



Million dollar questions: why deliberation is more than information pooling

Daniel Hoek¹ · Richard Bradley²

Received: 14 November 2020 / Accepted: 31 January 2022
© The Author(s) 2022

Abstract

Models of collective deliberation often assume that the chief aim of a deliberative exchange is the sharing of information. In this paper, we argue that an equally important role of deliberation is to draw participants' attention to pertinent questions, which can aid the assembly and processing of distributed information by drawing deliberators' attention to new issues. The assumption of logical omniscience renders classical models of agents' informational states unsuitable for modeling this role of deliberation. Building on recent insights from psychology, linguistics and philosophy about the role of questions in speech and thought, we propose a different model in which beliefs are treated as answers directed at specific questions. Here, questions are formally represented as partitions of the space of possibilities and individuals' information states as sets of and corresponding partial answers to them. The state of conversation is then characterised by individuals' information together with the questions under discussion, which can be steered by various deliberative inputs. Using this model, deliberation is then shown to shape collective decisions in ways that classical models cannot capture, allowing for novel explanations of how group consensus is achieved.

1 Introduction

Deliberation is a process by which agents, individually and/or collectively, use reasoning to reach opinions or judgements about what to believe, value and do. One way in which deliberation can serve this end is by bringing information to bear on issues. But if the communication of information were the sole aim of deliberation,

✉ Richard Bradley
r.bradley@lse.ac.uk

Daniel Hoek
dhoek@vt.edu

¹ Department of Philosophy, Virginia Tech, Blacksburg, VA, USA

² Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, London, UK

then one could only make a productive contribution to a deliberative exchange by offering a new piece of pertinent information. But in fact information is not the sole currency of deliberation, and fruitful contributions can be made to it even from a position of ignorance. Consider, for instance, the following exchange:

Example 1. Vacationer. Jane is about to leave for the airport to go on a far-away vacation with her roommate Phil. Unsure whether she has everything she needs, she asks Phil (who has not seen her pack her suitcase). They have the following conversation:

J: I'm not sure I have everything.

P: Well, did you bring your toothpaste?

J: Yup, it's right here.

P: Passport?

J: It's in the front pocket.

P: What about your phone?

J: Oh darn, I left it charging upstairs. Let me run up and get it.

In this exchange, Phil contributes no new information. He does not even have any information about the contents of Jane's luggage. His contribution to the deliberation consists of posing some good *questions*. These questions focus Jane's attention on various relevant pieces of information that she already possessed, and in this way, they figure out what needs to be done.

The growing formal literature on deliberation, spanning (primarily) economics, political science and philosophy, presents models of deliberative exchanges that differ along a number of dimensions: the types of opinion states they work with (beliefs, preferences, probabilities), the kinds of speech acts involved (assertions, judgements, arguments), the mechanisms for opinion revision they postulate (information updates, preference revisions, conditionalization), and the role that they assign to strategic incentives.¹ Many of these models assume that the aim of group deliberation is the pooling or aggregation of information that is distributed amongst different agents and that "the transmission, processing, and aggregation of the information that forms the basis of individual and collective decision making is the engine that drives the deliberative wheel" (Landa and Meirowitz 2009, p. 427; see also Stalnaker 1974; Heim 1988). Consequently, much of the debate around the value of deliberation has centred around evidence that shows the scope and limits of our ability to effectively exchange information, and of our ability to build consensus through information exchange (Sunstein 2006; Austen-Smith and Feddersen 2006; Lehrer and Wagner 1981).

In this paper, we challenge the notion that deliberation revolves around information exchange, by underlining a different way in which deliberation can shift opinion and forge consensus. In doing so, we build on recent insights about the role of *questions* in speech and thought. Those insights stem from linguistics (Roberts 2012; Ciardelli et al. 2019; Bledin and Rawlins 2020), psychology (Koralus and

¹ See List (2017) for a survey.

Mascarenhas 2013; Koralus 2014; Carruthers 2018) and philosophy (Friedman 2017; Yalcin 2018; Hoek forthcoming). Our model emphasises the central role that questions play in deliberative exchanges and in the way information is stored. It suggests that information distribution is not the sole or even the main aim of deliberation; an equally important role is to aid in processing and organising the information we already possess by drawing our attention to new issues. In this way, deliberation can help us see the relevance of old information to new predicaments. The *Vacationer* example is a case in point: Jane already has the information that her phone is upstairs. But she fails to attend to it and does not realise its relevance, until Phil asks her about it.

Besides directing our attention, questions have the power to bring new deductive consequences of our existing beliefs to light. They can do so by bringing disparate pieces of relevant information together, and also by shaping our conception of the possibilities in play. Here is one illustration:

Example 2. Dinner Danger. Alice and Bob are getting ready for a dinner party—they know Claire, Darian, Evelyn and Fiona will all be coming. Earlier in the day, Alice spoke to Gerry who said he would come as well. Alice and Bob both know that, unbeknownst to their host, Gerry has an acrimonious ongoing dispute with Darian, and the party will be a bust if both of them come. However, unlike Bob, Alice has not considered this possibility.

B: I'm not sure we should go.

A: Why not? It will probably be a lot of fun!

B: Darian is going to be there. What if Gerry comes too?

A: Oh you're totally right! Actually Gerry told me earlier today he is coming. Those two are just going to spend all night quarrelling. All right, we're staying home—I'll come up with an excuse.

Again, Bob does not bring any new information to the table. Prior to the conversation, Alice already has the following three pieces of information:

- (i) Gerry will be at the party.
- (ii) Darian will be at the party.
- (iii) The party will not be fun if Gerry and Darian are both there.

But she has not yet put two and two together: she has not linked the new information that Gerry will come to her prior knowledge of (ii) and (iii). Consequently she has not yet formed the belief that

- (iv) The party will not be fun.

She only arrives at that conclusion when Bob raises the issue of Darian and Gerry both coming. This at once forces Alice to bring the relevant beliefs together while considering their bearing on the question at hand: whether the party will be fun.

