Epistemic resistance moves *

Justin Bledin
Johns Hopkins University

Kyle Rawlins
Johns Hopkins University

Abstract This paper introduces and analyzes a new kind of non-acceptance, non-disagreeing move: *resistance*. We focus in particular on *attention-targeted* resistance facilitated by epistemic possibility claims. In this response type, we suggest, an agent draws attention to some subsidiary issue that they think might cause an interlocutor to withdraw a previous commitment. We develop a granularity model of attention where drawing attention in discourse can refine the space of possibilities under consideration and consequently lead to changes in view.

Keywords: discourse, resistance, rejection, disagreement, epistemic modals, attention

1 Resistance moves

The two most obvious ways that an agent can respond to an assertion are to agree or to disagree with it. This paper explores a response type that does not slot neatly into either of these categories: *resistance moves*, witnessed below.

- (1) **Party.** Alice and Bob are on their way to an opening party at the Whitney Museum. They hope to see their friend Charlie there who has recently broken up with his girlfriend Diana.

A: Charlie is coming to the party.

B: Okay. / That's not true! (agreement / disagreement)

B': Diana might be there. (epistemic resistance move)

In simple disagreement, a speaker directly indicates that they are unwilling for the proposed claim to be incorporated into the common ground—the claim is rejected outright. That is not what happens with resistance. Resistance moves instead indicate unwillingness to incorporate a proposal before some subsidiary issue has been explored. In (1), B wants to ensure that A is aware of some facts about Diana that (in this context) might bear on the truth of A's claim.

* For discussion of this topic we are grateful to audiences at JHU, ILLC, the New York Philosophy of Language Workshop, the Unstructured Conference at Rutgers in April 2016, SALT 26 at UT Austin, as well as María Biezma, Fabrizio Cariani, Lucas Champollion, Ivano Ciardelli, Jeroen Groenendijk, Daniel Hoek, Harvey Lederman, and Floris Roelofsen.

In this paper we focus on *epistemic resistance moves* where the resistance is mainly signaled with an epistemic possibility modal. The category of resistance, once defined, turns out to be much larger than this, as illustrated by a sampling of responses in this Bridges scenario.

- (2) **Bridges.** Charlie is supposed to meet Alice and Bob for dinner in Manhattan and he calls to tell them that he is just leaving his place in Brooklyn. A little later, Alice and Bob hear on the radio that westbound trains are delayed and there is heavy traffic on the Brooklyn Bridge.

A: Charlie is going to be late for dinner.

B: He might take the Manhattan Bridge. (epistemic resistance move)

B': What if he takes the Manhattan Bridge?

B'': Is there traffic on the Manhattan Bridge as well?

B''': The Manhattan Bridge isn't too far from his home.

Resistance moves are in no way incompatible with the resisted claim. Unlike corrections and other rejection moves, they do not put a conversation in a *crisis state* (cf. Farkas & Bruce 2010) where any issues under discussion can be settled only after the original speaker or rejector retracts one of their commitments, or they agree to disagree.¹ This is evident from follow-ups in a resistance discourse:

- (3) (Bridges continuations)

B: He might take the Manhattan Bridge.

A: Good point. If he takes the Manhattan Bridge, he might not be late.

A': I was overlooking that option. However, if Charlie takes the Manhattan Bridge, he's still going to be late.

A'': I already considered that option.

Our proposal about epistemic resistance moves is that they work by targeting speakers' *attention* (Ciardelli, Groenendijk & Roelofsen 2011; de Jager 2009; Franke & de Jager 2010). In particular, they draw attention to some issue that the resister thinks the resistee might have overlooked. We model attention using a Yalcin (2011)-style granularity model in which overlooking an issue makes an agent unable to 'see' propositions relevant to that issue; shifting attention onto that issue can therefore cause certain latent premises to become visible, leading (potentially) to a change in what the resistee is willing to publicly commit to. At the level of discourse dynamics,

¹ Compatibility differentiates them from Khoo 2015 on modal disagreement, as well as similar disagreement data in Klecha 2014. Also, while resistance moves may overlap with clarification requests (Ginzburg 1998, 2012), many resistance moves are not requests at all.

we suggest that (to borrow a term from [Roberts 1996](#)) resistance moves indicate a *strategy* for coordinating on whether to incorporate the content of some assertion into the common ground. We provide an implementation of this using a version of [Farkas & Bruce's 2010](#) notion of the *Table*. In what follows, we first develop the proposal about attention in §2, then incorporate this into a dynamics for resistance in §3, and conclude in §4.

2 Resistance and attention

In this section we develop a granularity model of attention based on ideas in [Hulstijn 2000](#), [Yalcin 2008, 2011](#), [de Jager 2009](#), [Franke & de Jager 2010](#), and [Fritz & Lederman 2015](#). But first, let us briefly mention a potentially simpler alternative.

In line with [Lewis 1979](#) and [Stalnaker 1984](#) (and more recently, [Rawlins 2010](#) and [Klecha 2014](#)), it is tempting to say that after Bob resists Party (1), Alice's sphere of *relevant alternatives* expands outward to include previously ignored possibilities in which Diana comes to the party. Inattention, on this view, amounts to tacitly assuming that certain maximal possibilities are out of the running: certain possible worlds are simply ignored (or *circumscribed*; [McCarthy 1980](#)) by the agent.

But this idea, while plausible in some cases, is not generally adequate. In Party (and other resistance examples), Alice needn't be making any tacit assumptions one way or the other about whether Diana is coming to the party.² Rather, Alice might simply not have factored the impact of this issue into her reasoning. This leads to a central question: how can someone fail to attend to an issue and at the same time still be willing in principle to accept any resolution of this issue?

2.1 Modeling attention

Our proposal here is that attention states for purposes of understanding resistance should be thought of in terms of *granularity*. The starting point is Lewis's (1988) theory of subject matters. Let $\mathcal{W}$ be a nonempty set of possible worlds.

- (4) **Subject matters.** A *subject matter* $M \subseteq \mathcal{W} \times \mathcal{W}$ is a total equivalence relation over $\mathcal{W}$.³ (Let M^* be the set of subject matters.)

Every subject matter $M \in M^*$ partitions $\mathcal{W}$ into mutually exclusive equivalence

² It might be strictly compatible with Alice's initial view of the world that both Charlie and Diana are coming since her knowledge about their breakup is inactive, and, in any case, this knowledge might make only possibilities in which both come less plausible without ruling them out completely.

³ I.e. M is reflexive, symmetric, and transitive. Like Lewis, we do not count $\mathcal{W} \times \mathcal{W}$ as a genuine subject matter. See [Yablo 2014](#) for a more general account where these needn't be total or transitive.

Figure 1 From left to right: inclusion, disjointness, and non-inclusive overlap.

classes of worlds that jointly exhaust this space.⁴ A cell $[w]_M = \{v : M_{wv}\}$ of the partition $\mathcal{W}/M$ determined by this subject matter groups together maximally specific possible states of reality that are exactly alike with respect to it. For example, we can model the subject of whether Diana is coming to the party with the relation M_D where M_{Dwv} just in case either w and v are both D -worlds in which Diana comes to the party or neither w nor v is a D -world. Someone attending to M_D distinguishes between worlds in different cells of the bipartition $\mathcal{W}/M_D$. In (5), we provide some helpful mereological notions for subject matters. This are illustrated visually in Figure 1.

