

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's [AM terms of use](#), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is forthcoming at <https://www.springer.com/journal/11229>.

Tweet Acts and Quote-Tweetable Acts

Chris Cousens

University of Stirling, chrisjcousens@gmail.com

Abstract

Online communication can often seem different to offline talk. Structural features of social media sites can shape the things we do with words. In this paper, I argue that the practice of ‘quote-tweeting’ can cause a single utterance that originally performed just one speech act to later perform several different speech acts. This describes a new type of illocutionary pluralism—the view that a single utterance can perform multiple illocutionary acts. Not only is this type more plural than others (if one utterance can acquire *many* kinds of illocutionary force), but it also shows how illocutionary forces can be accumulated over time. This is not limited to online utterances—some offline contexts are similarly structured, and so offline utterances can also come to perform many different speech acts.

1. How to do Things with Tweets

Speech acts are the things we do with words. We can assert, promise, warn, name, marry, or order an orange mocha frappuccino. The type of action performed with words is dubbed the *illocutionary force* of the utterance (Austin, 1962). The illocutionary force of ‘I hereby assert that P’ is that of an assertion.

Re-tweets are a feature of the social media site Twitter.¹ These allow users to broadcast *someone else's* tweet to their followers. *Quote*-tweets allow the user to add a comment to the re-tweet:

¹ Twitter has been renamed X; to avoid confusion, I use the original name.



But what happens to the original utterance once it is quote-tweeted? If it was a speech act with illocutionary force, does it retain this and remain the same type of speech act as when it was originally uttered? An ordinary quotation brings an illocutionary act from one discourse into another (Oishi, 2022). If I utter ‘Britney said “I promise to drive Christina to the airport”’, I import the normative obligations and expectations from Britney’s original conversation to this one. Others in the conversation should now take Britney to be obliged to drive Christina to the airport, despite not hearing the original promise.

But quote-tweeting does not only quote the original tweet—it also adds a comment. Marsili (2021) suggests that this might substantially alter the meaning of the tweet being quoted; here I argue that it may also alter its illocutionary force by changing how other people are supposed to treat the original tweet. As a tweet can be quote-tweeted many thousands of times, quote-tweeting could give the original tweet many different illocutionary forces. This is, in effect, an argument for a strong (and perhaps strange) kind of *illocutionary pluralism*—the view that a single utterance can be used to perform multiple speech acts (and thus have more than one type of illocutionary force).

This is not unique to online conversations and quote-tweeting. Philosophers have argued that online communication is far from ideal (Frost-Arnold, 2021; Goldberg, 2021), but we do not have to perfectly perform speech acts for them to still have illocutionary force (Johnson, 2020; see also Kukla, 2014). Speech acts are *temporally extended*—it takes time for them to be performed (Lance and Kukla, 2013; Kukla, 2014; Caponetto, 2020). As our online utterances can last a very long time—some persist for decades—I suggest that the speech acts they perform can not only change, but also multiply.

§2 outlines a puzzle—how a quote-tweeted utterance might seem to perform several different types of speech act. §3 considers whether existing accounts of illocutionary pluralism can explain this. §4 develops a new account whereby changes in conversational score alter the norms enacted by the original utterance. §5 extends this to offline speech, while §6 considers two potential objections.

This paper does not offer an analysis of the communicative properties of the act of quote-tweeting itself. The difference between a re-tweet and a quote-tweet is the user-added comment. That comment might be an assertion, a question, a directive, and so on. ‘This is so right!’ would *agree*; ‘What a lovely way to say this!’ would *compliment*; ‘Please, no one do this!’ would *instruct*. This may mean that no unified account of *the* communicative act of quote-tweeting is possible, unlike other online communicative acts such as liking (McDonald,

2021) or re-tweeting (Marsili, 2021). Instead, my focus is on how quote-tweeting affects the illocutionary force of the original tweet.

2. The Puzzle: Quote-Tweeting a Speech Act

A politician logs into Twitter and posts the following OP²:



They have not only conveyed some information about their intentions. They have also *done* something: performed an assertion. To use terminology coined by Austin (1962), their utterance has the *illocutionary force* of an assertion.³ A different utterance might have a different illocutionary force: ‘Should we cut taxes?’ would have the force of an enquiry, while ‘To my party members: cut taxes!’ would have the force of an instruction. Speech acts are the things we do with words. Illocutionary force is the concept we use to describe what, exactly, is being done. Austin distinguishes *illocution* (what is done) from *locution* (what is said) and *perlocution* (the effects of the utterance)—although of course utterances have content, effects, *and* illocutionary force.

There is nothing new about using speech act theory to analyse online utterances. Computational linguists have used it to create algorithms to identify assertions, questions, recommendations, and requests in online corpuses (Zhang et al., 2011; Vosoughi & Roy, 2016). However, as Lewiński (2021b) notes, this taxonomic approach does not capture the logic and patterns of actual conversation, online or offline. It is one thing to say that the OP constitutes an assertion, but another to explain what the politician asserts with it and what normative effects their assertion has.

The politician tweets something, but does not *quote-tweet*, our object of study. Their tweet is quote-tweeted by another user.

QT1:

² ‘OP’ stands for ‘original post’—the first post in a ‘threaded’ online discussion.

³ Readers might think this is instead a prediction, pledge, or promise. For now, allow that this utterance *may* have the illocutionary force of an assertion.



They just said they would cut taxes!



The comment attached to this quote-tweet seems relatively straightforward. The politician tweeted that they would cut taxes, and the quote-tweeter accurately reports this, with the embedded quotation of the OP providing evidence for the accuracy of the report. But lots of other comments might have been added, and these could be less straightforward.

QT2:



They just suggested they would cut taxes!



QT3:



They just agreed they would cut taxes!



The verbs here— ‘suggested’ and ‘agreed’—are less neutral than ‘said’. They *colour* the report. This might make the quote-tweet seem less accurate, if the speaker was not merely suggesting but was strongly asserting. The politician might object to this characterisation of their utterance. But instead, upon receiving notification that their OP had been quote-tweeted, the politician might reply in a way that tacitly endorses this characterisation— ‘Yes, I did; I think it’s a great idea!’. Does the quote-tweet mischaracterise the OP *after* the politician posts such a reply? My answer (no) will come later.

The puzzle arises because the politician *could* use the OP to assert, suggest, or agree. The same words, under slightly different conditions, can constitute a range of different speech acts. Depending on tone, previous utterances, or the political context, any of these options might seem most fitting. The speaker might have been ambiguous in the speech act they intended to perform, accidentally or ‘strategically’ (see Lewiński, 2021a). Utterances only

need to be *good enough* to successfully perform speech acts—they do not need to be *perfect* (Johnson, 2020). Importantly, all three options (assertion, suggestion, and agreement) incur similar normative commitments—that the speaker endorses cutting taxes. So, the speaker might be aiming at this group of speech acts and count it as successful if they perform one of them.⁴

But a quote-tweet could also invoke illocutionary force that the speaker is *not* aiming at:

QT4:



Paaarty Animal
TheElectricMayhem

They just promised they would cut taxes!



QT5:



From The Balcony
waldorfwisdom

They just threatened they would cut taxes!



Again, the politician might object to this interpretation of the OP— ‘I can’t *promise* to cut taxes because I need others to vote with me!’. And they may want to avoid being ‘on the hook’ for the obligation of a promise—all else being equal, people typically think worse of someone who breaks a promise than someone who makes an incorrect assertion. The politician might not mind being thought to say something wrong, but mind very much being thought to break promises.

Some politicians still seem to care about such things.

Assertions and promises both alter normative obligations between speaker and listener, but in different ways—and as suggested above, the politician may be very interested in precisely which norms are enacted by their utterance. Following Sbisà (2013) and Kukla (2014) I take this enacting of norms and altering of interpersonal commitments to be the central purpose of speech acts (see also Lance and Kukla, 2013; Sbisà, 2009).⁵ The politician may object to QT4

⁴ Green (2013) describes the ‘assertive family’ of speech acts (presuppositions, guesses, conjectures, and presumptions, as well as assertion proper); the speech acts in QT1-3 might be related to this family, while the speech acts in QT4-7 might only be grouped together at the level of class or order.

⁵ Sbisà (2013, p242-3) identifies the illocutionary act with the “production of a change in the conventional aspects of the interpersonal relationship among the participants”.

because they are only willing to risk the criticism that follows an incorrect assertion, and not a broken promise.

They might also reject QT5. While they do intend to cut taxes, they do not mean it as a threat. The clearest threats are conditional (Schiller, 2021; Fraser, 1998 calls these ‘direct’ threats). Do X, or I shall Y. Here, the politician says they shall Y, but nothing about X (which would, if the OP is a threat, lead to $\neg Y$). So, the politician might say that characterising the OP as a threat is inaccurate. Furthermore, they might say that they did not intend to intimidate anyone, so any feelings of intimidation are accidental (Fraser, 1998 takes this intention to be necessary for issuing a threat).