As we will explain in Sect. 2 below, classical models of belief make it difficult to capture this phenomenon because they operate under the assumption of *logical*

omniscience, which says that agents believe all the logical consequences of what they believe. This property makes these models unsuitable for capturing Jane's prior information state in *Vacationer* and Alice's prior information state in *Dinner Danger*. For instance, no probability function captures Alice's beliefs in *Dinner Danger* prior to her exchange with Bob, because no probability function that assigns a suitably high probability to each of (i-iii) simultaneously assigns a suitably low probability to (iv). Likewise, no probability function can capture Jane's beliefs at the start of *Vacationer*. For any probability function that assigns a middling probability to the proposition that *Jane has everything she needs* fails to assign a high probability to the proposition that *Jane does not have her phone (which she needs)*. This assumption of logical omniscience also renders classical models unsuitable for modelling the role that deliberation plays in decision making (Bradley 2009).

As shown in Sect. 3, our model avoids these issues by embracing the idea that beliefs should be viewed as answers directed at specific questions, where questions are formally modelled as partitions of the space of possibilities. We further hold that agents can possess such answers without necessarily bringing them to bear on every other question on which the agent has a view (even when the information is relevant). It is here that a new role for deliberation emerges: namely, the role of bringing old answers to bear on new questions, thereby allowing for new inferences.

Section 4 applies our model to *Vacationer* and *Dinner Danger*. Section 5 concludes by drawing out some of the implications of our treatment of deliberation for the understanding of collective attitude formation.

2 Classical models of deliberation

In this section, we present a standard model of conversation as information exchange, along the lines of Stalnaker 1974 and Heim 1988. Following Hintikka (1962), those models represent agents' information states as subsets of the event space Ω , rather than a probability function over Ω :

A (consistent) *classical information state* is a nonempty subset \mathbf{B} of the event space Ω . The corresponding proposition set $|\mathbf{B}|$ is the set of all propositions that hold true at every possible world in \mathbf{B} . A classical information state \mathbf{B} is *smaller than* another state \mathbf{B}' just in case $|\mathbf{B}| \subseteq |\mathbf{B}'|$.

This model of doxastic states can be understood as a simplification of the probabilistic representation: the points composing an agent's information state can be thought of as representing the possibilities to which they ascribe a non-zero probability. This simplification does not matter for our purposes, in that closely analogous problems arise for probabilistic models.

A classical agent with belief state \mathbf{B} believes all and only the propositions in $|\mathbf{B}|$. Here *propositions* can themselves be understood as subsets of the event space. On that understanding, $|\mathbf{B}| = \{p: \mathbf{B} \subseteq p \subseteq \Omega\}$. So on the classical picture, agents' beliefs are closed under multi-premise entailment. That is, if $p, q \in |\mathbf{B}|$ jointly entail r , then $r \in |\mathbf{B}|$. Note that, while stronger information states contain fewer worlds, they are *bigger* in the sense of encompassing more propositions. Note also that our definition of

an information state excludes the degenerate case of an inconsistent information state, which does not represent a possible doxastic state.

In a conversation, some information is shared between all participants, while other information is only had by some. We may represent this situation as follows:

A (*classical*) *conversational state* with interlocutors $1, 2, \dots, n$ is a tuple of classical information states $\langle \mathbf{C}, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_n \rangle$ where \mathbf{C} represents the *common ground* between the interlocutors and \mathbf{B}_i the belief state of interlocutor i . It is assumed that every proposition that is common ground is believed by each interlocutor—that is, $|\mathbf{C}| \subseteq |\mathbf{B}_i|$ for all i .

A participant i is *in a position to assert* a proposition p when $p \in |\mathbf{B}_i|$ but $p \notin |\mathbf{C}|$. It is possible that a proposition is believed by all participants and yet fails to be common ground. For instance, if we both know the time but I don't know that you know it, it still makes sense for me to assert that "It's two o'clock," to ensure we are on the same page. This is why it is not in general true that $|\mathbf{C}| = |\mathbf{B}_1| \cap \dots \cap |\mathbf{B}_n|$.

When an interlocutor asserts that p , they make public a piece of information that was not yet common ground. This effect is captured with the notion of a conversational update:

An **update** of an information state \mathbf{B} by a proposition p , written $\mathbf{B} + p$, is defined as the smallest classical belief state whose belief set includes $|\mathbf{B}| \cup \{ p \}$. A **conversational update** by the proposition p is defined thus:

$$\langle \mathbf{C}, \mathbf{B}_1, \dots, \mathbf{B}_n \rangle + p = \langle \mathbf{C} + p, \mathbf{B}_1 + p, \dots, \mathbf{B}_n + p \rangle$$

The updated information state $\mathbf{B} + p$ is only defined when p is consistent with \mathbf{B} . Whenever it is defined, $\mathbf{B} + p = \mathbf{B} \cap p$. A conversational update is defined just in case p is consistent with all information states in the prior state.

As a conversation progresses, more and more information comes to be shared, as more updates occur. When the belief states of the participants are consistent, this process eventually leads to an equilibrium in which all belief states are identical to the common ground, and no participant is in a position to contribute any further information.

For our purposes here, the key limitation of the classical model is that interlocutors can only contribute to the conversation when they possess information that is not yet common ground, and can only learn from those who possess information that they themselves do not yet have. However, in *Vacationer* we saw how Jane learnt something from Phil's questions, though Phil had no relevant information. Likewise in *Dinner Danger*, Bob's remarks are essential to the deliberation although he has less information than Alice has. The classical model cannot make sense of such deliberative exchanges. To capture them, we must enrich the model by incorporating the role of questions in deliberation and in cognition.

3 An inquisitive model of deliberation

The inquisitive model we propose enriches conversational states with a **question under discussion** or **QUD**. This question serves to pick out the body of shared information on which the interlocutors are presently focusing their joint attention (Koralus 2014) while at the same time representing the issues they aim to resolve (Roberts 2012; Ginzburg 2012; Ciardelli et al. 2019). Besides the addition of a QUD, our model also replaces classical information states with *inquisitive information states*, which represent a body of answers directed at specific questions (as in Yalcin 2011, 2018; Fritz and Lederman 2015; Bledin and Rawlins 2020; Hoek forthcoming). Since agents need not have views on every question, an inquisitive information state \mathbf{B} is associated with a domain $\mathcal{D}_{\mathbf{B}}$, which contains all, and only, the questions it addresses. This yields the following characterization of a conversational state, which will be unpacked in what follows:

An (*inquisitive*) *conversational state* for interlocutors 1, 2, ..., n is a tuple $\langle \mathbf{D}, \mathbf{C}, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_n \rangle$ where \mathbf{D} represents the *question under discussion*, and $\mathbf{C}, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_n$ are inquisitive information states representing the common ground and the belief states of the participants. It is assumed that $\mathbf{D} \in \mathcal{D}_{\mathbf{C}}$, and that $|\mathbf{C}| \subseteq |\mathbf{B}_i|$ for each interlocutor i .