- (5) a. **Inclusion.** M_1 (mereologically) *includes* M_2 just in case: $M_1 wv$ only if $M_2 wv$. We will also say this: M_1 *refines* M_2 and M_2 is a *coarsening* of M_1 .
- b. **Overlap.** M_1 and M_2 *overlap* just in case both of these subject matters (mereologically) include a common subject matter M_3 . If M_1 and M_2 do not overlap, they are *disjoint*.
- c. **Summation.** The *sum* $M_1 + M_2 + \dots$ is the least subject matter, should this exist, which includes each of $M_1, M_2, \dots$ and is included in any other subject matter that includes these matters.

Like Fritz & Lederman 2015, we assume that: (i) anybody attending to a subject matter M is also attending to every coarsening of M ; (ii) anybody attending to subject matters $M_1, M_2, \dots$ is also attending to the least subject matter $M_1 + M_2 + \dots$ (which exists); (iii) anybody engaged in conversation is attending to at least one subject matter. An agent's attention to subject matters can therefore be uniquely characterized by the sum Π of each of the subject matters that she is attending to. Whether this agent is attending to a subject matter M depends on whether a certain mereological relation holds between Π and M :

- (6) **Attention to subject matters.** An agent in attention state Π is *fully attending* to M if Π includes M , *fully overlooking* M if Π and M are disjoint, and *partially attending* to M otherwise.

⁴ This is the same kind of representation often used in the semantics literature for question denotations (Groenendijk & Stokhof 1984; Groenendijk 1999). We return to this comparison in §2.2.

Figure 2 Visibility and views.

To elaborate, note that some propositions will carve up $\mathcal{W}$ in a way that respects the granularity of $\mathcal{W}/\Pi$ while others will cut across cells in this partition. Visibility is illustrated graphically in Figure 2-a.

- (7) Proposition $P \subseteq \mathcal{W}$ is *visible* at M just in case: Mwv only if ($w \in P$ iff $v \in P$).

The definition of attention can now be restated in terms of visibility:

- (8) An agent in attention state Π is *fully attending* to M if every contingent proposition visible at M is also visible at Π , *fully overlooking* M if no such proposition is visible at Π , and *partially attending* to M otherwise.

Since each proposition P visible at M corresponds to a way that things might be like with respect to this subject matter, this is just to say that an agent is (fully) attending to M if for (every) some way in which the worlds in $\mathcal{W}$ can differ with respect to M , this agent distinguishes between any worlds in $\mathcal{W}$ that differ in this way.

To adapt some terminology from Yalcin 2011, an agent whose attention to subject matters is representable by Π will view reality at the “modal resolution” $\mathcal{W}/\Pi$. From this agent’s perspective, each cell $[w]_\Pi$ in $\mathcal{W}/\Pi$ is a kind of “pixelated” world, a distinct candidate for the way things are. Following the common practice of representing an agent’s knowledge and belief states by the set of open possibilities that are compatible with what she knows and believes, we define “views” as follows:

- (9) **Views.** The *view* of an agent in attention state Π is a set of possible worlds $\Delta \subseteq \mathcal{W}$ that is visible at Π . The cells in Δ/Π explicate the possibilities that are compatible with what this participant accepts modulo Π .

An example view, relative to some space of visible propositions, is illustrated in Figure 2-b. Applying this model to our Party case, Alice’s discourse-initial attention state and view can be explicated using a relation Π disjoint with M_D and a set Δ of C -worlds in which Charlie is coming that is visible at Π . After Bob resists, her attention state shifts to $\Pi + M_D$ and worlds in which Diana comes to the party

and worlds in which she does not come are differentiated in Alice's thought (more realistically, Alice's updated attention state will strictly include $\Pi + M_D$ since she will also become sensitive to a cluster of topics that she cognitively associates with M_D). Against the more fine-grained resolution $\mathcal{W}/\Pi + M_D$, Alice's updated view now leaves open the chunked possibility that Charlie does not come to the party.

How can we systematically derive this attention-induced change in her view? This is where things get tricky. In related work in progress, we argue that the granularity model of attention is especially well-suited for handling the kind of memory retrieval exhibited in cases like Party. The strategy pursued there is to assign agents caches of stored background information that are “filtered” through their awareness states to obtain their current views of reality (cf. [Franke & de Jager 2010](#)). In Party, for example, Alice's stored information might incorporate that Charlie is not one to miss a good party, that Charlie said he was coming to the party a few weeks ago, that Charlie and Diana subsequently had a messy breakup, that Charlie now wants nothing to do with Diana, etc. While the former information supporting Charlie's attendance is initially visible and informs Alice's view, the latter information about Diana is not. After Alice attends to whether Diana is coming, however, these additional ‘latent’ premises become accessible and percolate up to Alice's awareness. She changes her mind on their basis.

Fortunately, we can make substantial headway on our current project without getting into more details. As a placeholder for a fuller dynamic account of attention-induced change in view, we will here simply hardwire this into our model by working with a richer representation of an agent's cognitive state consisting both of a subject matter Π explicating her current attention to subject matters and of a *change function* δ mapping each subject matter $M \in M^*$ to a set of worlds visible with respect to it. This δ -function determines the agent's current view Δ at Π and also encodes how her view can potentially vary along with her attention:

(10) **Acceptance states**

An agent's *acceptance state* $\langle \delta, \Pi \rangle$ consists of a function $\delta : M^* \rightarrow 2^{\mathcal{W}}$ where $\delta(M) \subseteq \mathcal{W}$ is visible at M , and a subject matter $\Pi \in M^*$. The agent's current view is $\delta(\Pi)$.

We can now say this: if Alice is initially in the acceptance state $\langle \delta, \Pi \rangle$ and she comes to attend to M_D , then her updated state is $\langle \delta, \Pi + M_D \rangle$. Bracketing off any information that Alice acquires as a result of Bob's resistance move—about his own state of mind or anything else for that matter—her response to this challenge will depend on $\delta(\Pi + M_D)$. If $\delta(\Pi + M_D)$ includes only *C*-worlds in which Charlie comes, then Alice will stand by her original assertion. But if $\delta(\Pi + M_D)$ includes a world in which Charlie does not come, then Alice will conditionalize or retract.

2.2 Attention in discourse

Moving along, it is time to get our resolution-sensitive agents talking to one another. Our starting point for incorporating attention into context will be Stalnaker's (1978, 2002, 2014) classic theory on which a discourse unfolds against a background of publicly available information: the *common ground*. In its most recent incarnation (Stalnaker 2014), the common ground consists of propositions that are *commonly accepted* by the members of a conversation: P is common ground just in case everyone accepts P for the purposes of the conversation, everyone accepts that everyone accepts P , everyone accepts that everyone accepts that everyone accepts P , and so on ad infinitum.