But just as with QT2 and QT3, the politician might reply to QT4 and QT5 *endorsing* these interpretations. They could say ‘Yes, I promised I would cut taxes, so you should vote for me!’, or ‘Yes, if the opposition does not agree to decrease benefits, I will cut taxes!’. After such a reply, it seems wrong to say that QT4 or QT5 is *inaccurate*, or mischaracterises the OP. After all, the person who posted the OP just replied as though they did, in fact, promise or threaten.

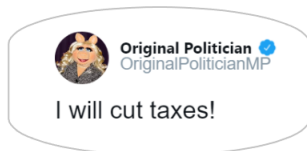
The speaker does not *need* to reply for this puzzle to occur. The discussion could proceed without them (see Maitra, 2012 on omission shaping illocutionary force; also, Sbisà, 2009). Online discussion is rife with miscommunication, so I include the speaker’s ratifying reply to simplify things for now, and to forestall potential objections (until I discuss them in §6).

The illocutionary possibilities extend even further:

QT6:




They just conceded they would cut taxes!



QT7:



They just warned they would cut taxes!



The politician might object to each of these, claiming that they have negative connotations ill-fitting the beneficence of cutting taxes. They might say ‘In saying that I will cut taxes, I

conceded nothing!’, attempting to portray themselves as strong, as powerful, as a winner. Or they might say ‘Warnings are for bad things, and cutting taxes is good!’. But again, they might reply to the quote-tweet in a way that seems to ratify this interpretation (‘I’m conceding this for a good reason!’; ‘If you’re worried about tax cuts, then you should consider this a warning!’).

So, the politician says in the OP that they will cut taxes. The OP is quote-tweeted several times, with each quote-tweet treating the OP as a different speech act. The politician replies to each quote-tweet in a way that endorses (or tacitly accepts) this characterisation of the illocutionary force of the OP. And—this is important—all of this could happen at once.

Imagine that the quote-tweets happen almost immediately. Within a few minutes, alerted to this activity by their Twitter notifications, the politician has replied to all of them. In one reply, they endorse the characterisation of the OP as an assertion. In another, as a promise. In another, as a warning. And so on.

This may sound strange. If the speaker was promising, could they have also been warning? If they were threatening, could they have also been suggesting? It sounds even stranger to say that the speaker *intended* to perform all these different speech acts when they initially tweeted, particularly when some seem at odds with others (although Saul 2018 shows that skilled speakers can intend to convey different content to different audiences, as with dog-whistles).

But even if the speaker did not originally intend their utterance to function as several different speech acts, their replies endorsing such inconsistent characterisations of their utterance need not be irrational. If each of the quote-tweeters is a community leader whose support the politician wishes to court, they may decide that it is better to agree with seemingly incompatible interpretations of their utterance than to (publicly) disagree with those leaders. This might be sneaky, two-faced, and duplicitous, but hardly irrational.

Now, hopefully, the puzzle is clear. After each quote-tweet (and reply), the OP seems to have performed a different speech act. In one conversation, it appears to be correctly treated as an assertion. In another, it appears to be correctly treated as a suggestion, in another as an agreement, a promise, a threat, a concession, or a warning—including by the politician who originally posted it. So, which (if any) of these speech acts is now performed with the OP? I will argue that it is *all* of them.

Within the discussion following each quote-tweet, it is conversationally appropriate for speakers to treat the OP as the type of speech act identified in that quote-tweet. I think that the characteristic goal of speech acts is to structure normative interpersonal relations (Sbisà, 2009, 2013; Kukla, 2014). What makes a promise distinctively a *promise* is the obligations it places upon the speaker and the expectations it licenses for the addressee; what makes an order distinctive is the obligations it places on the addressee and the expectations it licenses for the speaker.

So, by making it conversationally appropriate (which is a kind of normative relationship) to treat the OP as a different kind of speech act, these quote-tweets make it the case that the

norms enacted by the OP now indicate that it has seven different illocutionary forces. In other words, people within each conversation should treat the OP as constituting a different kind of speech act. I will argue that this means it *is* several different speech acts—or at least might as well be—even if it started as only one.

3. Varieties of Illocutionary Pluralism

The puzzle outlined in §2 suggests that the OP demonstrates *illocutionary pluralism*—the view that a single utterance can constitute (i.e., be used to perform) more than one illocutionary act (Sbisà, 2013; Johnson, 2019; Lewiński, 2021a).⁶ The alternative is *illocutionary monism*, a longstanding assumption that a single utterance can only perform a single speech act (setting aside conjunctions like ‘I hereby assert that it is raining, and order you to fetch my umbrella’). Sbisà (2013) argues that this is a received view in speech act theory which takes illocutionary acts to be an expression of the speaker’s intention (more on intention in §6; see also McDonald, 2022). In this section, I survey existing analyses of illocutionary pluralism and argue that an alternative account may better explain our puzzle.

Many utterances *do* only perform one speech act. ‘I promise to meet you tomorrow’ performs a promise. It is not also an assertion, nor a suggestion, nor an order for an orange mocha frappuccino. But other utterances seem to perform more than one. ‘Can you pass the salt?’ seems to perform both an enquiry and a request. The speaker asks about salt-passing ability, and in doing so indirectly requests that they be passed the salt (Searle, 1975). Alternatively, ‘You haven’t seen the last of us!’ might be both an assertion and an indirect threat (Fraser, 1998; Schiller, 2021).

But these utterances still perform one *primary* speech act. The purpose of ‘Can you pass the salt?’ is to request that the salt be passed. The phrasing as an enquiry is a sort of ‘lower-order’ speech act which is conventionally used to perform the ‘higher-order’ request. Lewiński (2021a) describes this as a *vertical* arrangement of multiple speech acts, with one more important than the other. ‘Can you pass the salt?’ might, then, perform multiple distinct speech acts, but the primary act is the one that enacts (at least most clearly) norms and obligations. Responding to ‘Can you pass the salt?’ as a genuine enquiry—‘Yes, I can pass the salt,’ without actually passing it—is inappropriate (albeit very funny). The ‘surface grammar’ of the utterance does not by itself determine the speech act it performs (Kukla, 2014; see also Clark, 1996, who suggests ‘layers of activity’ for language use). And, importantly for our purposes, the OP is not an indirect speech act—‘I will cut taxes’ does not have stable conventions surrounding its indirect use, whereas ‘Can you pass the salt?’ does.

However, utterances might have *horizontal* illocutionary pluralism. Rather than one speech act being used to indirectly perform another, a single utterance might be used to do multiple things concurrently. Sbisà (2013) argues that this is presupposed by Austin’s (1962) original

⁶ This is pluralism of *type* of speech act; Cappelen and Lepore (2008) and Egan (2009) suggest that a single utterance might perform lots of instantiations of the same type of speech act—see §4.

distinction between rhetic, phatic, and phonetic acts, as well as between locutionary and perlocutionary acts. If one utterance is already performing all these kinds of acts, it need not be limited to just one illocutionary act.

In addition to indirect speech acts, Sbisà (2013) identifies two more ways that utterance tokens might be illocutionarily pluralistic. The first is when an utterance is ambiguous between two illocutionary patterns: ‘There’s a bull in that field’ fits the pattern of an assertion and of a warning. The utterance could be either; speaker and addressee must ‘negotiate’ to settle the matter. But once it has been settled, Sbisà notes, this no longer seems pluralistic—the utterance token now only performs one illocutionary act.

However, in the puzzle outlined in §2, each conversation seems to settle the illocutionary force of the OP differently.⁷ Sbisà focuses on conversations between one speaker and addressee—but in cases where the utterance (through quote-tweeting) is part of multiple conversations, its force can be settled in multiple ways. So, in our puzzle, even though the OP has a *locally* settled force, if we ‘zoom out’ to look at all seven conversations at once, it still seems pluralistic.⁸

The second kind of pluralism Sbisà identifies is an utterance designed to conflate two ‘illocutionary patterns’. The following example I adapt, with kind permission, from a talk by Lewiński.⁹ Kelly says to her friend Michelle ‘Can you do me a favour in coming to my sister’s wedding?’. Michelle replies, ‘Of course!’. Has Kelly requested that Michelle do her a favour? Or has she invited Michelle to the wedding?

The answer is, I think, that she does both. If, after Michelle attends the wedding, and Kelly says ‘Remember, I owe you a favour!’, she would be surprised if Michelle replied ‘No you don’t, you invited me to the wedding, but didn’t ask me for a favour.’ And if Michelle did not attend the wedding, she could hardly say to Kelly ‘You asked me for a favour but didn’t actually invite me to the wedding!’