The first concept that needs to be formalised here is that of a *question*. Formally speaking, a question is a partition of the event space (Groenendijk and Stokhof 1984). The basic idea here is to characterise a question in terms of the information that is required to answer it exhaustively. For instance, the question *How many marbles are in this jar?* partitions the event space according to the number of marbles in the jar at each world. Polar questions, like *Will it rain this afternoon?*, partition the space into just two cells. Thus a question is much like a random variable, where each cell of the partition is the pre-image of a value in the domain of the corresponding random variable.

A *question* Q is a set of non-empty, disjoint subsets of the event space Ω which jointly cover the entire space. Every non-empty set of Q -cells is a (*partial*) *answer* to Q , and every singleton set is a *complete answer*. A question Q *contains* the question R as a part just in case Q is at least as fine-grained as R , in the sense that every R -cell is a union of Q -cells. The *conjunction* of two questions Q and R , written QR or $Q \wedge R$, is the coarsest common refinement of Q and R .²

Here are some examples to illustrate this definition. *Jane is over 21* is a partial answer to the question *How old is Jane?*, while *Jane is 26* is a complete answer. The polar question *Is Jane over 21 or not?* is part of the question *How old is Jane?*. The question *What's today's date?* is the conjunction of the question *What month is it?* and *What day of the month is it?*.

² The *coarsest common refinement* of two partitions Q and R is the partition $\{q \cap r : q \in Q, r \in R\} - \{\emptyset\}$.

In our model, belief contents are answers that are directed at specific questions. We formalise this as an ordered pair of a question and an answer in the sense just defined:

A **question-directed proposition**, or **quizposition** for short, is an ordered pair $\langle Q, A \rangle$, written A^Q , where Q is a question, and $A \subseteq Q$ is either a (partial) answer to Q or the empty set. A quizposition A^Q **entails** B^R just in case $\cup A \subseteq \cup B$. A quizposition A^Q **contains** B^R as a part just in case Q contains R and $\cup A \subseteq \cup B$. The **negation** $\neg A^Q$ of a quizposition A^Q is the quizposition $\langle Q, Q \setminus A \rangle$. The **conjunction** $A^Q \wedge B^R$ of two quizpositions A^Q and B^R is the quizposition $\langle QR, AB \rangle$, where AB is the weakest QR -answer that entails both A and B . (In other words, $A^Q \wedge B^R$ is the smallest quizposition containing both A^Q and B^R .)

Intuitively, quizposition containment, also known as analytic entailment, is a particularly direct form of entailment (Gemes 1994). For instance, the quizposition *Jill lives on Broad Street*, in answer to the question *Which street does Jill live on?*, is part of the quizposition that *Jill lives on 34 Broad Street*, in answer to *What is Jill's address?* The reason is that the latter quizposition answers a bigger, more comprehensive question than the former. By contrast, the quizposition *Either Jack or Jill lives on Broad Street*, in answer to *Who lives on Broad Street?*, is not part of *Jill lives on 34 Broad Street*. For while the latter claim entails the former, they concern distinct questions. On the model we propose, the problematic classical assumption of closure under entailment is replaced by the weaker assumption of closure under parthood. So while we may not believe every entailment of what we believe, we do believe every *part* of what we believe.

Closure under parthood is one aspect of the particular account of belief states we adopt here, which has been defended on independent, decision-theoretic grounds by Hoek (2019). According to this account, an agent's beliefs form a network of views on different questions, where their view on any given big question is composed of their views on its component questions. For instance, your view about *What date it is* incorporates your view on *What month it is* and your view on *What day of the month it is* as parts. Thus, agents' views on bigger questions harmonize with their views on its component questions, and their views on overlapping questions must harmonize on the overlapping part. In this way, an agent's beliefs are connected through their thematic links in a web-like structure. Crucially however, these links between agent's views are weak enough to allow for the possibility that agents' beliefs fail to be closed under entailment, and also for the possibility of inconsistent beliefs:

An **inquisitive domain** \mathcal{D} is a set of partition questions that is closed under parthood (coarsening). A (coherent) **inquisitive information state** with domain $\mathcal{D}_{\mathbf{B}}$ is a function $\mathbf{B}: \mathcal{D}_{\mathbf{B}} \rightarrow \mathcal{P}(\Omega)$ such that:

1. *Answerhood*. For all $Q \in \mathcal{D}_{\mathbf{B}}$, $\mathbf{B}(Q) = \cup A$ for some non-empty $A \subseteq Q$.
2. *Harmony*. For all $Q \in \mathcal{D}_{\mathbf{B}}$, if R is part of Q , then $R \in \mathcal{D}_{\mathbf{B}}$ and $\mathbf{B}(R) = \cup \{ r \in R : r \text{ is consistent with } \mathbf{B}(Q) \}$

The corresponding quizposition set is definable as follows:

$$|\mathbf{B}| := \{A^Q : Q \in \mathcal{D}_{\mathbf{B}} \text{ and } \mathbf{B}(Q) \subseteq \cup A\}$$

An inquisitive information state \mathbf{B} is *inconsistent* if there is no possible world in Ω at which every quizposition in $|\mathbf{B}|$ is true. An inquisitive information state \mathbf{B} is *smaller than* state \mathbf{B}' just in case $|\mathbf{B}| \subseteq |\mathbf{B}'|$.