- (11) **Frames** v. 1. A *frame* $\mathcal{F} = \langle \mathcal{W}, \{\Delta_a : a \in \text{Agt}\} \rangle$ consists of a nonempty set of worlds $\mathcal{W}$ and a function $\Delta_a : \mathcal{W} \rightarrow 2^{\mathcal{W}}$ for each agent $a \in \text{Agt}$ where $\Delta_a(w) \subseteq \mathcal{W}$ explicates the possibilities that are compatible with what a accepts in w .

Δ_a must satisfy **introspection**:

For all $w, v \in \mathcal{W}$, if $v \in \Delta_a(w)$, then $\Delta_a(v) = \Delta_a(w)$.

- (12) **Acceptance**. Agent a accepts $P \in 2^{\mathcal{W}}$ in w just in case $\Delta_a(w) \subseteq P$.

Introspection ensures both that if one accepts P then one accepts that one accepts P , and that if one does not accept P then one accepts that one does not accept P . The Stalnakerian notion of common ground can now be defined as follows:

- (13) **Accessibility**. A sequence of worlds $z_1, \dots, z_n$ in $\mathcal{W}$ is an *accessibility path* iff for each adjacent pair $\langle z_i, z_{i+1} \rangle$ in this sequence, $z_{i+1} \in \Delta_a(z_i)$ for some $a \in \text{Agt}$.

$\mathcal{R}^*_{wv}$ iff there is an accessibility path $z_1, \dots, z_n$ with starting point $z_1 = w$ and endpoint $z_n = v$.

- (14) **Context sets** v. 1. The *context set* in w is $CS(w) = \{v : \mathcal{R}^*_{wv}\}$.

- (15) **Common ground** v. 1. Proposition $P \in 2^{\mathcal{W}}$ is *common ground* among Agt in w just in case $CS(w) \subseteq P$.

It is easy to verify that these definitions formalize the above acceptance hierarchy. As for dynamics, Stalnaker 1978 famously proposes that the characteristic effect of a successful assertion is to add the proposition P communicated with this speech act to the common ground. The resulting frame is $\langle \mathcal{W}, \{\Delta_a + P : a \in \text{Agt}\} \rangle$ where $[\Delta_a + P](w) = \Delta_a(w) \cap P$. The updated context set at each $w \in \mathcal{W}$ is $CS(w) \cap P$.⁵

⁵ A more common tack is to remove non- P worlds from $\mathcal{W}$ (cf. Plaza's (1989) *Public Announcement*

In the rest of this section, we introduce attention into this Stalnakerian framework, culminating with notions of mutual, distributed, and public attention. Running with our resolution-sensitive model of acceptance states, a natural move is to replace the frames in (11) with the following structures:

- (16) **Frames** v. 2. A *frame* $\mathcal{F} = \langle \mathcal{W}, \{\delta_a, \Pi_a : a \in \text{Agt}\} \rangle$ consists of a nonempty set of worlds $\mathcal{W}$, a function $\delta_a : \mathcal{W} \rightarrow (2^{\mathcal{W}})^{M^*}$ for each agent $a \in \text{Agt}$ mapping each world $w \in \mathcal{W}$ to a function $\delta_a(w)$ from subject matters to visible sets of worlds, and another function $\Pi_a : \mathcal{W} \rightarrow M^*$ for each agent $a \in \text{Agt}$ mapping each world $w \in \mathcal{W}$ to a subject matter $\Pi_a(w)$.

Letting $\Delta_a(w) = \delta_a(w)(\Pi_a(w))$, δ_a and Π_a must satisfy **introspection**:
For all $w, v \in \mathcal{W}$, if $v \in \Delta_a(w)$, then $\Delta_a(v) = \Delta_a(w)$ and $\Pi_a(v) = \Pi_a(w)$.

Agents must now be introspective regarding both information and attention.

We also want to generalize the Stalnakerian statics in another respect. On the original picture, whether a discourse participant accepts that something is the case can always be understood along the lines of (12) in terms of whether she accepts a proposition. However, as we will discuss in §3 when we turn to epistemic modality in more detail, acceptance might not always be understandable in this way. Roughly, what we need to say instead is something more general: agent a accepts in w that such and such is the case if and only if her view $\Delta_a(w)$ in w instantiates a certain *property* $p \subseteq 2^{\mathcal{W}}$ which needn't be the property of set inclusion in some proposition (i.e., there needn't be a $P \subseteq \mathcal{W}$ such that $\Delta_a(w) \in p$ iff $\Delta_a(w) \subseteq P$). We can still stick with Stalnaker's latest characterization of common ground as common acceptance, but we must now say this: a property p of views (such as including a D -world) is common ground just in case everyone's view of the world instantiates p , everyone accepts that everyone has a view that instantiates p , and so on. Of course, much of the higher-level acceptance required for common ground status can and often will be tacit.

To accommodate this kind of account, we will need to revise the notions of common ground and context sets. The context set $CS(w)$ in w will now collect the view and attention state of each member of Agt in w , the view and attention state of each member of Agt at each world compatible with someone's view in w , and so on:

- (17) **Context sets** v. 2. Let $\mathcal{W}_w = \{w\} \cup \{v : \mathcal{R}^* wv\}$ designate the set containing w and any world reachable from w along $\mathcal{R}^*$. The *context set* in w is
 $CS(w) = \{ \langle \Delta_a(v), \Pi_a(v) \rangle : a \in \text{Agt}, v \in \mathcal{W}_w \}$.

Logic). The discourse updates in this paper will all leave $\mathcal{W}$ intact and alter only structures defined over this space (more closely resembling the dynamic systems in Gerbrandy & Groeneveld 1997; van Benthem & Liu 2007; Kooi & Renne 2011; van Benthem & Minică 2012; etc.).

A property p of views is common ground in w among the participants in Agt just in case their views all instantiate p , their views at worlds compatible with anyone's first-order view all instantiate p , and so forth. Equivalently, p is common ground just in case all of the views in $CS(w)$ instantiate it:

- (18) **Common ground** v. 2. Property $p \subseteq 2^W$ is *common ground* among Agt in w just in case for all $\langle \Delta_a(v), \Pi_a(v) \rangle \in CS(w)$, $p(\Delta_a(v))$.⁶

But why stop at common ground for information? Now that we are focusing on attention and have resolutions in the mix, we can also extend a kind of common ground status to subject matters. In fact, we can define a variety of collective (full) attention concepts by analogy with some standard kinds of group knowledge:⁷

- (19) **Mutual attention.** The members of Agt are *mutually attending* to M in w just in case for all $a \in Agt$, $\Pi_a(w)$ includes M .
- (20) **Distributed attention.** The members of Agt are *distributively attending* to M in w just in case $\Pi_a(w) + \Pi_b(w) + \dots$ includes M .
- (21) **Public attention.** The members of Agt are *publicly attending* to M in w just in case for each $\langle \Delta_a(v), \Pi_a(v) \rangle \in CS(w)$, $\Pi_a(v)$ includes M .⁸

That is, the participants in Agt are mutually attending to M iff everyone is sensitive to this issue, distributively attending to M iff their combined attention state is sensitive to this issue, and commonly attending to M iff they are mutually attending to M and this is commonly accepted.