This cuts both ways. Michelle would be surprised (and annoyed!) if Kelly refused to return the favour by saying ‘Actually, I invited you to the wedding but didn’t request a favour’. And it would be odd if Kelly was surprised when Michelle arrived at the wedding, saying ‘I only requested a favour, I didn’t actually invite you!’

This indicates that Kelly’s utterance has updated the normative obligations between her and Michelle in multiple ways. It becomes appropriate, after Michelle says ‘Of course!’, for her to expect Kelly to return the favour. And it becomes appropriate for Michelle to attend the wedding. While, as Lance and Kukla (2013) note, invitations and requests (such as for favours) function similarly, they are not the same. Kelly has not simply performed two different requests (to owe Michelle and favour and to have Michelle attend the wedding). Rather, there are “indefinitely many kinds of calls, with distinctive structures, whose

⁷ Or at least, each conversation settles on treating the utterance as having a different illocutionary force; as explained earlier, I take the way an utterance is treated and its illocutionary force to be closely connected.

⁸ An additional difference between my account (in §4) and Sbisà’s: I do not afford a central role to uptake (a contested notion, see McDonald, 2022 and further discussed in §6.2).

⁹ At the Lisbon Summer School on Speech Acts in Public Discourse, June 2022.

subtleties help to constitute a rich moral and social space” (Lance and Kukla, 2013, p458). Despite both asking something of the listener, begging a favour and issuing an invitation structure normative relationships in different ways. Kelly uses a single utterance to do both.

If the utterance seems to be intended as two speech acts, receives uptake as two speech acts, and changes normative obligations in two distinct ways appropriate to those two speech acts, then a single utterance can perform more than one speech act even in conversation between just two people.¹⁰ And this is not an indirect speech act, where the grammatical form of asking a favour is used to issue an invitation. The utterance aims at the successful performance of *both* speech acts, whereas ‘Can you pass the salt?’ only aims at getting the salt passed, not discovering information about the target’s salt-passing ability. However, this still does not explain how the OP comes to count as seven different speech acts. The OP does not have the complex construction of Kelly’s utterance, and within each conversation it is treated as having a single type of illocutionary force.

Illocutionary pluralism as discussed so far involves conversations between two people. But one of the most striking features of the puzzle in §2 is how many people it involves. When it comes to conversation online, particularly for famous speakers with many followers, their utterances might find *millions* of addressees.

Johnson (2019) and Lewiński (2021a) develop accounts of illocutionary pluralism that leverage the presence of multiple addressees to explain how a single utterance might perform two separate speech acts (here I conflate these two accounts). If I am speaking to Mel (my boss) and Geri (my subordinate), I say ‘Please finish the reports by the end of the day’. For Mel, this is intended to be a request. For Geri, this is intended to be an order. Furthermore, I have no institutional authority with which to order Mel to do anything; on the other hand, I *do* have the authority required to give orders to Geri. Johnson (2019, p. 1156-7) argues that intention is not required for illocutionary force; I include it here so that this example is compatible with Lewiński’s (2021a, p6689) analysis which *is* restricted to intended speech acts.

Mel recognises that with this utterance I intend to make a request of her and responds appropriately, trying to finish the reports but not feeling bound to do so. She gives the utterance *uptake* as a request. Geri, on the other hand, recognises that I intend to order her and consequently stays back late to finish the reports, giving the utterance uptake as an order.

According to illocutionary monism, my utterance can only have performed a single speech act, so I must have either performed a request *or* an order. Either Mel or Geri is wrong about which I have performed. But this would ignore standard stories of the roles of intention, uptake, and normative obligation in the performance of speech acts.

¹⁰ See also Lewiński’s (2021a, §4.1) discussion of other ways that illocutionary pluralism might occur in a dialogue. There, examples show speakers using illocutionary pluralism to keep open the ‘space of illocutionary possibilities’, either through prompting dilemmatic deliberation or through strategic ambiguity. But rather than keeping illocutionary possibilities open, Kelly’s utterance makes two (requesting a favour and issuing an invitation) concrete.

It is not just that Mel incurs a weak obligation to finish the reports and Geri incurs a strong obligation—as Kukla (2014) notes, different kinds of speech acts bring about different kinds of normative changes. For example, Geri faces potential remonstrance if she fails to complete the reports, whereas Mel does not. It looks like my utterance has successfully performed two distinct speech acts, one for each listener, based on the normal requirements for the performance of a speech act. The only thing that might stop this from working is a commitment to illocutionary monism. Therefore, Johnson (2019) argues, we should be suspicious of illocutionary monism. I agree.

But even though this kind of illocutionary pluralism involves multiple addressees, it still does not explain the puzzle in §2. In the example request/order for Mel and Geri, I can perform two different illocutionary acts because of the relevant institutional relationships between myself and the two addressees. The politician posting the OP (and subsequent quote-tweeters) do not know who will see their tweet, and so cannot be aware of the relevant relationships between themselves and their audience (Marwick and boyd, 2011; Frost-Arnold, 2021). And little seems to turn on the institutional affiliations of the quote-tweeters, whereas for Mel and Geri, it is our specific institutional roles that underpin the dual intentions (and uptake) of my utterance. Perhaps the politician is ‘strategically ambiguous’ in their OP, saying something that *could* be interpreted in many ways and hoping that listeners will resolve this ambiguity to the advantage of the politician (see Lewiński, 2021a). But after the politician replies to the quote-tweet, such ambiguity disappears (Saul, 2018 discusses intended ambiguity further).

There are other ways that an utterance might appear to be used to perform multiple speech acts. Elliot-Maksymowicz et al. (2021) identify political tweets that function as both assertions and expressives (e.g., ‘I don’t want to see that Trump garbage!’)—but the OP does not have the hallmarks of this kind of expressive utterance. McGowan (2019) shows that a single utterance can count as a move in multiple activities (see also Sbisà, 2013 describing a single utterance as a step in multiple ‘interactional tracks’, invoking Levinson, 2011). Hate speech might be both a conversational move (like an assertion) *and* a move in an ‘activity of oppression’, although McGowan characterises the non-conversational move as a ‘parallel act’ rather than a speech act (see Lewiński 2021b, for further discussion).

This is all to say that while it might seem conceptually tidy for a single utterance to have a single illocutionary force (and thus perform a single speech act), illocutionary practices out in the wild are messier than that. We can perform more than one illocutionary act with the same few words. But the varieties of illocutionary pluralism surveyed here do not fully explain the puzzle in §2. This is not a weakness of previous accounts of illocutionary pluralism; they are examining different (albeit related) phenomena. In the next section, I develop an account of illocutionary pluralism better suited to explaining the multiple illocutionary forces of online utterances with large audiences, such as the OP. In §5 I extend this to similarly structured offline utterances. §6 considers two potential objections.

4. Illocutionary Pluralism and Conversational Score

I have suggested that a single utterance can end up with many different illocutionary forces. How can we do so many things with so few words? §3 outlined the explanatory power of illocutionary pluralism, but previous accounts of this phenomenon do not quite fit the puzzle in §2. What is required is an account of an utterance starting out as a single speech act (in the OP, an assertion) before later counting as several different speech acts—*without* drawing on defined interpersonal relationships which are often absent on Twitter.

Egan (2009) suggests a similar problem, with one utterance eventually performing many different speech acts. A billboard reading ‘Jesus loves you!’ seems to perform a different assertion for each passer-by who reads it. To Shania, it says that Jesus loves Shania. To Alanis, it says that Jesus loves Alanis. To Kelis, it says that Jesus loves Kelis. The ‘you’ is directed at the individual reader, not at a group or collective, and so the utterance performs a different assertion for each passer-by (the same could happen if uttered by a preacher to their congregation). This is dubbed a ‘shotgun’ assertion, indiscriminately spraying out speech acts.

But the pluralism here (and in Cappelen & Lepore’s 2008 defence of semantic minimalism) is in instantiations of the same *type* of speech act. The utterance performs lots of different assertions, but it only performs assertions. The puzzle in §2, on the other hand, involves a single utterance performing different types of speech act. After being quote-tweeted seven times, the OP seems to assert, suggest, agree, promise, threaten, concede, and warn. While Twitter posts *could* be phrased as a ‘shotgun’ (‘I will cut *your* taxes!’), the OP is not.

I briefly sketch out a solution, before adding further detail.