Note that, although belief states can be inconsistent, an agent's view on any particular question is always consistent. An inquisitive belief state is fully characterized by the corresponding quizposition set.³ In particular $Q \in \mathcal{D}_{\mathbf{B}}$ if and only if $Q^Q \in |\mathbf{B}|$. An agent in belief state \mathbf{B} is said to *believe* the quizposition A^Q if and only if $A^Q \in |\mathbf{B}|$. As in the classical case, inquisitive belief states can be viewed as simplified versions of a probabilistic description of the agent's doxastic state.⁴

In the inquisitive model, participants can contribute by posing new questions or by contributing information:

In the conversational state $\langle D, C, \mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_n \rangle$, interlocutor i is *in a position to ask Q* just in case $Q \in |\mathbf{B}_i|$ and Q is not part of D . They are *in a position to answer A^Q* just in case Q is part of D , $A^Q \in |\mathbf{B}_i|$ and $C(D) \not\subseteq \cup A$.

The characteristic effect of posing a question is to draw participants' attention to that question, thereby refining the QUD (Roberts 2012; Ciardelli et al. 2019). This can be captured with the notion of a question update:

A *question update* of the information state \mathbf{B} by the question Q , written $\mathbf{B} + Q$, is the smallest inquisitive information state bigger than \mathbf{B} whose domain contains Q . A conversational *question update* by Q is defined accordingly:

³ Alternatively, an inquisitive information state can be defined as a quizposition set rather than a function from questions to views. An inquisitive information state is then a set of quizpositions subject to the following closure constraints:

1. *Closure under parthood*: if $A^Q \in \mathbf{S}$, R is part of Q , and $B \subseteq R$ is such that $\cup A \subseteq \cup B$, then $B^R \in \mathbf{S}$.
2. *Partial conjunctive closure*: if $A^Q, B^R \in \mathbf{S}$ and R is part of Q , then $AB^Q \in \mathbf{S}$.

This is analogous to the classical definition of an information state as a set of propositions closed under entailment and conjunction (see Hoek 2019).

⁴ Let $\mathcal{Q}(\mathcal{D})$ denote the space of all quizpositions about a question in \mathcal{D} , and let $[0, 1]$ be the real unit interval. Then an *inquisitive probability* with domain \mathcal{D} is a function $\Pr: \mathcal{Q}(\mathcal{D}) \rightarrow [0, 1]$ satisfying the following constraints:

1. *Normality*: For all $Q^Q \in \mathcal{D}$, $\Pr(Q^Q) = 1$.
2. *Partial Additivity*: if $Q, R \in \mathcal{D}$, R is part of Q , and A^Q is inconsistent with B^R , $\Pr(A^Q \vee B^R) = \Pr(A^Q) + \Pr(B^R)$.

Here $(A^Q \vee B^R) = \neg(\neg A^Q \wedge \neg B^R)$. The set of all quizpositions with probability 1 form a coherent inquisitive information state (just as the set of propositions with classical probability 1 always form a classical information state). The crucial difference with a classical probability is that additivity only holds in some cases. For instance, *Vacationer's* Jane is highly confident that *her phone is upstairs* (probability above 0.9), and is uncertain that *she has everything* (probability above 0.3). So her subjective probabilities in these inconsistent quizpositions sum to a value greater than 1. Partial additivity admits that possibility in this case, while classical (total) additivity would not. (For a decision-theoretic motivation of this definition of credence states, see Hoek 2019, Ch. 3–4.).

$$\langle D, C, \mathbf{B}_1, \dots, \mathbf{B}_n \rangle + Q = \langle DQ, C + DQ, \mathbf{B}_1 + DQ, \dots, \mathbf{B}_n + DQ \rangle$$

Much as before, the characteristic effect of giving a (partial) answer to the QUD is that both the participants' views of that question and the view in the common ground are updated with the new information (Stalnaker 1974, 2014; Heim 1988):

An *update* of the information state \mathbf{B} by the quizposition A^Q , written $\mathbf{B} + A^Q$, is the smallest inquisitive information state whose belief set includes $|\mathbf{B}| \cup \{A^Q\}$. A conversational *answer update* is defined accordingly:
 $\langle D, C, \mathbf{B}_1, \dots, \mathbf{B}_n \rangle + A^Q = \langle D, C + A^Q, \mathbf{B}_1 + A^Q, \dots, \mathbf{B}_n + A^Q \rangle$
 where Q is any part of D .

In addition to posing questions and answering questions that have already been posed, it is also possible to combine both these conversational moves in one by asserting a quizposition that answers a new question. In *Dinner Danger*, for instance, Bob reminds Alice that Darian is coming to the party. This assertion has the effect of both raising the question whether Darian will attend and of answering that question.

For an arbitrary quizposition A^Q , the conversational *update* by A^Q is defined thus:

$$\langle D, C, \mathbf{B}_1, \dots, \mathbf{B}_n \rangle + A^Q = \langle DQ, C + A^{DQ}, \mathbf{B}_1 + A^{DQ}, \dots, \mathbf{B}_n + A^{DQ} \rangle$$

Both question and answer updates are special cases of this generalized notion of a conversational update, with a question update by Q being equivalent to an update with the tautological quizposition Q^Q . Conversation updates thus capture the effect on the state of opinion of both the new attention on some issue and the new information acquired.

Besides the common ground, it will be useful to distinguish the *focused common ground*, the common ground's view $C(D)^D$ on the QUD—this is the part of the common ground on which the participants are focussing their joint attention. Or to put it differently, while the regular common ground models the complete stock of shared beliefs passively held by the interlocutors, the focussed common ground highlights those beliefs of which they are actively *aware*.⁵

A conversational update by A^Q makes A^Q part of the focused common ground, and it is inert just in case A^Q is not part of the focussed common ground. Thus it is natural to say an interlocutor is *in a position to assert* A^Q whenever A^Q is not part of the focused common ground. In particular, this means one can sometimes felicitously make an assertion of a quizposition that is already in the common ground, provided only that it is not in the focused part. As we will see in the next section,

⁵ We could elaborate the model by adding questions indexed to the participants, modelling their individual states of attention (subject to the constraint that the QUD should be part of each individual state). This would then allow us to distinguish passive beliefs and active awareness at the individual level as well—roughly in line with Fagin and Halpern 1988; Schipper et al. 2015. For our purposes here, we have no need for this complication: it is sufficient to capture collective attention. We largely eschew the language of awareness in this paper, which we think can sometimes obscure rather than clarify important cognitive distinctions (see Fritz and Lederman 2015). But those who are used to thinking in these terms should bear in mind that the individual belief states in our model are intended to capture passive belief and not active awareness.

even when this does not affect the common ground or any of the belief states involved, such assertions can still expand the QUD and thereby steer the subsequent conversational dynamics.