These collective attention concepts might bring to mind the concept of a Question Under Discussion (QUD; Roberts 1996, 2012; Ginzburg 1996, 2012), a contextually salient discourse topic that interlocutors are collaboratively answering.⁹ But while we think that QUDs are intimately related to our group attention concepts, especially to public attention, there are good theoretical reasons to distinguish these concepts from one another. Subjects under public attention—let us call these ‘SUPs’ for

⁶ It is worth stressing that if every property expressible via assertion amounted to set inclusion in some proposition, then there would be little point in revising the original definitions: $CS_{v.1}(w) = \bigcup CS_{v.2}(w)$ so $CS_{v.1}(w) \subseteq P$ just in case $\Delta_a(v) \subseteq P$ for each $\Delta_a(v) \in CS_{v.2}(w)$. However, ‘unpacking’ the Stalnakerian model along the lines of (17) and (18) allows us to handle other kinds of properties.

⁷ See Fagin, Halpern, Moses & Vardi 1995 for a good survey of knowledge in multi-agent settings.

⁸ Ciardelli & Roelofson 2015 defines related notions of public *entertainment* and *wonderment*. Wondering about an issue (in this sense) involves having a preference to reach a state in which this issue is resolved, and entertaining is a technical notion that amounts to wondering or knowing. We think that attending to an issue is more basic; it needn't involve having a preference to resolve the issue but is rather a precondition for having such a preference. One must distinguish between the distinctions between worlds that an agent is sensitive to from those distinctions that she cares about.

⁹ For instance, Fritz & Lederman (2015) suggest that a QUD might be the join of the individual resolutions of the discourse participants.

short—might be regarded as the questions that could be asked given the distinctions that the interlocutors are aware of (in a strong iterated sense). But unlike mere SUPs (i.e., SUPs not under discussion), QUDs are typically taken to carry a public discourse goal of resolving or answering them (or otherwise *dispelling* them (Isaacs & Rawlins 2008); for e.g., by denying a presupposition or determining that they are practically unanswerable). Given their attendant commitments, QUDs constrain the flow of future discourse in a way that mere SUPs do not.

3 The dynamics of resistance

With the machinery for attention developed in previous sections, we are now in a position to turn to the account of resistance. The key idea is that epistemic resistance moves signal an unwillingness to immediately accept by drawing attention to some relevant issue. We formalize this using the notion of a *Table* from Farkas & Bruce 2010. The table is a holding-slot in context where assertions (and other moves) that are candidates for being accepted into the common ground (‘proposals’) are kept before they are accepted (or retracted). Resistance, on our proposal, amounts to a strategy for coordinating on whether a proposal should be accepted. The sketch is that the resisted move stays on the table while a related issue is discussed, and the agent who originally made the proposal has a chance to reaffirm or retract their proposal in light of new information. Relative to the present paper, the new information is always revealed by a shift in attention, though there may be other mechanisms for resistance. (This proposal for the dynamics of resistance, a generalization of Rawlins 2010, is independent of the attention mechanisms we have developed here, though for epistemic resistance moves we suggest that both pieces are necessary.)

To keep things focused, let us now narrow in on the following variant of Party:

- (22) A: Charlie is coming to the party.
 B: Diana might be coming.
 A: Good point. Then Charlie might not come.

Before formalizing the dynamics of resistance, we first informally walk through how we take this discourse to proceed. Suppose that (22) takes place in a background context where the following facts obtain: (i) Alice and Bob are inquiring into whether Charlie is coming; (ii) Alice accepts that Charlie is coming; (iii) Alice is not attending to whether Diana is coming; (iv) Alice knows that Charlie and Diana just had a messy breakup, but this is initially off her radar. If Alice weren’t overlooking Diana, then Alice would accept that Charlie might not come; (v) Bob is not sure whether or not Charlie is coming; according to what Bob accepts, Charlie might come but he might not. Bob is also not sure if Diana is coming (unlike Alice, he is

considering this issue); (vi) Bob is not sure whether or not (ii) holds; (vii) Bob is not sure whether or not (iii) holds.

First phase. Alice asserts that Charlie is coming. As a result, ‘Charlie is coming to the party’ is added to the assertion stack of the table. In making her assertion, Alice intends for Bob’s acceptance state to come to share the property, instantiated by her own acceptance state, of excluding the possibility that Charlie is not coming. The sentence on the stack encodes this coordination effort. After her assertion, Alice is publicly committed to Charlie coming; Bob accepts that Alice accepts this.

Second phase. Bob resists Alice’s claim by asserting that Diana might be coming, temporarily deferring acceptance or rejection of her claim while the context reaches an equilibrium with respect to attention states. This resistance move adds ‘Diana might be coming [to the party]’ to the assertion stack above Alice’s earlier utterance. Note that ‘Charlie is coming to the party’ must remain on the stack in order to make sense of alternative Party continuations like this:

- (23) A: Charlie is coming to the party.
 B: Diana might be coming.
 A: No, she isn’t.
 B: I see. Do you know what time he’ll show up?

In this exchange, Alice does not reassert that Charlie is coming after Bob’s resistance move. However, Bob’s final question presupposes that Charlie is coming, so Bob must accept that he is coming after Alice responds to the resistance.

How exactly should we understand Bob’s epistemic modal claim? We want to remain fairly noncommittal in this paper about how modal speech works. As a baseline account of epistemic *might*, though, we will tentatively adopt some ideas from [Veltman 1996](#) and [Yalcin 2008, 2011](#). Yalcin argues that for Bob to *accept* that Diana might be coming is not for him to be in a “second-order” state of mind where he accepts that it is compatible with some contextually salient information (his evidence, some group’s evidence, etc.) that Diana is coming. On Yalcin’s “first-order” model, Bob is rather in an acceptance state that leaves open the possibility that Diana will attend. Accepting that Diana might be coming is not a matter of accepting this or that proposition; it is rather a matter of being in a state of mind compatible with the “prejacent” proposition that Diana is coming.¹⁰ As Veltman

¹⁰ As Yalcin himself recognizes, compatibility with the proposition that Diana is coming is intuitively not enough to accept that Diana might be coming. More is needed. But what? According to Yalcin, Bob must also be sensitive to the question of whether Diana is coming (this amounts to being “appropriately receptive”, in some sense, to information which speaks to the question of whether Diana will come; it has nothing to do with attention) and Bob’s state of mind with respect to this question must leave open the possibility that Diana is coming. We are not sure that this is the right

would then put it, in asserting that Diana might be coming, Bob is inviting Alice to perform a “test” on her acceptance state to check whether it shares this property. This proposal is encoded by the sentence now lying at the top of the assertion stack.