The discussions taking place after each quote-tweet are parts of different conversations (maybe divergent extensions of an original conversation, or maybe brand-new ones).¹¹ Conversations are governed by rules and norms—some contributions are appropriate, others are not. What is appropriate in one conversation might not be in another (for example, a professor might ask a postgraduate class, but not an undergraduate class, to use their first name). Within the conversations following QT1-7, it is appropriate to treat the OP as the speech act indicated in the quote-tweet. After QT1, it is appropriate to treat the OP as an assertion. After QT2, it is appropriate to treat it a suggestion, and so on. Speakers could challenge that interpretation of the OP—but unless they do, the utterance of QT1-7 makes it the case that it is conversationally correct to treat the OP as having this (new) illocutionary force. In that conversation, the OP *now* counts as a suggestion, or a promise, or a threat, even if it used to count as an assertion. As it is the current state-of-play that shapes normative relations between conversational participants, and (as suggested in §2) I think that shaping normative obligations is what speech acts are all about, we should say that the OP now *has*

¹¹ It can be difficult to identify the beginning or end of conversation. For some, such as telephone conversations (Schegloff, 1968) or those between strangers (Goffman, 1963), it will usually be easier—but not for online conversations (Marwick and boyd, 2011).

these new illocutionary forces. In other words, it now constitutes (or at least counts as) seven different speech acts.¹²

That is the sketch—now for the detail.

What makes an utterance appropriate or not? It depends on the rules and norms governing the conversation at hand. McGowan (2009, 2019) argues that some norms will be very general, such as the rules of grammar and syntax. Others are situational and specific, updated locally. When a professor says to a student ‘Just call me Whitney’, they change the norms governing this conversation here and now. This student may, in this conversation, use the professor’s first name—the extent to which this ‘carries over’ to future conversations may vary.

While general rules governing communication will almost always be in play, specific and situational norms are constantly being updated as new conversational contributions are made (McGowan, 2019). Because what counts as an appropriate contribution depends on what has been said before, participants need to keep track of this to know how to proceed. For example, whether a student may permissibly call their professor Whitney depends on whether permission to do so was given.

A useful way to conceptualise this is Lewis’s (1979) notion of a *conversational score*. Lewis argues that just like in a game of baseball, correct play in conversations is governed by both rules and the ‘score’. The rules tell players how to play the game, while the score tracks what has happened in *this* game so far. A batter in baseball is out after their third strike, but not their second. So, after striking, to know if they should stop batting (i.e., if this is correct play), they need to know not only the rules of baseball but also the score in this game at this moment—specifically, how many strikes they are currently on.

Lewis applies this distinction between rules and score to conversations. To work out which utterances would be appropriate (i.e., count as ‘correct play’), speakers need to know not only the general rules for conversation—rules of logic, syntax, Gricean maxims, and so on—but also the conversational score: what has been said so far in this conversation? For example, after I have said ‘I took my dog for a walk,’ it is no longer correct play for you to ask if I have any pets. I have already said something that presupposes that I *do* have a pet, so the ‘score’ is updated to include this information (for further discussion of presuppositions and conversational score, see Langton and West, 1999; McGowan, 2004, 2009, 2019; Langton, 2018; Witek, 2019).

So, when someone posts a quote-tweet that says the OP *promised*, or *threatened*, or *conceded*, they contribute to the conversation in a way that presupposes that the OP had that type of illocutionary force. They do not need to assert the proposition that it constitutes such-and-such a speech act, just as I do not need to say ‘I have a pet’ to add the fact that I have a pet to the conversational score—instead, I might say something that presupposes that I have a pet. The score is updated with both the content of the utterance *and* the presuppositions required to make sense of it. Unlike baseball, Lewis argues, when someone in conversation

¹² I think that the OP *has* new illocutionary force rather than only *counting as* having new illocutionary force—I outline further advantages of this framing in §6.1.

says something that does not make sense (‘incorrect play’), the score updates automatically to make it *count as* correct. To make sense of ‘I took my dog for a walk’, which would only be correct play if I have a pet, the score updates to *accommodate* the utterance by adding the presupposition that I have a pet.

So, even though the OP might have looked like an assertion, QT1-7 only make sense if the OP is a different type of speech act. The presupposition that it is a promise, a threat, a warning, and so on, is then added to the score of each conversation to accommodate the quote-tweet so that it now counts as ‘correct play’. Someone might object (‘Actually, I don’t think they were promising!’), but accommodation happens automatically (according to Lewis), and so without an objection it becomes correct play, from then on, to treat the OP as the new type of speech act. And, as this happens concurrently in seven different conversations, it is now correct play to treat the OP as being seven different speech acts. Participants in each conversation should treat it as the type of speech act the score in their conversation says it is.

Conversational score is not the only way to think about tracking what has happened in conversation. Another option is to describe this as ‘common ground’ shared by participants, rather than conversational score. The ‘common ground’ describes the knowledge, beliefs, and suppositions speakers take each other to share, and they refer to these when trying to work out what each other is saying (Clark, 1996; Stalnaker, 2002, 2014). This has been discussed in further work on presupposition and accommodation (e.g., Langton, 2012, 2018; Witek, 2019; see also Camp’s 2018 related notion of conversational record).

McGowan (2019) argues that for some purposes, the Lewisian notion of scorekeeping may be more suitable than that of common ground; for the questions under discussion in this paper, I agree. McGowan says that common ground tracks psychological facts, whereas the conversational score also includes non-psychological facts not shared by all participants, and so *encompasses* common ground. If a professor has permitted students to use their first name, this affects conversational permissibility and ‘correct play’ even if a late arrival does not realise it (and so it affects permissibility even though not all participants take it to be shared knowledge). In this way, the score is an objective standard of permissibility—perhaps supervening on the psychological facts of common ground—at least, ‘objective’ in the sense that there is a fact of the matter about what is conversationally permissible, and participants can be wrong about this (more on this in §5).¹³ I think there is also a methodological reason to prefer conversational score here: it more sharply reflects the notion of ‘correct play’ that informs conversational appropriateness.

Online conversations do not follow all the norms of offline conversation (which conversational score was developed to describe). Online posts are permanent; people can read what had been said before to ‘catch up’ and work out what the score is. On the other hand, participation in an online conversation is notoriously hard to track (this is ‘context collapse’,

¹³ Camp (2018) also argues for a more ‘objective’ or ‘normatively constrained’ understanding of conversational score or record. See also Witek (2015) on the distinction between objective and subjective accommodation—although I do not take my view to presuppose uptake externalism (see §6.2).

see Marwick and boyd, 2011; Frost-Arnold, 2021). But despite these differences, online conversations still have norms governing correct play. A response to the OP such as ‘Great, I hate taxes!’ would be appropriate. On the other hand, ‘Why are you saying you will raise taxes?’, or ‘I make \$30 an hour working from home, here’s how you can too!’ would not. A response can only be inappropriate if there are norms governing appropriateness for speakers to transgress. Online communication is clearly *conversational*, despite not always being a ‘good’ conversation (Goldberg, 2021). So, I refer to the appropriateness-tracking mechanism for both online and offline conversation as ‘conversational score’, even though precisely *how* it is tracked might be slightly different in each medium.

Does all this mean that the OP merely ‘counts as’ each of these different speech acts for the purposes of conversation? I think so—but there is nothing ‘merely’ about it. After QT4, the OP should be treated as a promise by participants within the subsequent conversation. This does not only cover how it should be spoken about. There are lots of norms enacted by a promise. We are entitled to expect the promiser to attempt to complete the promised action and they are liable to be criticised if they fail. Once the conversational score is updated to count the OP as a promise, it become appropriate (i.e., correct play) to treat the politician as *having promised*: to act as though one expects them to cut taxes and to criticise them if they do not at least attempt to do so. The politician (within this conversation) is now ‘on the hook’ for the obligations and expectations of a promise. A hold-out participant might think that QT4 is mistaken, and that the OP *just is* an assertion (and that the score is wrong). But even for them, if they do not actually voice their objection (and thus again change the conversational score), it is correct play according to the rules of the conversation to act as though the OP is a promise.¹⁴

In short, once the quote-tweet has re-scored the OP as a promise, participants should treat it as a promise. It affects the normative landscape—the obligations people have to each other—in the same way that ‘I hereby promise to cut taxes’ would have. So, the OP does not ‘merely’ count as a promise—it counts as a promise! At least, within this conversation and so long as there is no objection to QT4.

Does ‘counting as’ a promise mean that the OP now *actually has* the illocutionary force of a promise? That will depend on one’s prior speech-act-theoretic commitments. I have outlined mine—I think that enacting normative effects is the core business of speech acts (invoking Sbisà, 2009, 2013; Lance and Kukla, 2013; Kukla, 2014), and so as the OP now (locally) enacts the normative effects of a promise, we should say that it is now a promise (within this conversation)—albeit perhaps an imperfect promise. In §6.1 I will consider whether the speaker’s original intention might over-rule the normative effects of an utterance in determining its illocutionary force. For now, though, I hope to at least have shown that QT4 makes it the case that within this conversation, the OP *might as well as* be a promise.

¹⁴ It might be tempting to say that as listeners can object, or go along with, the QT’s characterisation of the OP, this is a process of ‘joint negotiation’ (Clark, 1996) or collaboration (McDonald, 2022). This is not quite what I am describing; instead, speakers change what is conversationally appropriate without the assistance of the listener.