As in the classical case, not all updates yield a coherent information state. For instance, $\mathbf{B} + A^Q$ is not a coherent information state when $\neg A^Q \in \mathbf{BI}$. But question updates, too, can result in incoherence when the added question brings a previously hidden inconsistency to light. Our simple update rule models what happens when an agent learns something that coheres with their extant beliefs, but not what happens when an agent learns something that requires a *revision* of their extant beliefs. Incorporating the possibility of belief revision into the model is possible, but adds complications that need not deter us here—see Berto 2019 and Bledin and Rawlins 2020 for relevant discussion.

4 Applications

In *Vacationer*, the initial question under discussion is *Does Jane have everything?* We treat this as a polar question, with complete answers $Yes^{Everything?}$ and $No^{Everything?}$. Phil's questions are all of the form *Do you have X?*, and again we will think of those as polar questions with answers $Yes^{X?}$ and $No^{X?}$. Agnosticism about any of these questions is modelled as possession of a tautologous view *Yes or No^X*. Now we use \mathbf{B}_J to denote *Vacationer* Jane's belief state prior to the conversation with Phil. At this point, Jane's views are as follows:

- $\mathbf{B}_J(Everything?) = Yes\ or\ No$
- $\mathbf{B}_J(Toothpaste?) = Yes$
- $\mathbf{B}_J(Passport?) = Yes$
- $\mathbf{B}_J(Phone?) = No$

Even if she is not attending to it, Jane already has all this information from the start—Phil does not give her any information about her luggage. All Phil does is direct her attention, bringing these various pre-existing beliefs into focus. In particular Jane's answer $\mathbf{B}_J(Phone?)$ entails that the answer to $\mathbf{B}_J(Everything?)$ should be "No." But at this stage, Jane has failed to draw this inference. It is important to note that this tension is possible only because Jane has failed to consider the questions *Everything?* and *Phone?* in conjunction. In the model this is represented by excluding the conjunctive question *Does Jane have everything and does Jane have her phone?* from \mathbf{B}_J 's domain. Since Phil has no information, his belief state \mathbf{B}_P answers *Yes or No* to all these questions, and hence the same must be true for the initial common ground \mathbf{C}_0 .

The initial state of Jane and Phil's conversation is the tuple $\langle Everything?, \mathbf{C}_0, \mathbf{B}_J, \mathbf{B}_P \rangle$. As Jane and Phil's conversation progresses, each of Phil's questions results in a question update. Thus Phil successively refines the question under

discussion, thereby expanding the range of issues to which Jane and Phil are jointly attending:

$$\begin{aligned} D_0 &= \textit{Everything?} \\ D_1 &= \textit{Everything?} \wedge \textit{Toothpaste?} \\ D_2 &= \textit{Everything?} \wedge \textit{Toothpaste?} \wedge \textit{Passport?} \\ D_3 &= \textit{Everything?} \wedge \textit{Toothpaste?} \wedge \textit{Passport?} \wedge \textit{Phone?} \end{aligned}$$

Here D_n represents the question under discussion after Phil has asked his n th question. As Jane processes each of these questions, her belief state first progresses from \mathbf{B}_J to $\mathbf{B}_{J,1} = \mathbf{B}_J + D_1$, then to $\mathbf{B}_{J,2} = \mathbf{B}_J + D_2$, and finally to $\mathbf{B}_{J,3} = \mathbf{B}_J + D_3$. In the first two of these transitions, nothing notable happens: the question domain of Jane's belief state expands, but she does not gain any information about her initial question, *Everything?*.

The final update, however, settles that question for Jane: $\mathbf{B}_{J,3}(\textit{Everything?}) = \textit{No}^{\textit{Everything?}}$. Intuitively, the reason for this is that this question makes Jane consider her views about *Everything?* and *Phone?* together for the first time. This causes her to see the relevance of the latter to the former. The model captures this intuitive idea. Because D_3 is included in $\mathbf{B}_{J,3}$'s domain, and the question *Phone?* is part of the question D_3 , it follows from the *Harmony* constraint on inquisitive information states that, since $\mathbf{B}_{J,3}(\textit{Phone?}) = \textit{No}^{\textit{Phone?}}$, the alternative answer *Yes*^{*Phone?*} must be ruled out by $\mathbf{B}_{J,3}(D_3)$. But then, since *Everything?* is also part of D_3 , we get that.

$$\mathbf{B}_{J,3}(\textit{Everything?}) = \cup \{E \in \textit{Everything?} : E \text{ is consistent with } \mathbf{B}(D_3)\} = \textit{No}.$$

In short, Jane's acquisition of a view about D_3 forces her to bring her views about the parts of D_3 into harmony, including her views about *Phone?* and *Everything?* Since she preserves her prior view *No*^{*Phone?*}, she must therefore acquire a belief in its consequence, *No*^{*Everything?*}.

Thus, our model explains how Phil's questioning can lead Jane to new beliefs without giving her any new information. When we consider the common ground, we observe a parallel pattern. At each point, Jane's assertions settle the additional issue that Phil has just raised. Thus, the successive common grounds after each of Jane's answers come to reflect Jane's successive views on the QUD:

$$\begin{aligned} C_0(D_0) &= \mathbf{B}_J(D_0) = \textit{Yes or No}^{\textit{Everything?}} \\ C_1(D_1) &= \mathbf{B}_{J,1}(D_1) = \textit{Yes or No}^{\textit{Everything?}} \wedge \textit{Yes}^{\textit{Toothpaste?}} \\ C_2(D_2) &= \mathbf{B}_{J,2}(D_2) = \textit{Yes or No}^{\textit{Everything?}} \wedge \textit{Yes}^{\textit{Toothpaste?}} \wedge \textit{Yes}^{\textit{Passport?}} \\ C_3(D_3) &= \mathbf{B}_{J,3}(D_3) = \textit{No}^{\textit{Everything?}} \wedge \textit{Yes}^{\textit{Toothpaste?}} \wedge \textit{Yes}^{\textit{Passport?}} \wedge \textit{No}^{\textit{Phone?}} \end{aligned}$$

Here $C_n(D_n)$ is the focused part of the common ground after Jane has answered Phil's n th question. Crucially, the final focused common ground C_3 not only settles the question that preceded it, *Phone?*, but also the initial question *Everything?* that Phil and Jane started out with.