In any case, for present purposes what are most interesting are the attentional effects of Bob’s modal claim. Before the resistance, Bob was neutral (given the scenario) as to whether Alice was attending to whether Diana is coming; even if Alice was attending to this issue, she and Bob were still not publicly attending to it. But after Bob resists, he accepts that Alice is attending to whether Diana is coming (and Alice accepts that Bob accepts that Alice is attending to this issue, and so forth). Furthermore, as a result of becoming aware of Diana, Alice changes her mind about Charlie. She now accepts that Diana might come and keep Charlie away.¹¹

Third phase. Rather than rejecting or counter-resisting Bob’s assertion, Alice agrees that Diana might come (this is signaled with the verbal marker ‘Good point’; in fact, this expression further signals that Alice thinks Bob’s modal assertion is *relevant* to the goals of their exchange). After this confirmation, ‘Diana might be coming’ is removed from the stack and the corresponding property of leaving open the possibility that Diana is coming becomes common ground.

Fourth phase. Alice now retracts her initial assertion and concedes that Charlie might not be coming to the party (there is no explicit retraction marker like ‘I take that back’; the retraction is rather signaled by the continuation that is incompatible with Alice’s earlier assertion). This retraction+concession sequence removes ‘Charlie is coming to the party’ from the assertion stack and replaces it with ‘Charlie might not come [to the party]’. At this point, Bob realizes that Alice has changed her mind about Charlie. Bob now accepts that Alice accepts that Charlie might not come.

Fifth phase. Bob silently accepts Alice’s claim. As a result, ‘Charlie might not come’ is removed from the stack and the property of leaving open the possibility that Charlie is not coming becomes common ground.

Things have gotten complex: unpacking our superficially straightforward Party

thing to say but we will bracket off this issue. Since we are concerned primarily with the attentional effects of epistemic possibility claims, we will just work with the simple compatibility account in the main text in place of a richer account.

11 Though we will not explicitly model this, it is worth noting that Bob’s resistance move has the further attentional effect that Alice comes to focus on the proposition that Diana is coming which is “visible” to Alice only after Bob’s modal claim. Alice takes Bob to be challenging her assertion only over the *D*-worlds where Diana attends the party. *Pace* Ciardelli et al. 2011; Ciardelli, Groenendijk & Roelofsen 2014; Roelofsen 2013 who present an *attentive semantics* for ‘might’ where the potential for an epistemic possibility claim to draw attention to the possibilities for its prejacent is hardwired into the semantics, we think that ‘might’ sentences can bring about attention to propositions through a pragmatic mechanism; the attention change potential for ‘might’ needn’t be built into its semantic entry. Unfortunately, we do not have space to go into details here.

example has revealed an amalgam of informative and attentive updates, discourse structuring mechanisms, and evolving higher-order attitudes. In the rest of this section, we develop the formal machinery underpinning this complexity.

3.1 Discourse models

To begin, consider a formal language $\mathcal{L}$ with a countable set $At_{\mathcal{L}}$ of sentence letters ‘A’, ‘B’, ‘C’, ..., negation ‘ $\neg$ ’ and conjunction ‘ $\wedge$ ’, and epistemic possibility ‘ $\Diamond$ ’. Assume that $\mathcal{L}$ has the usual grammar and let $S_{\mathcal{L}}$ designate the set of all well-formed sentences of $\mathcal{L}$. To interpret $\mathcal{L}$, we use a ‘textbook’ dynamic semantics.

- (24) **Discourse models.** A *model* $\mathcal{M} = \langle \mathcal{W}, \{\delta_a, \Pi_a : a \in \text{Agt}\}, \mathcal{V} \rangle$ consists of a frame $\mathcal{F}$ and an interpretation function $\mathcal{V} : At_{\mathcal{L}} \rightarrow 2^{\mathcal{W}}$ mapping each atom $\varphi_{At} \in At_{\mathcal{L}}$ to the set $\mathcal{V}(\varphi_{At})$ of φ_{At} -worlds in $\mathcal{W}$.
- (25) **Updates on information states.** The following recursive clauses extend $\mathcal{V}$ to the *update function* $+: S_{\mathcal{L}} \rightarrow (2^{\mathcal{W}})^{2^{\mathcal{W}}}$ (implicitly relativized to $\mathcal{M}$):

$$\begin{aligned} i + \varphi_{At} &= i \cap \mathcal{V}(\varphi_{At}) \\ i + \neg \varphi &= i \setminus i + \varphi && \text{Heim 1983} \\ i + \varphi \wedge \psi &= (i + \varphi) + \psi && \text{Heim 1982} \\ i + \Diamond \varphi &= \{w \in i : i + \varphi \neq \emptyset\} && \text{Veltman 1996} \end{aligned}$$

We also define updates over subject matters. These updates are quite simple and could in fact be stated in a more general way: each sentence letter φ_{At} appearing in the update rule $+\varphi$ introduces the polar issue of whether φ_{At} .

- (26) **Updates on subjects.** Letting $M_{\varphi_{At}} = \{\langle w, v \rangle : w \in \mathcal{V}(\varphi_{At}) \text{ iff } v \in \mathcal{V}(\varphi_{At})\}$, the *update function* $+: S_{\mathcal{L}} \rightarrow (M^*)^{M^*}$ is defined as follows (for a cleaner presentation, we use the same notation for both of our update functions):

$$\begin{aligned} M + \varphi_{At} &= M \cap M_{\varphi_{At}} & M + \varphi \wedge \psi &= (M + \varphi) + \psi \\ M + \neg \varphi &= M + \varphi & M + \Diamond \varphi &= M + \varphi \end{aligned} \quad (\text{cf. Yalcin 2011})$$

We can now define *incorporation* for both information states and subject matters. The idea is the same in both cases: an information state or subject matter incorporates φ (notation: $\triangleright \varphi$) just in case it is a fixed point under update with φ .

- (27) **Incorporation.** $i \triangleright \varphi$ iff $i + \varphi = i$
 $M \triangleright \varphi$ iff $M + \varphi = M$

Having defined incorporation for information states and subject matters, we can be more precise about *acceptance* than we were in §2:

- (28) **Acceptance.** An agent $a \in \text{Agt}$ in state $\langle \delta_a(w), \Pi_a(w) \rangle$ *accepts* in w that φ just in case $\Delta_a(w) \triangleright \varphi$. This agent *actively accepts* that φ if it is also the case that $\Pi_a(w) \triangleright \varphi$ and only *tacitly accepts* that φ otherwise.

3.2 Discourse contexts

We assume that discourse contexts involve three kinds of information: information about the subject matter of the discourse, information about the acceptance states of the individual participants in the discourse, and information about the discourse itself. We limit ourselves to two kinds of information of the latter sort: a (short) history of assertions made by discourse participants, and a history of questions that are currently under discussion (QUDs).

- (29) **Tables.** A *table* $T = \langle A, Q \rangle$ consists both of an *assertion stack* A of sentences in S_L and a *Question Under Discussion stack* Q of subject matters in M^* .¹²

Given a table, a model, and a designated actual world, we can construct what we will use as a discourse context:

- (30) **Discourse contexts.** A *context* $C = \langle \mathcal{W}^C, @^C, \{\delta_a^C, \Pi_a^C : a \in \text{Agt}\}, \mathcal{V}^C, T^C \rangle$ consists of a pointed model $\mathcal{M}^C, @^C$ with table T^C .

A discourse context is a time-slice of a conversation happening at $@^C$. We will subscript our contexts with (integer) indices to indicate an order for these slices. Since $\mathcal{W}$, $@$, and $\mathcal{V}$ will not change from context to context, we will omit the C -superscript altogether from these parameters going forward.

