, we fix the range of w to be $[5] \times [5]$. We also fix $\lambda_1 = 0.9$ and $\lambda_2 = 0.5$. Thus the global discounting λ_1 is more or less gradual whereas the penalty discounting λ_2 for disastrous moves by either player is pretty severe. Let us also assume that the recovery constant $c = 5$. As we will show, these values fit the facts of the conversational sequences we have analyzed.

After the CDU introducing the question of political experience to R, R's response has 3 CDUS: 1) he talks about his record; 2) he argues that years of experience is not sufficient; years of experience aren't necessary either; 3) Obama with little experience knows exactly what he's doing (not necessary). We'll call (3) the *Obama CDU*. This seems to be a perfectly adequate response; it is responsive to the question and internally coherent. The audience applauds politely, and we could fix $w(\epsilon, \langle 1\rangle\langle 2\rangle\langle 3\rangle) = (3,1)$. That is R (Player 0) gets a score of 3 for his points 1,2 and 3 which satisfactory but not overwhelming and C (Player 1) reaps only a minimal reward of 1 at this stage.

The moderator then invites C to comment on R's prior response. 4) C mounts a direct attack on R's record. 5) C also picks up on R's reference Obama but uses Obama as an example of disastrous government on the part of an inexperience one time senator, which indirectly attacks R as well. There are two points at which the audience applauds so we might set $w(\langle 1\rangle\langle 2\rangle\langle 3\rangle, \langle 4\rangle\langle 5\rangle) = (1,4)$. C has a forceful reply and R gains only minimally from C's response.

Now R in (6) briefly responds with an attack on C's record as a problem solver but then in (7) returns to the Obama CDU. The problem is that the Obama CDU does not cohere with (6). R flubs the connection between the attack by implicating contrast (*"but let me add this"*), when he should have made an explicit reference back to C's use of Obama's record. While the point could have been effective, it wasn't rhetorically crafted in the right way, and the Obama CDU seems just to hang there, in addition to (7)'s being an almost verbatim repetition of (3). We could even imagine that C actually gains from R's dubious move. So here we let $w(\langle 1\rangle\langle 2\rangle\langle 3\rangle\langle 4\rangle\langle 5\rangle, \langle 6\rangle\langle 7\rangle)) = (1,2)$. This inept response nevertheless does not kick in the penalty discount λ_2 for R yet, as λ_2 only makes a difference if there are moves evaluated with 0.

R's inept rhetorical connection and reuse of the Obama CDU gives C a crucial opening; C characterizes R's attack and the incoherently linked Obama CDU in a devastating way in (8). That is, (8) has the rhetorical function of commenting on the Obama CDU, not its content but its representation. With (8), C provides an evaluation of R's turn that capitalizes on its inept rhetorical structure. The audience sees the aptness of the characterization and roars its approval. Their evaluation coincides with C's, which means: $w(\langle 1\rangle\langle 2\rangle\langle 3\rangle\langle 4\rangle\langle 5\rangle\langle 6\rangle\langle 7\rangle, \langle 8\rangle) = (0,5)$.

λ_2 now kicks in and since it is relatively low (0.5), R would have to do very well for the rest of the debate while C has to do very badly in order for R to win. We do allow that a long sequence of very good moves resets λ_2, but this seems to happen rarely. Actually, things get worse for R. In (10) R starts to deliver the Obama CDU again. Given (8), we can set $w(\langle 1\rangle\langle 2\rangle\langle 3\rangle\langle 4\rangle\langle 5\rangle\langle 6\rangle\langle 7\rangle\langle 8\rangle, \langle 9\rangle\langle 10\rangle)) = (0,5)$, that is, it is a disastrous move for R while C's reputation is not hampered in any way. Moreover, C in (11) reuses his characterization again on R's contribution in (10), making $w(\langle 1\rangle\langle 2\rangle\langle 3\rangle\langle 4\rangle\langle 5\rangle\langle 6\rangle\langle 7\rangle\langle 8\rangle\langle 9\rangle\langle 10\rangle, \langle 11\rangle)) = (0,5)$. At this point the contribution of the penalty discount, λ_2, is cubed ($= 0.125$), which is terrible for R. This makes C's characterization of his performance stick and affects the audience's (Jury) evaluations for the rest of R's turns.