As noted in Sect. 3, the question updates resulting from Phil's interrogation are equivalent to updates by tautological quizpositions. Nonetheless, Jane doesn't just

add tautologies to her stock of beliefs—she acquires new contingent beliefs as well. In particular she acquires the quizposition $No^{Everything?}$ —that is, she realises that she does not yet have everything. In this way, our model captures the way that posing new questions can bring out previously unseen deductive consequences of our beliefs. Deductive inquiry may thus be understood in terms of the posing of new questions (Hoek forthcoming). At the theoretical limit of this process is an agent that has views on every question. On our model, such an agent would be logically omniscient: seeing all the deductive consequences of their beliefs, they are just like a classical agent.⁶ Contrapositively, failures to see the consequences of one’s beliefs always correspond to a failure to have a well-defined view on certain questions, on this way of understanding things.

We can see these same principles at work in more complex deliberative scenarios like *Dinner Danger*. Here, the main questions in play are *Will the party be fun?*, *Will Darian come to the party?*, and *Will Gerry come to the party?*. These three questions, abbreviated *Fun?*, *Darian?* and *Gerry?*, may be treated as polar questions as before. We will also need to consider a fourth question to connect them, *Conditional Fun?*. This question has three complete answers: *Darian and Gerry both come and the party will be fun^{Conditional Fun?}*, *Darian and Gerry both come and the party will not be fun^{Conditional Fun?}* and *Darian and Gerry won’t both come^{Conditional Fun?}*. Alice and Bob’s shared view on this final question is that *the party will not be fun if Darian and Gerry both come*, and we shall treat this as a disjunction of the latter two possibilities.

We can assume that these four questions are included in the domain of all three information states in the initial state $\langle Fun?, C, B_A, B_B \rangle$ of Alice and Bob’s conversation. In particular, Bob’s initial belief state B_B includes only what is common ground about these four questions:

$B_B (Darian?)$	$= C (Darian?)$	$= Yes$
$B_B (Gerry?)$	$= C (Gerry?)$	$= Yes\ or\ No$
$B_B (Conditional\ Fun?)$	$= C (Conditional\ Fun?)$	$= If\ they\ both\ come,\ no\ fun$
$B_B (Fun?)$	$= C (Fun?)$	$= Yes\ or\ No$

Alice has a bit more information, namely $B_A(Gerry) = Yes$, but otherwise her views on these four questions match Bob’s. While Alice’s beliefs entail the answer to the fourth question as well, she has not yet put the pieces together. This means, in particular, that \mathcal{D}_{B_A} does not contain the big question

$$Q = Fun? \wedge Darian? \wedge Gerry?$$

⁶ *Proof.* Suppose an agent X has beliefs on every question. We show that X’s belief state B_X is closed under multi-premise entailment. First, let’s see why B_X must be closed under conjunction. Suppose A^Q and B^R are in X’s belief state B_X . Since \mathcal{D}_{B_X} contains every question, it must contain the conjunction QR. Since Q and R are both part of QR, it follows by *Harmony* that $B_X(QR)$ must rule out every QR-answer that is inconsistent with either A^Q or B^R . Hence $B_X(QR)$ must entail the conjunction AB^{QR} , whence $AB^{QR} \in B_X$. Now we show closure under single-premise entailment. Let A^Q be any quizposition in B_X , and let B^R be some quizposition that is entailed by A^Q . Since Q is part of QR, it follows from *Harmony* that $B_X(QR)$ rules out every possibility ruled out by A^Q . Hence $B_X(QR)$ entails AB^{QR} , and $AB^{QR} \in B_X$. But B^R is part of AB^{QR} , and so $B^R \in B_X$. ■

which would bring all the pertinent pieces together. Hence $\mathcal{D}_C \subseteq \mathcal{D}_{\mathbf{B}_A}$ does not contain Q either. On the other hand, it is natural to think that Bob, who is worried about this specific scenario, has considered this question, so that $Q \in \mathbf{B}_B$. Given the Harmony constraint, Bob's view on Q must be the conjunction of the views listed above, namely the following:

$$\mathbf{B}_B(Q) = \textit{Darian will come and if Gerry comes too, the party won't be fun}^Q$$

Thus Bob has beliefs that Alice does not have, even if he does not have information that she lacks. Unlike Alice's, Bob's views on the questions *Fun?* and *Gerry?* are linked. (If Alice's views on *Fun?* and *Gerry?* were linked in the way Bob's views are, she would believe *No*^{*Fun?*} since she believes *Yes*^{*Gerry?*}.)

The *Dinner Danger* conversation starts with Alice and Bob raising the question whether the party will be fun, bringing us to the initial state $\langle \textit{Fun?}, \mathbf{C}, \mathbf{B}_A, \mathbf{B}_B \rangle$. Bob then asserts that *Darian will come* — i.e. that *Yes*^{*Darian?*}. Note that, from the classical perspective, this is a peculiar thing to assert, since this information is already common ground. The main intended effect of this assertion is not to change the common ground or to affect Alice's belief state, but rather to bring this public belief into the focus of attention by expanding the QUD to include the question of Darian's attendance. Indeed, if we assume that the posterior QUD $\textit{Fun?} \wedge \textit{Darian?} \in \mathcal{D}_C \subseteq \mathcal{D}_{\mathbf{B}_A}$, this change in focus is the only effect of the assertion:

$$\langle \textit{Fun?}, \mathbf{C}, \mathbf{B}_A, \mathbf{B}_B \rangle + \textit{Yes}^{\textit{Darian?}} = \langle \textit{Fun?} \wedge \textit{Darian?}, \mathbf{C}, \mathbf{B}_A, \mathbf{B}_B \rangle$$

The shift in attention sets the stage for Bob's question, *Gerry?* It is only because he has mentioned *Darian?*, that this question will bring Bob and Alice's shared view of *Conditional Fun?* into focus. That's because *Conditional Fun?* is part of $Q = \textit{Fun?} \wedge \textit{Darian?} \wedge \textit{Gerry?}$ but not part of $\textit{Fun?} \wedge \textit{Gerry?}$. So without the stage-setting, the question *Gerry?* might have seemed irrelevant. As it is, the result is the conversational state $\langle Q, \mathbf{C} + Q, \mathbf{B}_A + Q, \mathbf{B}_B \rangle$.