3.3 Discourse moves

With the context structure in place, we can now define various discourse moves. We earlier decomposed Party into five stages. Formally, these phases can be modeled with updates, summarized in (31). In the remainder of this section we unpack the effect of each of these moves, though we will have to leave the detailed formal version of the running example for a longer paper.

- (31) $C_5 = C_0 + [\text{Assert}_{\text{Alice}} C] + [\text{Assert}_{\text{Bob}} \Diamond D] + [\text{Confirm}_{\text{Alice}}] + [\text{Retract}_{\text{Alice}}] + [\text{Assert}_{\text{Alice}} \Diamond \neg C] + [\text{Confirm}_{\text{Bob}}]$

We begin with an assertion move. A felicitous assertion that ϕ must do (at least) three things: (i) it must put ϕ on the table, (ii) it must lead to it being publicly accepted that the author of the assertion accepts that ϕ is true, and (iii) it must ensure that the issues associated with ϕ come under public attention (this can potentially change agents' views). Effect (iii) is the novel component of our framework.

¹² We assume some familiarity with stacks here; see e.g. Kaufmann 2000; Isaacs & Rawlins 2008 for similar uses. We will use the following (standard) notation: $\text{push}(x, s)$ is the stack obtained by adding x to top of stack s , $\text{pop}(s)$ is the stack obtained by removing the top element of s , and $\text{top}(s)$ is the top element of s .

Let us focus on dialogues involving two participants $Agt = \{s, h\}$. Though we assume that these discourse participants utter only sentences in $\mathcal{L}$, it will be useful to introduce acceptance and attention operators that we as modelers can use to describe their mental states. Read ' $Acc_a^C(\varphi)$ ' as 'agent a actively accepts in C that φ is true', ' $Att_a^C(\varphi)$ ' as 'agent a 's attention state in C supports φ ', and ' $Con_a^C(\varphi)$ ' as 'agent a 's view in C can be consistently updated with φ '. For continuity, we interpret these operators in an update semantics:

(32) **Updates for acceptance and attention operators**

$$\begin{aligned} i + Acc_a^C(\varphi) &= \{w \in i : \Delta_a^C(w) \triangleright \varphi \text{ and } \Pi_a^C(w) \triangleright \varphi\} \\ i + Att_a^C(\varphi) &= \{w \in i : \Pi_a^C(w) \triangleright \varphi\} \\ i + Con_a^C(\varphi) &= \{w \in i : \Delta_a^C(w) + \varphi \neq \emptyset\} \\ M + Acc_a^C(\varphi) &= M \cap M_{Acc_a^C(\varphi)} \text{ (etc.)} \end{aligned}$$

where $M_{Acc_a^C(\varphi)} = \{\langle w, v \rangle : \{w\} \triangleright Acc_a^C(\varphi) \text{ iff } \{v\} \triangleright Acc_a^C(\varphi)\}$. It will also be useful to isolate an agent's acceptance states across worlds:

(33) **Agent states.** $\mathcal{F}_a^C = \{\delta_a^C, \Pi_a^C\}$ is the *agent state* for a in context C .

The set $\mathcal{F}_a^C$ encodes a 's contribution to a discourse context.

We are now ready to state our default assertion move (in what follows, we provide conditions only for components of a context that are affected by a move):

(34) **Assertive update.** $C_n + [Assert_s \varphi] = C_{n+1}$ s.t. $A^{C_{n+1}} = push(\varphi, A^{C_n})$

Felicity condition: $\{@\} \triangleright Acc_s^{C_n}(\varphi)$.

The only thing hardcoded into an assertion move is the update to the assertion stack. The felicity condition incorporates a variant of Searle's (1969) *sincerity* condition for assertions into our model of rational cooperative communication: a speaker can assert that φ is true only if she (actively) accepts this for the purposes of the conversation (of course, this might be violated by non-cooperative agents).¹³

We turn now to publication of acceptance and attention: (ii) and (iii) above. Rather than building these additional effects directly into our assertive update, we prefer to see them as arising from a couple of more general communicative principles. The first principle formalizes how the utterance of any declarative or interrogative sentence automatically draws attention to its associated issues before any further processing occurs, and it is commonly accepted that this happens.

(35) **Drawing attention.** If a speaker utters φ , then the agent state of any other discourse participant is immediately updated with $\oslash \varphi$

¹³ While sincerity will do here, many philosophers will want to strengthen the felicity condition along epistemic lines. Williamson (1996, 2000) famously argues that the standard for assertion is *knowledge*. For critical responses, see Weiner 2005; Lackey 2007; Hill & Schechter 2007; Kvanvig 2009.

The *attentive operator* $\odot$ appearing in this principle is defined as follows:

$$(36) \quad \mathcal{F}_a^C \odot \varphi = \begin{cases} [\delta_a^C \odot \varphi](w)(M) = \delta_a^C(w)(M) \\ [\Pi_a^C \odot \varphi](w) = \Pi_a^C(w) + \varphi \end{cases}.$$

The update (36) accomplishes the goal of bringing the issues associated with φ to a hearer's attention. Applying $\odot \varphi$ to h 's agent state $\mathcal{F}_h^{C_n}$ in C_n updates her attention state $\Pi_h^{C_n}(w)$ with φ at each $w \in \mathcal{W}$ and shifts her view to $\delta_h^{C_n}(w)(\Pi_h^{C_n}(w) + \varphi)$.¹⁴

The second communicative principle formalizes how hearers will typically infer that felicity conditions are satisfied, and it is commonly accepted that this happens (our formulation of this principle is crude; we leave its development for the future).

- (37) **Felicity inferences.** If a discourse move has a felicity condition of the form $\{@\} \triangleright \varphi$ where φ explicitly concerns the state of the acting participant and another discourse participant does not have good grounds for believing that this felicity condition is violated, suspended, or altered in the context of use, then this other participant's agent state is updated with $\oplus \varphi$

where the *informative operator* $\oplus$ appearing in this principle is defined as follows:

$$(38) \quad \mathcal{F}_a^C \oplus \varphi = \begin{cases} [\delta_a^C \oplus \varphi](w)(M) = \begin{cases} \delta_a^C(w)(M) + \varphi & \text{if } M \triangleright \varphi \\ \delta_a^C(w)(M) & \text{otherwise} \end{cases} \\ [\Pi_a^C \oplus \varphi](w) = \Pi_a^C(w) + \varphi \end{cases}$$

Of course, it would be possible to hardwire the result of (38) directly into our assertive update, but we take the fact that felicity conditions can be violated or suspended as evidence against such a maneuver.