We can now compute the discounted payoff to R and C respectively after these 3 rounds of p_{RC} as:

$$R : 3 + (0.95)\cdot 1 + (0.9)^2 \cdot 1 + (0.9)^3(0.5)\cdot 0 + (0.9)^4(0.5)^2 \cdot 0 + (0.9)^5(0.5)^3 \cdot 0 = 4.76$$

$$C : 1 + (0.9)\cdot 4 + (0.9)^2 \cdot 2 + (0.9)^3 \cdot 5 + (0.9)^4 \cdot 5 + (0.9)^5 \cdot 5 = 16.10$$

Thus we see that after just 6 turns C has a overwhelming advantage over R in terms of his discounted payoff. Now suppose R tries to recover by playing brilliant moves (so as to neutralize the penalty discounting λ_2). That is, suppose he scores 5 for each of the subsequent 15 turns. Since $c = 5$, after each set of 5 turns the index of λ_2 will reduce by 1. A simple calculation shows us that the payoff to R after these 15 turns (that is a total of 6+15=21 turns) would be 9.63. After that, the penalty discounting λ_2 would disappear. But from then on the global discount λ_1 itself would start contributing heavily to the weights of the moves and we can show that even if R keeps playing brilliant moves forever, the maximum payoff he can receive from then on is just 5.47. Thus his total payoff in the infinite game after the initial slump is 9.63+5.47=15.10 which is still less than what C has amassed in the first 6 rounds (16.10). This justifies Proposition 1 and shows that the Jury can already offer the win to C (which it implicitly does).

What is crucial here is that C's attack on R's delivery rings true, and the fact that R could have attempted to rebut C's commentary but did not, confirms C's characterization of it. This affects the rest of the debate's evaluation; R's subsequent moves never mattered. In other words, the fate of R's evaluation was sealed after this initial exchange of 3 rounds. Thus, not responding to an attack on either the style or the substance of ones contributions forces the evaluation to go negative as in (Asher et al., 2016)'s general constraint.

(Asher and Paul, 2013) gives another example of a disastrous debate move. Though (Asher and Paul, 2013) does not use a weighing function and discounted payoffs, we can still apply our formalism to that example. The example concerns Senator Dan Quayle's (Q) reply to a similar question about his experience to be President in the 1988 Vice-Presidential debate, in which he drew a parallel between his own experience and that of President John Kennedy (K). His opponent, senator Lloyd Bentsen (B), took a weak implicature from Q's response, that Q had the potential to be a similar president to K, and attacked it forcefully, drawing a roar of appreciation from the audience, giving Q a score of 0 for that move. Q's subsequent rejoinder *"that was unfair Senator, unfair,"* was a comment that did not take issue with B's drawing of the implicature concerning Q and K. This amounted to a tacit acceptance of the implicature. Given that B had refuted that implicature, Q was saddled with having conveyed an implicit content that he was unable to defend but accepted, which netted him a second zero, which was enough to sink his performance for the rest of the debate. B's attack move, though different from C's in (8) in that it attacked content not presentation, also colored Q's performance for the rest of the debate. Q's evaluation went to the bottom of the scale for the rest of the debate and stayed there, making B the clear winner.

We have modeled the consequences of disastrous moves on evaluations of a conversational play. But what about brilliant moves that are not attacks, how do they function? One memorable line used over by Ronald Reagan during the 1980 US Presidential campaign was *"Are you better off than you were four years ago?"* In one question, Reagan was able to remind Americans that they were worse-off under the incumbent Carter; inflation and unemployment had dramatically risen under Carter and purchasing power has waned. Carter himself described the American mood as a *"malaise"* during his Presidency. This one move set the tone for the discussion and put Reagan in a winning position, as Carter could not convincingly counter the obvious *"no"* answer to Reagan's question.