With these two conversational moves, Bob has in effect put his own view about Q on the table for public consideration, making it the focused part of the common ground:

$$(\mathbf{C} + Q)(Q) = \textit{Darian will come and if Gerry comes too, the party won't be fun}.$$

Alice, meanwhile, is now led to combine her views on *Darian?*, *Gerry?* and *Conditional Fun?* into a single view about Q:

$$(\mathbf{B}_A + Q)(Q) = \textit{Darian and Gerry will come, and the party won't be fun}.$$

Consequently, Alice's focused beliefs outstrip the focused common ground for the first time in the conversation, putting her in a position to contribute additional information on the question under discussion by asserting that *Yes*^{*Gerry?*} (i.e. that Gerry will come). She does so, and thus we finally arrive at an equilibrium state where all relevant information, and all the questions, have been shared:

$$\langle Q, \mathbf{C} + \textit{Gerry comes}^Q, \mathbf{B}_A + Q, \mathbf{B}_B + \textit{Gerry comes}^{\textit{Gerry?}} \rangle.$$

As the notation shows, Bob gained some information from Alice, while Alice gained only the question that Bob was worried about. Both the information and the question became part of the common ground. Crucially, resolution of Alice and Bob's deliberation was mutually beneficial even though only Alice contributed information. By explaining his worry to Alice, Bob combined some information that they already shared at the outset, thereby showing a link between *Fun?* and *Gerry?* that Alice did not previously see.

The consensus Alice and Bob achieve in this exchange is the joint result of the information Alice contributes to the common ground and the questions that Bob adds to the QUD. In the classical update model, agents with jointly consistent beliefs arrive, in the limit, at reflective equilibrium where all their beliefs have been shared, and all agents have the same beliefs. In the inquisitive model, agents with jointly consistent beliefs likewise converge to a limit of consensus. To reach this limit point it is not sufficient that all participants provide all the answers they are in a position to give. They must also pose every question they are in a position to ask, since these questions can draw out further answers, and tease out further consequences of the view on the table. When reflective equilibrium is eventually reached, all agents believe every part of the conjunction of all the beliefs they started out with.⁷

It is also worth considering how the model behaves in a less idealised setting, where the initial beliefs are not assumed to be consistent. In such a situation, there is no convergence to a predetermined limit. Depending on which questions become salient, different outcomes may be reached. This point can be illustrated with a variant of the *Dinner Danger* scenario:

Example 3. Dinner Delight. As in *Dinner Danger*, Alice and Bob are getting ready for a dinner party—and in addition to the information mentioned there, Alice also believes that Fiona got a puppy last week, and that she is the kind of person who would take a puppy to a party.

B: I'm looking forward to the party.

A: Why? I don't know if it will be any good.

B: Well, Fiona is coming, and didn't she get that puppy last week?

A: You're right, I hadn't thought of that! What could be more fun than a puppy? This party is going to be a blast, we *have* to go.

We might model this exchange much as we did before, if we make sure to add to Alice's initial belief state \mathbf{B}_A the quizpositions that *Fiona is coming*, that *Yes, Fiona got a puppy last week* and that *If Fiona is coming, and got a puppy last week, the party will be fun*. Note that, in doing so, we will have rendered \mathbf{B}_A inconsistent, in that Alice's beliefs now entail both that the party will be fun and also that it won't

⁷ Of course this theoretical limit is not practically achievable. Realistically, the scope of our attention is limited (e.g. Koralus 2014), and so it is not possible for agents to attend to the totality of all their beliefs. One upshot is that real agents are forced to make strategic decisions about which questions they do and do not pose. See Rothe et al. (2018) for an empirical study of how we choose questions; for theoretical work on good questions, see Friedman 2018, Pérez Carballo 2018.

be fun. Nonetheless this is an admissible, coherent inquisitive belief state, because Alice is unaware of the conflict.

Dinner Danger and *Dinner Delight* illustrate the path-dependence of deliberative exchanges, in that different lines of questioning can lead to directly opposite conclusions. Our model captures this path-dependence. This also raises an intriguing possibility. For suppose Bob in fact has full knowledge of Alice's antecedent belief state B_A . Then he is in a position to predict how these two different lines of questioning will lead to different outcomes, and to select one or the other depending on which conclusion he would like for Alice to draw.

Thus our model explains how it is possible for a skilled interlocutor with sufficient knowledge of their audience's priors to manipulate the opinions of their audience through strategic questioning. Of course, this possibility has always been known to rhetoricians. This question-based form of belief manipulation has the dubious advantage of avoiding the liability that results from lies or cherry-picked facts. Its success moreover does not require the audience's trust. After all, as long as the speaker is merely asking questions, they are ostensibly allowing their audience to make up their own minds.

5 Forming collective attitudes

Our focus in this paper has been on substantiating the claim that deliberation provides non-informational pathways for individual attitude change and highlighting the role of questions in inducing shifts in individuals' attention. But much of what we have said could be reframed as an exploration of how collective opinion changes through deliberation and, in particular, of the effect of deliberation on the focused common ground. To do this, and to bring out in an informal way the implications for collective decision making, it suffices to make explicit (or to specify) the preferences of our deliberators.