Just as for ordinary promises, the normative effects of a re-scored promise can extend beyond the current conversation, carried forward by the beliefs and behaviour of participants (for a related ‘carry-over effect’, see Popa-Wyatt (forthcoming)). If Natalie says to Martie ‘I promise to take you to the airport’, and one of them later reports this to Emily, it becomes appropriate for Emily to hold Natalie accountable for promising even though she did not hear the promise itself (see Oishi, 2022). Similarly, if Emily is told that a politician tweeted a promise to cut taxes, even if this reports a conversational re-scoring rather than an ordinary promise, it would (I think) become appropriate for her to criticise the politician if they failed to follow through. From Emily’s perspective, there is little to distinguish between the report of an intended promise and that of a re-scored promise—and it is the report that gives Emily access to the norms enacted by the promise. However, the new normative effects of a re-scored promise will be less reliable than those of an ordinary promise. Some participants might play along with the re-scoring but privately believe that no ‘real’ promise occurred; they are unlikely to report that the politician promised. And for this type of public utterance, some third parties might already think that the politician’s utterance is a different kind of speech act, disagreeing with reports of their ‘promising’. Emily’s attempt to hold the politician to their ‘promise’ would then be more likely to be challenged. More will be said about how such communication breakdown might be resolved in §5.

I finish this section with some further comments on the relationship between a quote-tweeted utterance and its illocutionary force.

Quote-tweeting is not the only way that conversational participants might try to sneakily place obligations onto someone without their knowledge. Someone might say ‘They promised to cut taxes!’ about a politician who had never committed to cutting taxes. This would be an outright lie—but without any objection, the score should update such that it is now correct play to treat that politician as having so promised. Of course, other participants could object (‘I doubt they said anything of the sort!’), but until they do, it becomes conversationally appropriate to hold the politician to that (confected) promise. However, while the lie and the quote-tweet might both be able to (conversationally) obligate the politician in similar ways, they do not enact identical updates to the conversational score.

Both do conversationally attribute the obligations of the promise to the politician, and accommodating both makes it the case that it becomes correct play to treat the politician as having promised. But the QT also adds the embedded OP itself to the conversational score, and so is tied directly to the OP. The lie is not tied to any utterance at all, leaving it ‘untethered’. This leads to an important difference between the norm-changes enacted by the QT and the lie.

A little tentatively: the norms enacted by the QT are *stronger* than those enacted by the lie; the obligations conversationally enacted with the lie can more easily be *undone*. A response to the lie such as ‘I doubt that!’, or ‘Unlikely!’, or even (in face-to-face conversation) an expressively arched eyebrow could call the lie into question; to accommodate these responses, the obligations enacted by the lie would be either undone or suspended until participants have worked out whether the lie will stand (this might ‘block’ the lie, undoing its normative effects after the fact, invoking Langton, 2018). Similar responses would probably

not undo the re-scoring brought about by the QT. ‘I doubt that!’ is going to get the respondent nowhere because the QT brings along with it (and adds to the conversational score) evidence (the OP itself) that should remove the respondent’s doubts.

While two different mechanisms might be used to bestow the status of promiser upon the politician, the normative effects of one are more easily removed from the conversational score than the other.¹⁵ As the formatting and mechanics of quote-tweeting tie the QT to an actual utterance, it is more robust, and more easily able to withstand pushback. This shows that there are limits to how easily and effectively we can use accommodation to change the illocutionary force of an utterance.

But *how* is this force-shifting limited? Which conversational moves might be able to change the force of an utterance, and into which other speech acts? Certainly, not all of them. I cannot say, in reply to the OP, ‘Yes, I accept your proposal of marriage!’ and in doing so become engaged to the politician. I have thus far focused on the conversational *score*, but changes in the score are governed by *rules* (Lewis, 1979).

For example, the rules of soccer detail how a goal may be scored. Once all the conditions for scoring a goal are satisfied, the score changes. This can be undone later—a team might think that a goal should not have been awarded, and increasingly, recording technologies enable appeals and revisions to the score. A referee can, within a limited timeframe, change their previous rulings and thus the score (and what constitutes ‘correct play’)—some sports also have multiple referees who disagree and negotiate to determine the score. The rules of soccer allow the referee to change the status of a certain movement of the ball from a goal into not-a-goal, but not into *two* goals. The rules of the game limit how the score can be changed. Similarly, in conversation, accommodation can re-score some illocutionary acts into other illocutionary acts but does not have unlimited illocutionary power.

I do not here try to give an exhaustive account of the precise limitations of conversational re-scoring, but rather try to outline where these will be so that future investigations may explicate them. These limitations will generally be grounded in the rules governing the performance of individual speech acts; different speech acts can be governed by very different rules. Some ceremonial speech acts will have specific and codified rules, while the rules governing (at least some) conversational speech acts may be less rigid (to borrow Bach and Harnish’s 1979 distinction). The limitations of conversational re-scoring will (in large part) be derived from the (disunified) rules governing speech acts themselves.

Some illocutionary acts can more easily be turned into certain others—the more similar the acts, the easier this will be. Caponetto (2020) describes something like this when explaining how speech acts might be *amended*. With ‘P... actually, I guess that P’, the assertion is amended into a guess. Assertions and guesses are (fairly) similar. Both commit the speaker (to some degree) to the accuracy of P. Both can be performed using similar locutions (such as ‘P’). And both enact similarly structured obligations and expectations for speaker and listener

¹⁵ Simpson (2013) describes a similar phenomenon, where the harmful norms enacted by hate speech are easier to introduce than remove, as ‘asymmetric pliability’. Here, it is certain methods of norm enactment (quote-tweeting rather than lying) that create norms which are harder to remove.

(although there are often subtle differences in type of obligation, see Lance and Kukla's 2013 discussion of *calls*).

So, 'P' might be easily amended into an assertion, or a guess, or perhaps a presupposition, a conjecture, or an assumption. Green (2013) calls this the 'assertive family'. Other speech acts might not share *family* resemblance but still be sufficiently similar for re-scoring to occur. Depending on the content of P, it might be used to perform a promise. 'I will meet you at the airport tomorrow' might be part of a 'predictive family', more closely related to predictions, pledges, and promises than assertions proper. But this might still be *close enough* for the rule of accommodation to re-score an utterance from an assertion into a promise (and vice-versa). If my suggestion is right, this will be harder to do than re-scoring an assertion into a (more similar) guess, but if the locution is sufficiently *inexplicit* about its performative purpose, it will still be possible.

We can find similarities between speech acts along different axes, and several of these might help to enable re-scoring (or amendment). We might look to the conventions governing the correct performance of each candidate act, the obligations they enact (their conventional effects), and the types of locutions that might be used to perform them. We would essentially be looking for similar 'illocutionary patterns' (Sbisà, 2013). 'I *will* meet you at the airport tomorrow' and 'I *might* be able to meet you at the airport tomorrow' might both belong to the predictive family, but the former has a level of certainty and commitment that suits the illocutionary pattern of a promise; the latter does not. If both received the re-scoring response 'Thanks for promising to meet me at the airport tomorrow!', the former ('I *will*') is more likely to be successfully turned into a promise than the latter ('I *might*').

The OP in §2 is phrased very inexplicitly; 'I will cut taxes!' could be used to perform a wide range of speech acts including an assertion, suggestion, agreement, threat, promise, warning, and concession. These all, in varying ways, commit the speaker to cutting taxes. And the politician has the relevant social position to perform each of these (unlike if the OP was uttered by, say, a baseball player). As a result, the OP is open to being re-scored into (at least) these seven different speech acts.

The limitations of re-scoring through accommodation should hopefully now be a little clearer, albeit described in preliminary detail. If there is sufficient similarity between the conventions governing the relevant acts, the obligations they would enact, and the locution the speaker uses, then illocutionary force can be changed through accommodation. Explicit performatives (see Austin, 1962), with more defined locutions, will be harder to change: 'I hereby promise to meet you tomorrow' would be *very* difficult to accommodate as something other than a promise. Ceremonial speech acts often have such formalised conditions that they may be impossible to change. Pronouncements of marriage, for example, often require not only uttering specific words but also formalised uptake (the signing of witness statements). But most of our illocutionary activity is not so clear-cut, and many of our imperfectly performed speech acts are nonetheless successful (Johnson, 2020).

I have not said what counts as *sufficiently similar* for re-scoring to be achieved through accommodation. Re-scoring should seem very plausible in the examples described in §2,

whereas it should seem *implausible* that the OP could be re-scored into a proposal of marriage. For more difficult cases, we would need to determine the precise requirements of the acts under consideration; for some speech acts, philosophers have done (or are doing) this (Fraser, 1998 on threats; Green, 2013 on assertions; Lance and Kukla, 2013 on calls). We would also need to work out what counts as sufficient similarity (my guess is that the higher the stakes, the higher the degree of similarity required). Hopefully, further philosophical attention will shed light on these issues; for now, these initial comments should show that QT1-7 are within the limits of re-scoring illocutionary force through accommodation.