The informative operator $\oplus$ is a bit more complicated. Unlike with the attentive update $\odot \varphi$ where the change in an agent's attention state is primary and any change in her view is derivative, the aim of $\oplus \varphi$ is to update an agent's view with φ . This being so, $\oplus \varphi$ must also adjust the agent's attention state to ensure that updating her view returns a visible set of possible worlds. Although $\delta_a^C(w)(M) + \varphi$ needn't be visible at M when $M \not\triangleright \varphi$, it can be shown that $\delta_a^C(w)(M) + \varphi$ will be visible

¹⁴ To ensure that $\odot \varphi$ always returns a genuine agent state satisfying Introspection, we must further constrain δ_a and Π_a somewhat. The following strengthening of Introspection would do the trick:

- (i) **Strengthened Introspection (SI).** For all $w, v \in \mathcal{W}$, if $M \in M^*$ is a refinement of $\Pi_a^C(w)$ and $v \in \delta_a^C(w)(M)$, then $\delta_a^C(v)(M) = \delta_a^C(w)(M)$ and $\Pi_a^C(v) = \Pi_a^C(w)$.

Since each subject matter trivially refines itself, SI entails Introspection (and if one is already willing to buy into Introspection, it is a short step to buy into SI). Moreover, it is easy to prove that $\odot \varphi$ preserves SI (and so Introspection as well). We leave this as an exercise for the reader.

at M so long as $M \triangleright \varphi$. Moreover, we take an informative update $\oplus\varphi$ to percolate through an agent's acceptance state so that this update doesn't just disappear after subsequent shifts in her attention. When updating $\delta_a^C(w)(M)$ with φ is guaranteed to return a visible information state (i.e., when $M \triangleright \varphi$), we update it.^{15,16}

Following the felicity inference $\oplus Acc_s^{C_n}(\varphi)$ in particular, h 's attention state and view at each $w \in \mathcal{W}$ shift to $\Pi_h^{C_n}(w) + Acc_s^{C_n}(\varphi)$ and $\delta_h^{C_n}(w)(\Pi_h^{C_n}(w) + Acc_s^{C_n}(\varphi)) + Acc_s^{C_n}(\varphi)$ respectively. It becomes commonly accepted that h accepts that s actively accepts that φ is true.

We will henceforth use $C_n + [\text{Assert}_s \varphi]$ to designate the context C_{n+1} obtained from C_n by first applying Drawing Attention, then Assertive Update, and then Felicity Inferences based on the felicity condition $\{@\} \triangleright Acc_s^{C_n}(\varphi)$; we bracket off the possibility that h has good grounds for believing that this condition is violated, suspended, or altered. Assuming in addition that the felicity condition is actually met, it is straightforward to prove that in C_{n+1} , (a) for each sentence letter φ_{At} appearing in φ , the corresponding subject matter $M_{\varphi_{At}}$ is an SUP among $\{s, h\}$ in $@$, and (b) the property $\triangleright Acc_s^{C_{n+1}}(\varphi)$ of views is common ground among $\{s, h\}$ in $@$. The combination of assertive update, drawing attention, and felicity inferences thus accomplishes all three of our stated goals (i)-(iii).

Turning now to hearer response moves, let us consider what can happen following an assertion. The most straightforward reaction is when a hearer *confirms* the assertion (this is the default or unmarked response since an assertion “projects” its confirmation (Farkas & Bruce 2010)). If h confirms s 's assertion, then h publicly accepts that the top element of the stack is true, and this sentence is removed from the stack. Confirmation is defined only if the assertion stack is nonempty and has a felicity condition that this move will not lead to inconsistency.

(39) **Confirmation**

$$C_n + [\text{Confirm}_h] = C_{n+1} \text{ s.t. } \begin{array}{ll} \text{(i)} & A^{C_{n+1}} = \text{pop}(A^{C_n}) \\ \text{(ii)} & \mathcal{F}_h^{C_{n+1}} = \mathcal{F}_h^{C_n} \oplus \text{top}(A^{C_n}) \end{array}$$

Felicity condition: $\{@\} \triangleright \text{Con}_h^{C_n}(\text{top}(A^{C_n}))$. Defined in C_n only if $A^{C_n} \neq \langle \rangle$.

Again, we overload ‘+’ and henceforth use $C_n + [\text{Confirm}_h]$ to designate the context C_{n+1} obtained from C_n by applying Confirmation followed by Felicity Inferences. It is straightforward to prove that in C_{n+1} , the property $\triangleright Acc_h^{C_{n+1}}(\text{top}(A^{C_n}))$ of views

¹⁵ It is also straightforward to prove that $\oplus\varphi$ preserves Strengthened Introspection.

¹⁶ While our informative update works well in many cases, it should be regarded as a first pass. An alternative approach would be to have $\oplus\varphi$ update only the agent's view at $\Pi_a^C(w) + \varphi$ and then have the agent reprocess or reconsider earlier discourse moves whenever her attention shifts (Franke & de Jager 2010 floats a similar proposal). However, this makes things considerably more complicated; for one thing, it would require discourse participants to carry around a conversational record.

is common ground among $\{s, h\}$ in $@$. It is also straightforward to show that if s felicitously asserts that ϕ is true, this is subsequently confirmed by h , and it remains commonly accepted that s accepts that ϕ is true after this confirmation, then the property $\triangleright \phi$ of views acquires common ground status among $\{s, h\}$ in $@$.

At the other end of the spectrum of reactions that pop the assertion stack, we have the non-canonical discourse move of *retraction*.

(40) **Retraction.** $C_n + [\text{Retract}_s] = C_{n+1}$ s.t. $A^{C_{n+1}} = \text{pop}(A^{C_n})$

Defined in C_n only if $A^{C_n} \neq \langle \rangle$.¹⁷

Felicity condition: $\{@\} \triangleright \neg \text{Acc}_s^{C_n}(\text{top}(A^{C_n}))$.

Under the current setup, one's view must shift before retraction; if it does not, then the felicity condition for retraction will not be met. Suppose that s felicitously asserts that ϕ is true and later felicitously retracts this assertion. By the felicity condition for assertion, s must have actively accepted that ϕ is true at the time this sentence was used. But her retraction signals a shift in her view.

Besides confirmation and retraction, of course, there are a variety of other response options. A hearer might *reject* a prior move by making an assertion that updates the stack in such a way that it can no longer be emptied only by confirmation moves (arguably, this is just what it is to *disagree*—to be in a conversational context where the assertion stack can be fully popped only after at least one of the participants retracts an earlier public commitment or the participants agree to disagree (cf. [Farkas & Bruce 2010](#))). As we have been stressing all along, a hearer might alternatively just *resist* a prior move by making an assertion or asking a question with the aim of generating further discussion about a relevant subordinate issue.

4 Conclusions

In this paper we introduced *resistance moves*, a new type of move in the space between agreement and disagreement. We homed in on a particular kind of resistance move signaled by an epistemic possibility modal, and argued that this species of resistance targets *attention*: $\Diamond p$ draws attention to the issue of whether or not p holds. If an agent was not previously attending to this issue, there may be propositions that are not ‘visible’ that would cause them to retract a claim, and drawing attention to such an issue can trigger retraction. Our proposal for the dynamics of resistance is that it involves coordinating on subsidiary moves while the original (resisted) claim is still in limbo; once this sub-coordination is complete, agents again consider whether to incorporate the original claim into the common ground.

¹⁷ Since one can retract only what one has placed on the assertion stack, this move should be defined only if $\text{top}(A^{C_n})$ was previously uttered by s . Articulating this condition in our formal model would require us to keep track of who adds particular items to the table (cf. [Gunlogson 2008](#)).