We can model the above in our setting of WME games with w assigning a 5 to this move by Reagan and a 0 to Carter. Carter's inability to respond convincingly saddles him with another 0 and this colors the evaluation by w of the ensuing debate, heavily favoring Reagan. Reagan continues to get high scores for all his moves while Carter fares badly, which accords with history: Reagan was pronounced a clear winner of the exchange.

5 Related Work

As alluded to in the introduction, game theory has been used before in the literature for the analysis of strategic message-exchange. The focus for the purpose has mostly been on the use of signaling games (Spence, 1973). However, signaling games lack the necessary tools to model situations where the interests of the players are opposed, as is the case in the current setting. Noteworthy also is the work on persuasion games (Glazer and Rubinstein, 2001; Glazer and Rubinstein, 2004) which has the setup similar to that of signaling games where a 'speaker' is trying to persuade an uninformed 'listener' about the current state of the world. Despite being hugely successful in modeling many different economic and strategic situations, signaling games have certain drawbacks which restricts their applicability to dynamic strategic conversations, as in the current setting. This issue has been extensively discussed in (Asher et al., 2016).

Our notion of evaluation makes use of discourse structural moves and depends on work on discourse structure and rhetorical relations like that of (Asher and Lascarides, 2003); to our knowledge, we are the first to model evaluations of conversational success by exploiting ideas of discourse coherence and discourse structure, along with techniques of discounting from game theory. Our account also makes at least informal use of the notion of an attack, and is thus related to work on argumentation (Dung, 1995; Besnard and Hunter, 2008). (Besnard and Hunter, 2008) also considers a definition for evaluating an argument by an audience. They structure arguments as trees, which roughly parallels the notion of a discourse graph in SDRT (Stede et al., 2016). They also use a discounting function, so that more deeply embedded arguments (responding to prior attacks) are weighted less than the main arguments and counterarguments at the top. This discounting function is similar to our λ_1. However, there is nothing in the argumentation literature of the form of our penalty discount λ_2 for convincing attacks and very bad moves. And to our knowledge, no one in the argumentation literature, or anywhere else, has tried to formalize an evaluation of attacks and refutations over the course of a dialogue. The analysis of argumentation in game theoretic terms, which is a consequence of our approach, is also the first of its kind to our knowledge.

Evaluations of conversational success are also related to linguistic work on predicates of taste (Lasersohn, 2005; Glanzberg, 2007; Crespo and Fernández, 2011), in that our evaluations are relative to the standards of a person or group. It may be that two people may disagree over a evaluation of i's contributions, because they have incompatible views of what constitutes conversational success for i, just as people may disagree about whether say blood sausage is tasty or not. The received wisdom about predicates of taste, however, is there is 'no fact of the matter' as to whether blood sausage is tasty or not. We do not believe this carries over to evaluations of conversational success. Given that players in a political debate have the goal of convincing the public, it is really the public's evaluation that counts and gives an 'objective' evaluation of the player's success in terms of their own interests. Work on automatic debate evaluation in terms of an audience's reactions has attracted interest in NLP (Prabhakaran et al., 2012; Prabhakaran and Rambow, 2013; Prabhakaran et al., 2013), for which weighted ME games provide a formal framework.

6 Conclusions

We have presented a model of the evaluation of conversational success, WME games. Extending the framework of infinite ME games for modeling conversations introduced in (Asher et al., 2016), we have shown how a Jury can concretely evaluate a player's conversational success. We have illustrated how such evaluations depend upon the structure and content of a person's contributions as well as on discounting functions, and we have analyzed at length one sample conversation to show an evaluation process at work. Our discounting functions entail: (i) it is best to get one's very good moves in early, (ii) a sequence of moves that are bad by Player i affects the evaluation of future moves, and in particular, (iii) a failure by i to respond effectively to a convincing attack on i's earlier moves is disastrous, because λ_2 becomes very significant.