This also shows that even when quote-tweeted many times, the illocutionary force of an utterance does not entirely float free of the original utterance. The conventions governing various speech acts, and the rules governing the evolution of conversational score, keep the utterance and its force tethered together, ‘limiting the space of illocutionary possibilities’ (to borrow Lewiński’s 2021a phrase). One cannot make an 8th strike in baseball; ‘I will cut taxes!’ cannot be re-scored as a marriage proposal. But it can be re-scored as a promise, a suggestion, a threat, a warning, an agreement, and a concession. And the QT also adds the OP itself to the conversational score, anchoring the normative changes it makes in place.

By contrast, even though a lie may conversationally attribute the status of ‘having promised’ to the politician and change the conversational score to make it ‘correct play’ to treat them as having so promised, it does not add the OP to the score as an anchor. And so, the lie does not re-score anything; there is nothing to re-score. Unlike the norms enacted by the quote-tweet, those brought about by the lie are untethered to any utterance, and so can easily drift away.

To sum up so far: QT1-7 re-score the OP, making it the case that within each conversation the OP should be treated as having a different illocutionary force. If speech acts are individuated, and perhaps constituted, by their normative effects, this means that the OP now *has* that new illocutionary force (within this conversation). From a third-person perspective viewing all seven quote-tweets and their subsequent conversations at once, it would look like the OP no longer has just one illocutionary force, but rather seven.

5. Offline Pluralism and Communication Breakdowns

This account of illocutionary pluralism would be useful even if it only applied to online speech; the more communication that occurs online, the more philosophers of language should pay attention to it. But happily (perhaps), the structural features of quote-tweeting conversations that enable illocutionary pluralism (as described in §4) can occur offline too.

Imagine that the politician gives a speech at a fancy fundraiser. The final line of their speech uses the same words as the OP: ‘I will cut taxes!’. After the speech, canapes and drinks are served, and the audience breaks up into many small groups, each having their own conversation governed by its own conversational score. Some of these conversations discuss the speech, and seven begin with comments about the final line, the first with:

‘They just said they would cut taxes!’

The second with:

‘They just promised they would cut taxes!’

The third with:

‘They just threatened they would cut taxes!’

And so on, mirroring QT1-7. Each initial utterance treats the final line of the speech as a different type of speech act. This is tracked by the conversational score, making it correct play for other speakers to behave as though the final line *is* an assertion, a promise, a threat, and so on. Whatever the original illocutionary force of the utterance, in each conversation it has become correct to treat it as this new type of speech act.

Next, the politician begins ‘working the room’, joining each conversation to shake hands and accept congratulations. And within each conversation, they say something that ratifies the different characterisations of their final line:

‘Yes, I did; I think it’s a great idea!’

‘Yes, I promised I would cut taxes, so you should vote for me!’

‘Yes, if the opposition does not agree to decrease benefits, I will cut taxes!’

As noted earlier, the ‘re-scoring’ of the illocutionary force of the final line could be *blocked* with something like ‘I don’t think they really meant it as a promise.’ Participants would then need to decide how to treat the final line of the speech before continuing. The politician might be able to block more forcefully: ‘I didn’t intend that as a promise’ might resolve the matter if speakers are privileged with regards to determining the illocutionary force of their own utterances. But without an objection, it becomes correct play to treat the final line as the type of speech act mentioned at the start of the conversation (as per §4). As the change in conversational score alters norms governing the behaviour of participants in the conversation, the re-scoring changes the obligations and expectations that participants have for each other, and for the politician. So, the structural features of Twitter that enable my account of illocutionary pluralism can also occur offline, at least for public speech with a large audience.

There are some important differences in score-tracking (including of the illocutionary forces ascribed to an utterance) between each medium. Online, it is difficult to keep track of the participants in the conversation. Some people join part way through, notified of the discussion because one of their friends replied to the quote-tweet. Others leave part way through, while others might ‘lurk’, reading each post without commenting themselves (McDonald, 2021). Because the attention or presence of other participants cannot be tracked, speakers do not always know who they are (or might later be) talking to (boyd et al., 2010; Frost-Arnold, 2021; Goldberg, 2021).

By contrast, participants can literally see (or hear) who is involved in the conversations taking place over canapes and drinks. But people still join or leave these conversations. Mixing and mingling is part of the purpose of this kind of event. So, while the participants in

an offline conversation may be easily identifiable, they can *change*. When a new participant joins the conversation, they will not be aware of the current conversational score, including the re-scoring of the final line of the speech (whereas online, they might read the previous posts to work out what the score is). The new participant will have a different ‘mental scoreboard’ to the people they are talking to (Lewis, 1979). They might then say something that does not count as correct play, treating the final line as a threat when the current conversational score was tracking it as a promise. This is a kind of communication breakdown, or a ‘defective’ conversational score (see Stalnaker, 2014 on ‘defective’ common grounds).

When communication is working well, the ‘mental scoreboard’ of each participant and the conversational score itself will closely match. But the conversational score is not simply an amalgamation or tally of each participant’s ‘mental scoreboard’. Rather, it is an attempt to describe a kind of ‘objective’ score (Witek, 2015 and Camp, 2018 outline advantages of understanding accommodation and score as objective, at least for certain purposes). To return to Lewis’s metaphor of baseball: each player (and spectator) will have a mental scoreboard, comprised of beliefs about the score in the game. But what the score *actually is* does not depend on those beliefs. If a batter hits a fair ball, but the umpire calls it foul, the beliefs of the players (and outraged spectators) do not settle the score; the umpire’s call does (according to the rules of the game). They might appeal to the umpire to change the call, but the score is in this way ‘objective’.

Conversational score is somewhat similar. Participants will act in accordance with what they think the score to be, but they can be wrong. A new participant joining a conversation might have incorrect beliefs about the score; it would then be more likely that their utterances will be incorrect play. A student late to class might not know that it has been made permissible to use the professor’s first name, but the score would nonetheless track first name use as ‘correct play’. The score can change even if some (or many) participants did not realise it.

Ideally, when participants conversing after the politician’s speech have different beliefs about what the score is, they would recognise the communication breakdown and attempt to redress it. There are several possible solutions. One of the other participants could *correct* the new speaker— ‘They weren’t threatening, they promised to cut taxes!’. The new speaker could play along—they would recognise that it is now correct play to treat the final line as a promise, regardless of their private beliefs about it. Beliefs about what the score is would then converge with each other, and with the actual score, even if beliefs about what the politician intended remain at odds.

The newcomer might instead make a blocking move— ‘I’m pretty sure they were threatening the pro-tax politicians!’. Participants would then need to *negotiate* to come to agreement about how to treat the illocutionary force of the final line (see also Lewiński, 2021a on ‘dilemmatic’ deliberation). Going forward, the group might agree to treat the final line as a threat or a promise or leave it undetermined (agreeing to disagree). Correction, and negotiation, require conversational contributions and in contributing in this way, change the score itself as well as the beliefs participants have about the score, hopefully bringing all into alignment. The breakdown is repaired. Alternatively, the breakdown might go unrecognised.

An utterance like ‘I was so happy with that final line!’ might get agreement from all participants, even though the original participants were happy that the politician promised to cut taxes, while the newcomer was happy that the politician was taking a strong stance in threatening their opponents. Participants might then end up ‘talking past’ each other.

Importantly, this does not turn on who is *right* about the score. The original participants might ‘play along’ with the newcomer’s utterance that treats the final line as a threat when the score tracked it as a promise. Their playing along will update the score, and so it becomes correct play to treat the final line as a threat even though the person who made this conversational move was ‘wrong’ about the score when they spoke.

In baseball, there is an umpire who decides if a move is ‘correct play’ or not. In conversation, there is no umpire—but we should still think about conversational score from an umpire-like perspective (Sbisà, 2013 refers to something like this as ‘the analyst’). If conversations *did* have an umpire, equipped with a handbook outlining all the rules of conversation, and a recording of the conversation thus far, they would be able to work out what the score actually is, even when participants have different beliefs about it. So, when there is disagreement about the score, whose mental scoreboard matters? Just as in baseball, it is the umpire’s—although usually, the conversational umpire is hypothetical (at times, such as a court case that turns on whether an utterance constitutes a bribe, there may be someone—the judge—who adopts this umpiring role). And, as a hypothetical being, they could track the scores of all seven conversations described in §2 (or §5) at once. From their perspective, the OP (or final line) would appear to function as seven different speech acts.

The gap between what the score *is* and what participants believe it *should be* allows for other communicative oddities. For example, a participant might leave one conversation to join a new one, and then contribute in a way that *conflicts* with the score of the conversation they left. Perhaps they ‘played along’ with their original conversation, acting as though the final line was a promise while privately thinking that it was a threat. Once they are in a new conversation, they decide to treat the final line as a threat, rather than a promise (‘I’m surprised they threatened to cut taxes!’)—matching their private beliefs rather than prior conversational behaviour.