References

- van Benthem, Johan & Fenrong Liu. 2007. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics* 17(2). 157–182.
- van Benthem, Johan & Stefan Minică. 2012. Toward a dynamic logic of questions. *Journal of Philosophical Logic* 41(4). 633–669.
- Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2011. Attention! *Might* in inquisitive semantics. In Ed Cormany, Satoshi Ito & David Lutz (eds.), *Semantics and Linguistic Theory (SALT) 19*, 91–108. eLanguage.
- Ciardelli, Ivano, Jeroen Groenendijk & Floris Roelofsen. 2014. Information, issues, and attention. In Daniel Gutzmann, Jan Köpping & Cécile Meier (eds.), *Approaches to Meaning: Composition, Values, and Interpretation*, 128–166. Brill.
- Ciardelli, Ivano & Floris Roelofsen. 2015. Inquisitive dynamic epistemic logic. *Synthese* 192(6). 1643–1687.
- Fagin, Ronald, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi. 1995. *Reasoning about Knowledge*. Cambridge: MIT Press.
- Farkas, Donka & Kim Bruce. 2010. On reacting to assertions and polar questions. *Journal of Semantics* 27(1). 81–118.
- Franke, Michael & Tikitú de Jager. 2010. Now that you mention it: awareness dynamics in discourse and decisions. In Anton Benz, Christian Ebert, Gerhard Jäger & Robert van Rooij (eds.), *Language, Games, and Evolution*, 60–91. Springer.
- Fritz, Peter & Harvey Lederman. 2015. Standard state space models of unawareness. In R. Ramanujam (ed.), *Proceedings of TARK XV: Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, 163–172.
- Gerbrandy, Jelle & Willem Groeneveld. 1997. Reasoning about information change. *Journal of Logic, Language, and Information* 6(2). 147–169.
- Ginzburg, Jonathan. 1996. Dynamics and the semantics of dialogue. In Jerry Seligman & Dag Westerståhl (eds.), *Logic, Language, and Computation 1 (CSLI Lecture Notes 58)*, 221–237. CSLI Publications.
- Ginzburg, Jonathan. 1998. Clarifying utterances. In Joris Hulstijn & Anton Nijholt (eds.), *Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogue*, 11–30.
- Ginzburg, Jonathan. 2012. *The Interactive Stance: Meaning for Conversation*. Oxford: Oxford University Press.
- Groenendijk, Jeroen. 1999. The logic of interrogation. In Tanya Matthews & Devon Strolovitch (eds.), *Semantics and Linguistic Theory (SALT) 9*, 109–126.
- Groenendijk, Jeroen & Martin Stokhof. 1984. *Studies on the semantics of questions and the pragmatics of answers*: University of Amsterdam PhD dissertation.

- Gunlogson, Christine. 2008. A question of commitment. *Belgian Journal of Linguistics* 22. 101–136.
- Heim, Irene. 1982. *The semantics of definite and indefinite noun phrases*: University of Massachusetts, Amherst PhD dissertation.
- Heim, Irene. 1983. On the projection problem for presuppositions. In Michael Barlow, Daniel Flickinger & Michael Wescoat (eds.), *West Coast Conference on Formal Linguistics (WCCFL)* 2, 114–125. Stanford University.
- Hill, Christopher S. & Joshua Schechter. 2007. Hawthorne’s lottery puzzle and the nature of belief. *Philosophical Issues* 17(1). 102–122.
- Hulstijn, Joris. 2000. *Dialogue models for inquiry and transaction*: University of Twente PhD dissertation.
- Isaacs, James & Kyle Rawlins. 2008. Conditional questions. *Journal of Semantics* 25(3). 269–319.
- de Jager, Tikitu. 2009. *Awareness, attention, assumption*: University of Amsterdam PhD dissertation.
- Kaufmann, Stefan. 2000. Dynamic context management. In Martina Faller, Stefan Kaufmann & Marc Pauly (eds.), *Formalizing the Dynamics of Information*, 171–188. Stanford, CA: CSLI.
- Khoo, Justin. 2015. Modal disagreements. *Inquiry* 58(5). 511–534.
- Klecha, Peter. 2014. *Bridging the divide: scalarity and modality*: University of Chicago PhD dissertation.
- Kooi, Barteld & Bryan Renne. 2011. Arrow update logic. *The Review of Symbolic Logic* 4(4). 536–559.
- Kvanvig, Jonathan. 2009. Assertion, knowledge, and lotteries. In Duncan Pritchard & Patrick Greenough (eds.), *Williamson on Knowledge*, 140–160. Oxford University Press.
- Lackey, Jennifer. 2007. Norms of assertion. *Noûs* 41(4). 594–626.
- Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(3). 339–359.
- Lewis, David. 1988. Statements partly about observation. *Philosophical Papers* 17(1). 1–31.
- McCarthy, John. 1980. Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence* 13. 27–39.
- Plaza, Jan A. 1989. Logics of public communication. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic & Z.W. Ras (eds.), *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, 201–216. Oak Ridge National Laboratory.
- Rawlins, Kyle. 2010. Conversational backoff. In Nan Li & David Lutz (eds.), *Semantics and Linguistic Theory (SALT)* 20, 347–365.
- Roberts, Craige. 1996. Information structure in discourse: towards an integrated for-

- mal theory of pragmatics. In Jae Hak Yoon & Andreas Kathol (eds.), *OSUWPL: Papers in Semantics* 49, Ohio State University.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5(6). 1–69.
- Roelofsen, Floris. 2011. Information and attention. Unpublished manuscript, ILLC.
- Roelofsen, Floris. 2013. A bare bone attentive semantics for *might*. In Maria Aloni, Michael Franke & Floris Roelofsen (eds.), *The Dynamic, Inquisitive, and Visionary Life of ϕ , $?\phi$, and $\Diamond\phi$: a Festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman*, 190–215. ILLC.
- Searle, John. 1969. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Syntax and Semantics 9: Pragmatics*, 315–332. Academic Press.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge: MIT Press.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25. 701–721.
- Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press.
- Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25(3). 221–261.
- Weiner, Matthew. 2005. Must we know what we say? *Philosophical Review* 114(2). 227–251.
- Williamson, Timothy. 1996. Knowing and asserting. *Philosophical Review* 105(4). 489–523.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Yablo, Stephen. 2014. *Aboutness*. Princeton: Princeton University Press.
- Yalcin, Seth. 2008. *Modality and inquiry*: Massachusetts Institute of Technology Ph.D. dissertation.
- Yalcin, Seth. 2011. Nonfactualism about epistemic modality. In Andy Egan & Brian Weatherson (eds.), *Epistemic Modality*, 295–332. Oxford University Press.

Justin Bledin
Department of Philosophy
281 Gilman Hall
3400 N. Charles Street
Baltimore, MD 21218
jbledin@jhu.edu

Kyle Rawlins
Department of Cognitive Science
237 Krieger Hall
3400 N. Charles Street
Baltimore, MD 21218
kgr@jhu.edu