There are many ways in which we wish to extend this work. First, we want to explore further the space of weighting and discounting functions; different functions will yield new and potentially interesting evaluation schemes. Secondly, we wish to enrich our model with an epistemic framework by introducing imperfect information (Harsanyi, 1968). In the present abstract, as remarked, we assume that the players are unaware of the parameters of the Jury. Elaborating on this, we might assume that a Jury can be of different 'types'. For instance, it may be 'biased' towards a particular player or may be 'fair' to everybody. It may be 'patient' (with high λ_1) or 'impatient' (with low λ_1); 'strict' (with low λ_2) or 'lenient' (with high λ_2). In addition, the players might themselves be of different types: risk-takers, risk-aversers, rational, irrational etc. Players are aware of their own types but are uncertain about the types of the other players and that of the Jury; they hold certain 'beliefs' about these unknown types. A player's strategy now depends not only on the history but her own type and her beliefs about the types of the other players and that of the Jury. Such an approach is standard in epistemic game-theory and we believe that augmenting the current framework of WME games with it will lead to a much more complete analysis of the behavior of conversationalists and evaluations of conversations.

Finally, we wish to explore the existence of Nash equilibria and other solution concepts in our WME games and explore rationality criteria.

References

Nicholas Asher and Alexandra Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.

Nicholas Asher and Soumya Paul. 2013. Conversations and incomplete knowledge. In *SEMDIAL 2013: Proceedings of the 17th Workshop on the Semantics and Pragmatics of Dialogue*, pages 173–176. Citeseer.

Nicholas Asher, Soumya Paul, and Antoine Venant. 2016. Message exchange games in strategic conversations. *Journal of Philosophical Logic*. In press.

Philippe Besnard and Anthony Hunter. 2008. *Elements of argumentation*, volume 47. MIT press Cambridge.

Christie-Rubio-debate. 2016. `http://www.realclearpolitics.com/video/`

2016/02/06/christie_vs_rubio_the_memorized_30-second_speech_where_you_talk_about_how_great_america_is_doesnt_solve_anything.html.

Christie-Rubio-transcript. 2016. https://www.washingtonpost.com/news/the-fix/wp/2016/02/06/transcript-of-the-feb-6-gop-debate-annotated/.

Vincent Crawford and Joel Sobel. 1982. Strategic information transmission. *Econometrica*, 50(6):1431–1451.

Inés Crespo and Raquel Fernández. 2011. Expressing taste in dialogue. In *SEMDIAL 2011: Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue*, pages 84–93. Citeseer.

Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357.

Michael Glanzberg. 2007. Context, content, and relativism. *Philosophical Studies*, 136(1):1–29.

Jacob Glazer and Ariel Rubinstein. 2001. Debates and decisions, on a rationale of argumentation rules. *Games and Economic Behavior*, 36:158–173.

Jacob Glazer and Ariel Rubinstein. 2004. On optimal rules of persuasion. *Econometrica*, 72(6):119–123.

Herbert Paul Grice. 1975. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics Volume 3: Speech Acts*, pages 41–58. Academic Press.

John Harsanyi. 1968. Games with incomplete information played by 'bayesian' players, parts i, ii & iii. *Management Science*, 14(7):486–502.

Peter Lasersohn. 2005. Context dependence, disagreement, and predicates of personal taste*. *Linguistics and philosophy*, 28(6):643–686.

Yehuda Levy. 2013. Discounted stochastic games with no stationary nash equilibrium: Two examples. *Econometrica*, 81(5):1973–2007, September.

Vinodkumar Prabhakaran and Owen Rambow. 2013. Written dialog and social power: Manifestations of different types of power in dialog behavior. In *IJCNLP*, pages 216–224.

Vinodkumar Prabhakaran, Owen Rambow, and Mona Diab. 2012. Predicting overt display of power in written dialogs. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 518–522. Association for Computational Linguistics.

Vinodkumar Prabhakaran, Ajita John, and Dorée D Seligmann. 2013. Who had the upper hand? ranking participants of interactions based on their relative power. In *IJCNLP*, pages 365–373.

Lloyd S. Shapley. 1953. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095–1100.

Michael A. Spence. 1973. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374.

Manfred Stede, Stergos Afantenos, Andreas Peldszus, Nicholas Asher, and Jeremy Perret. 2016. Parallel discourse annotations on a corpus of short texts. In *Proceedings of LREC*.