This, I think, should seem strange. A friend who followed the speaker from one conversation to the next would likely be surprised or confused when the speaker suddenly changed their treatment of the politician’s utterance. To return to another example, if a professor tells a class of students to use their first name, this permission will still govern discussion in the class the following week. There are restrictions on this—the context matters. Students would rightly use the professor’s first name when in class the next week, but not in more formal settings. The point here is that norms enacted during a conversation do not necessarily end with the conversation itself and can follow us into conversations we join in the future. Precisely when and how this occurs is going to be governed by other rules and norms, and there is not space here to explain in detail how this might work—I hope to have shown only that this kind of lingering norm enactment is plausible, and compatible with my analysis of conversational re-scoring.

In these dynamic offline conversations, it is difficult for participants to perfectly track the score—but there is still a fact-of-the-matter about what the score is and thus what counts as correct play. Sometimes it may take a hypothetical analyst to work out exactly what it is, while real-world speakers muddle through with imperfect information and (at times) mistaken beliefs about the score. Nonetheless, this section has explained how the communicative structures enabling illocutionary pluralism as described in §4 can also occur offline.

6. Potential Objections

6.1 Intentionalism

One might worry that this is not illocutionary pluralism, but rather a case of mistaken identity. Whether or not the audience (on Twitter or at the fundraiser) realised it, maybe the speaker *intended* to make an assertion. On some influential accounts of speech act theory, an utterance can only constitute the intended speech act (Sbisà, 2009 calls this the ‘received’ view). If this is right, the audience have simply made a mistake about what speech act the OP performs—or in this case, several different mistakes. These mistakes from the audience might mean that the speaker fails to perform the speech act of assertion, but do not mean that the speaker has in fact performed a different speech act.

This mirrors a long-running debate in speech act theory. What settles the illocutionary force of an utterance? *Intentionalists* argue that the speaker’s intention determines the illocutionary possibilities of an utterance (Sbisà, 2009 attributes this to Strawson, 1964; Searle, 1975; Bach and Harnish, 1979; see McDonald, 2022 for further discussion). If the speaker did not intend to make a promise, what happens afterwards cannot turn their utterance into a promise. *Conventionalists* argue that the speaker’s intention does *not* determine the force of their utterance; instead, if the utterance (and the audience’s response to it) satisfies the conventions associated with a speech act (Sbisà, 2013 calls these ‘illocutionary patterns’), then that is the act the utterance performs. If an utterance, and the response to it, seems to match the conventions of a promise, then it constitutes a promise even if it was intended as an assertion (e.g., Lance and Kukla, 2013; Sbisà, 2013; Kukla, 2014—although the description here glosses over differences in their views).¹⁶

An intentionalist might argue that despite the conversational score being updated to make it correct play to treat the OP as a different type of speech act, the OP retains the illocutionary force of an assertion because that is what the speaker intended. The audience has made a mistake, and they do not get to refashion the illocutionary force of the OP to make it right.

Here, I do not intend to argue for either side of this debate. My (untested) hypothesis is that there *is* no unified set of conditions governing the performance of all speech acts. In some

¹⁶ While the account of quote-tweeting in this paper is roughly on the conventionalist side of the debate, it is not the audience’s uptake that determines illocutionary force but rather the conversational score.

cases, the speaker's intention might determine illocutionary force; in others, it might be the uptake of the audience. Perhaps for perfect performance, intention and uptake must match—but as Johnson (2020) argues, imperfectly performed speech acts can still be successful. Maybe the quote-tweets in §2 turn the OP into seven imperfect, but nonetheless successful, speech acts. But for now, what might I say to the intentionalist who holds that the speaker's intention determines the illocutionary potential of their utterance regardless of the conversational score?

A preliminary response: in the scenarios described in §2 and §5, the politician responds in a way that seems to ratify the re-scoring of their OP. This would suggest to participants that it is the speaker's intention (at least now) that their utterance constitutes a suggestion, a warning, a concession, and so on. Maybe it *was* intended as an assertion, but not anymore. This mirrors retraction and amendment (Caponetto, 2020). If I make a promise, then retract it, it is no longer my intention that my original utterance be taken to function as a promise. If the retraction is successful, I am no longer bound by the obligations enacted by a promise. In the conversations considered in this paper, it is someone else who re-scores the utterance, but as the speaker indicates their agreement, the outcome, I have suggested, is the same.

Perhaps the intentionalist will be satisfied with this—as the speaker seems to intend that their audience *now* treat their utterance as performing a different type of speech act, then it does in fact perform that new speech act. But that only gets us so far. After all, I argued that this illocutionary pluralism happens *whether or not* the speaker responds (although they may be in a privileged position to block the re-scoring of their utterance).

I think that the intentionalist could (and should) agree with a lot of my account anyway. They could accept that following the quote-tweet, it becomes conversationally appropriate to treat the OP as a new type of speech act (although it is technically still an assertion). They could agree that the norms it now enacts are those of a promise, a warning, a threat, and so on (despite it being technically still an assertion). There are many similar descriptions they could endorse: the OP now 'functions as', 'is near enough to', 'influences permissibility like', etc., the speech act characterised by the quote-tweet (but technically is still an assertion). The unlucky participants in this conversation mistakenly treat the OP as the wrong type of speech act but are nonetheless acting in accordance with the rules of conversation.

This seems promising, at least to me. We might be at an impasse regarding the source of illocutionary force, but the intentionalist could agree that the OP now pretty much does the job of its re-scored speech act (within this conversation). I would have shown how a quote-tweet (or similar offline comment) can reshape the norms enacted by an utterance, altering the permissibility conditions governing a conversation such that participants are supposed to respond to it as if it were a new type of speech act (even if they are mistaken about what speech act the utterance 'really' is). That said, if intentionalism allows that the illocutionary force of an utterance can be so far removed from its normative effects, it seems (to me) a less helpful description of how we do things with words.

So, in the scenario outlined in §2, the politician signals that their intention is compatible with the re-scoring of the OP. When they are *not* part of the subsequent conversation, the quote-

tweet still changes how participants are supposed to respond to the OP. Intentionalists may prefer an alternative term to describe this rather than *illocutionary pluralism* (functional pluralism? conventional pluralism?) but will hopefully agree that even if the audience is mistaken about the illocutionary force of the utterance, they are correctly following the rules governing this conversation.

6.2 Illocutionary Relativism

But might this not, in fact, be illocutionary *relativism*, rather than *pluralism* (Johnson, 2019, Lewiński, 2021a)? I have suggested that the illocutionary force of an utterance is, in a way, *relative* to the conversational score. I recognise the disadvantageous branding of ‘relativism’, but if the shoe fits, I shall wear it. Lewiński (2021a) goes as far as to try to ‘block’ analysis of illocutionary pluralism in terms of illocutionary relativism, where illocutionary force might be ‘relative’ to the uptake of idiosyncratic or uncooperative listeners. While I do argue that utterances may be ascribed with unintended illocutionary force, this is not done through unusual patterns of uptake but rather the (perhaps) more ‘stable’ mechanism of conversational score. So, is the illocutionary pluralism I have described actually illocutionary relativism in disguise?

Johnson (2019) and Lewiński (2021a) distinguish illocutionary pluralism from ‘illocutionary relativism’. The idea here is that if, following Sbisà (2013), the conventional effects of an utterance are dependent on the addressee’s uptake, then the illocutionary act performed with an utterance token will be *relative* to its addressee. If I say ‘Take the rubbish out’, Tionne might give it uptake as an order, while Lisa gives it uptake as a request. If uptake settles illocutionary force (see also Kukla, 2014), then the illocutionary force of my utterance would be relative to the uptake it secures from each target—in other words, illocutionary relativism (rather than pluralism).¹⁷

But there is a crucial difference between this and my account of the illocutionary force of the OP (and between my account and Sbisà’s (2013) second type of illocutionary pluralism). The illocutionary force is not altered by the uptake it secures, but rather its status in the conversational score. A listener might give the OP uptake as a promise, but if the conversational score counts it as an assertion, it would be incorrect play to treat it as a promise. In the same way, once the score changes in a game of football, even if the losing team thinks that the goal should not have been allowed, the score nonetheless changes, and it is now correct play to act as though the goal was scored. Some sports allow for rulings to be reviewed; but unless this happens, the goal stands, and play continues accordingly. Similarly, a conversational participant needs to *actually object* if they disagree with the characterisation of the OP.

¹⁷ Philosophers disagree about the nature of uptake, extending debate between intentionalists and conventionalists (see Austin, 1962; Strawson, 1964; Sbisà, 2009; Kukla, 2014; McDonald, 2022). The analysis here is, I hope, general enough to capture different theories of uptake.

So, while I suggest that illocutionary force is ‘relative’ to the conversational score, this is not how illocutionary relativism has previously been described. From the point of view of conversational participants involved in a single conversation, the OP only has one illocutionary force, even if this is not the force it had a minute ago. But from the hypothetical analyst’s view of all seven conversations, the OP now appears to have seven different illocutionary forces. Online, it is possible to participate in all seven conversations at once. A perceptive participant would, hopefully, realise that the score treats the OP as a different type of speech act in each conversation and modify their contributions accordingly (either to fit the score of each conversation, or by attempting to re-score the OP in each conversation so that they align). Within each conversation, there is a fact-of-the-matter about how participants should treat the OP. As a result, I think that ‘illocutionary pluralism’ is a better label than ‘illocutionary relativism’, although the importance of this distinction might be more a matter of marketing than deep theoretical implications.

6. Conclusion

What chaos has been unleashed? A speaker tweets (or utters) a seemingly simple assertion. Later, they appear to have performed seven different, and in some cases possibly incompatible, speech acts. This is a far cry from Austin’s (1962) stable and measured speech act theory. But alas, communication can often be messier than philosophers of language might wish. My solution to the puzzle outlined in §2 may not seem ideal; all the better, perhaps, for explaining non-ideal speech acts in the non-ideal online world.

A quote-tweet can refashion the normative force of an original post, changing *correct play* for responses to it. A conversational contribution can likewise change the illocutionary force that participants ascribe to a previous utterance. Reshaping the normative landscape is, I think, the core business of speech acts. They enact, alter, and remove normative obligations—and this process is temporally extended. So, it should not surprise us that an utterance might begin life as a single speech act but end up counting as more than one. Social media, and Twitter especially, are designed in a way that enables this kind of illocutionary pluralism to occur at a viral scale. And, for better or worse, the same can happen offline.¹⁸

¹⁸ Thanks to Marcin Lewiński and Ten-Herng Lai, as well as three anonymous reviewers, for their helpful feedback, suggestions, and comments.

References

- Austin, J. L. (1962). *How to do things with words*. Oxford University Press.
- Bach, K. & Harnish, R. M. (1979). *Linguistic communication and speech acts*. Harvard University Press.
- boyd, d., Golder, S., & Lotan, G. (2010, January). Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *2010 43rd Hawaii international conference on system sciences* (pp. 1-10). IEEE.
- Camp, E. (2018). Insinuation, common ground, and the conversational record. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 40-66). Oxford University Press.
- Caponetto, L. (2020). Undoing things with words. *Synthese*, 197(6), 2399-2414.
- Cappelen, H., & Lepore, E. (2008). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. John Wiley & Sons.
- Clark, H. (1996). *Using Language*. Cambridge University Press.
- Egan, A. (2009). Billboards, bombs and shotgun weddings. *Synthese*, 166(2), 251-279.
- Elliott-Maksymowicz, K., Nikolaev, A., & Porpora, D. (2021). How much can you say in a tweet? An approach to political argumentation on Twitter. *Humanities and Social Sciences Communications*, 8(1), 1-9.
- Fraser, B. (1998) Threatening Revisited. *Forensic Linguistics*, 5(2), 159-173.
- Frost-Arnold, K. (2021). The Epistemic Dangers of Context Collapse Online. In J. Lackey (Ed.), *Applied epistemology* (pp. 437-456). Oxford University Press.
- Green, M. (2013). Assertions. In M. Sbisà & K. Turner (Eds.), *Pragmatics of Speech Actions, Vol. II of the Handbook of Pragmatics* (pp. 387-410). De Gruyter.
- Goffman, E. (1963). *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Free Press.
- Goldberg, S. C. (2021). The Promise and Pitfalls of Online ‘Conversations’. *Royal Institute of Philosophy Supplements*, 89, 177-193.
- Johnson, C. R. (2019). Investigating illocutionary monism. *Synthese*, 196(3), 1151-1165.
- Johnson, C. R. (2020). Mansplaining and illocutionary force. *Feminist Philosophy Quarterly*, 6(4).
- Kukla, R. (2014). Performative force, convention, and discursive injustice. *Hypatia*, 29(2), 440-457.
- Lance, M., & Kukla, R. (2013). Leave the gun; take the cannoli! The pragmatic topography of second-person calls. *Ethics*, 123(3), 456-478.

- Langton, R. (2012). Beyond Belief: Pragmatics in Hate Speech and Pornography. In I. Maitra & M. K. McGown (Eds.), *Speech and harm: Controversies over free speech* (pp. 72-93). Oxford University Press.
- Langton, R. (2018). Blocking as counter-speech. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 144-164). Oxford University Press.
- Langton, R., & West, C. (1999). Scorekeeping in a pornographic language game. *Australasian Journal of Philosophy*, 77(3), 303-319.
- Levinson, S. C. (2011). Action formation and ascription. In T. Stivers & J. Sidnell (Eds.), *Handbook of conversation analysis* (pp. 103-130). Wiley-Blackwell.
- Lewiński, M. (2021a). Illocutionary pluralism. *Synthese*, 199(3), 6687-6714.
- Lewiński, M. (2021b). Speech Act Pluralism in Argumentative Polylogues. *Informal Logic*, 41(3), 421-451.
- Lewis, D. (1979). Scorekeeping in a language game. In R. Bäuerle, U. Egli, & A. Stechow (Eds.), *Semantics from different points of view* (pp. 172-187). Springer.
- Maitra, I. (2012). Subordinating speech. In I. Maitra & M. K. McGown (Eds.), *Speech and harm: Controversies over free speech* (pp. 94-120). Oxford University Press.
- Marsili, N. (2021). Retweeting: its linguistic and epistemic value. *Synthese*, 198(11), 10457-10483.
- Marwick, A. E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society*, 13(1), 114-133.
- McDonald, L. (2021). Please Like This Paper. *Philosophy*, 96(3), 335-358.
- McDonald, L. (2022). Reimagining Illocutionary Force. *The Philosophical Quarterly*, 72(4), 918-939.
- McGowan, M. K. (2004). Conversational exercitives: Something else we do with our words. *Linguistics and Philosophy*, 27(1), 93-111.
- McGowan, M. K. (2009). Oppressive speech. *Australasian Journal of Philosophy*, 87(3), 389-407.
- McGowan, M. K. (2019). *Just words: On speech and hidden harm*. Oxford University Press.
- Oishi, E. (2022). Illocutionary-act-type sensitivity and discursive sequence: An examination of quotation. *Intercultural Pragmatics*, 19(3), 381-406.
- Popa-Wyatt, M. (Forthcoming). Norm-Shifting through Oppressive Acts. In K. Jones, G. Restall, F. Schroeter, & L. Schroeter (Eds.), *Mind, Language, and Social Hierarchy: constructing a shared social world*. Oxford University Press.

- Saul, J. (2018). Dogwhistles, political manipulation, and philosophy of language. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 360-383). Oxford University Press.
- Sbisà, M. (2009). Uptake and Conventionality in Illocution. *Lodz Papers in Pragmatics*, 5(1), 33-52.
- Sbisà, M. (2013). Some remarks about speech act pluralism. In A. Capone, F. Lo Piparo, & M. Carapezza (Eds.), *Perspectives on pragmatics and philosophy* (pp. 227-244). Springer.
- Schiller, H. I. (2021). Is that a threat? *Erkenntnis*, 86(5), 1161-1183.
- Schegloff, E. A. (1968). Sequencing in Conversational Openings. *American Anthropologist*, 70(6), 1075-1095.
- Searle, J. R. (1975). Indirect speech acts. In P. Cole & J. Morgan (Eds.), *Speech acts* (Syntax and Semantics Volume 3, pp. 59-82). Brill.
- Simpson, R. (2013). Un-ringing the Bell: McGowan on Oppressive Speech and the Asymmetric Pliability of Conversations. *Australasian Journal of Philosophy*, 91(3), 555-575.
- Stalnaker, R. (2002) Common Ground. *Linguistics and Philosophy*, 25(5/6), 701-721.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Strawson, P. F. (1964). Intention and convention in speech acts. *Philosophical Review*, 73(4), 439-460.
- Vosoughi, S., & Roy, D. (2016, March). Tweet acts: A speech act classifier for twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, 10(1), 711-714.
- Witek, M. (2015). Mechanisms of illocutionary games. *Language and Communication*, 42, 11-22.
- Witek, M. (2019). Accommodation in linguistic interaction: On the so-called triggering problem. In P. Stalmaszczyk (Ed.), *Philosophical Insights into Pragmatics* (163-192). De Gruyter.
- Zhang, R., Gao, D., & Li, W. (2011). What are tweeters doing: Recognizing speech acts in twitter. *Analyzing Microtext*, WS-11-05 of AAAI Workshops.