Consider *Vacationer*, in which Jane and Phil are about to depart on vacation together. Phil's questions serve, as before, both to bring distributed information into the common ground and to focus Alice's attention on questions which she, but not Phil, is able to address. The upshot of their conversation is collective recognition that they don't have everything they need and presumably a collective decision not to depart immediately. Similarly, in *Dinner Danger* Alice and Bob must decide whether *they* want to go to the party. Collectively they hold all the information they need in order to agree not to, but deliberation is required to bring this into the common ground. Indeed, initially Bob is unsure whether they should go to the party while Alice thinks they should. The effect of Bob's questions is to both get Alice to realise that, given what she knows, they should not go and for her to communicate this information to Bob.

In both of these cases the deliberators have a mutual interest in the truth that derives from shared preferences over the possible consequences of the collective decision they must take. This suffices to motivate them to not only share whatever private information they hold, but also to draw the attention of other deliberators to questions that they know to be pertinent to the decision they must take together.

More generally, it is to be expected that rational deliberators with shared interests will, in the course of deliberation, simply pool their attentions as well as their information (at least regarding any propositions relevant to their collective decision).

By contrast, consider a variant of *Dinner Delight* in which we suppose that Alice would, all things considered, prefer not to go to the party if both Damian and Gerry are going, even if Fiona is bringing her puppy, while Bob would, all things considered, prefer to go even if both Damian and Gerry will be there. As we noted before, Bob is able to exploit what he knows about Alice's preferences, and how they depend on what she knows, to focus her attention on the reasons she has for going rather than those she has for not. This enables Bob to steer deliberation in such a way as to achieve consensus on the action that best serves his preferences.

These simple examples illustrate how a collective decision can be arrived at by building a consensus in the focused common ground. Deliberation does not always lead to consensus, of course, and the availability of another channel for inducing revision of individual beliefs does not change this fact. But our model does point to the possibility of changing the distribution of individual opinion in a manner which facilitates collective decision-making even in the absence of consensus; for instance by making possible compromises more visible or, more generally, allowing for the application of an aggregation rule.

That deliberation might play this role has long been a matter of conjecture in the literature on deliberative democracy (Elster 1986; Miller 1992; Estlund and Lademore 2018). It has been suggested, for instance, that deliberation serves to increase proximity to single-peakedness: a property of profiles of individuals' opinion that suffices for the existence of an aggregation schema that avoids Arrow's famous impossibility result (Dryzek and List 2003). There is some empirical evidence in support of this hypothesis (List et al. 2013; Luskin et al. 2002) and various models based on preference convergence have been developed to explain it (Rad and Roy 2020; Perote-Peña and Piggins 2015). The exact mechanisms by which this might occur have yet to be spelled out, however. Information exchange alone evidently does not suffice. But the process of drawing deliberators' attention to particular issues when forming their preferences, in combination with information exchange, does provide a plausible candidate for such a mechanism.

The notion of the focused common ground is crucial here. To illustrate, suppose that in *Dinner Delight* Alice prefers going to the party if there will be a puppy but not if both Damian and Gerry are there. Suppose that Bob, on the other hand, prefers to go to the party in either case. Now consider two possible cases that vary with respect to Bob's preferences regarding Gerry's presence at the party. Suppose firstly that he prefers that Gerry not be there. Then although Alice and Bob have conflicting preferences over whether to stay at home or to go to the party in the event that Gerry is going to be there, their preferences over the fully specified possible outcomes are nonetheless single-peaked; for instance, over the ordered set {stay at home, party without puppy or Gerry, party with puppy but not Gerry, party with both puppy and Gerry, party with Gerry, but puppy is absent}. And, loosely, this means that there is a possibility of reaching a compromise even if deliberation brings the information about Gerry into the focused common ground.

Now suppose instead that Bob prefers Gerry to be there. In this case Alice and Bob's preferences over the fully specified possible outcomes are not just in conflict, but there is in fact no way of arranging these outcomes in such a way that both have single-peaked preferences over them so-arranged. A deliberative exchange between them in which the QUD was too inclusive, when both the question of whether *Gerry will be at the party* and the question of whether *the puppy will be at the party* are raised, would risk creating a deadlock. In contrast, when only the question about the puppy's presence is part of the focused common ground, Alice and Bob's preferences *will* be single peaked on the set {stay at home, party with puppy, party and puppy is absent}. If Bob prefers the puppy to be at the party, their preferences will also have a common peak (go to the party and the puppy is there), but their preferences will be single-peaked even if he does not.

This second case exhibits the possibility of resolving collective decision problems by focusing the attention of deliberators on the 'right' dimensions of comparison. As such it sheds lights on political issues where such dimension sensitivity is present. The issue of whether immigration should be restricted might, for instance, be approached via the question of what its implications are for the availability of public goods such as housing or schools. Equally it might be approached via the question of what its implications are for economic growth. Quite plausibly for some people the attitude they take to immigration will depend to some degree on which of these questions that they focus on.⁸ As in *Dinner Delight* this fact opens up the possibility of steering opinion formation, and thereby collective decision making, by controlling what questions are brought to deliberators' attention. Arguably, indeed, one of the functions of political leadership is to build support for collective decisions by focusing attention on dimensions on which agreement can be found, a function that can be fulfilled by simply asking the right questions.

To illustrate this, suppose that the majority of the electorate hold the belief that bringing skills into the country spurs economic growth and that a majority hold the belief that increases in the population (above the norm) lead to temporary overgrounding in schools. Suppose also that initially both beliefs are not connected in most people's mind with the issue of whether immigration should be restricted. Simply by asking the question *Do immigrants bring skills into the country?* might lead many to draw the inference that immigration is beneficial to the economy. Likewise asking the question *Does immigration lead to a population increase?* might lead many to draw the inference that immigration puts pressure on public services. If this is so, then asking one question rather than another would serve to move deliberation towards agreement on a particular view on whether or not immigration should be restricted simply by focusing attention in a particular way.

Further progress in identifying the conditions under which deliberation creates a focused common ground that supports collective decision-making can only be made by specifying more precisely the settings in which deliberation occurs. A deliberative setting is characterised by the deliberative speech acts available to deliberating agents (assertions, questions, etc.), the rules of deliberation (who can speak, in what

⁸ See Koralus and Alfano 2017 for evidence that moral judgement is question-sensitive.

order, etc.) and the preferences of deliberators over any consequences of actions that might be taken as a result of deliberation. Together with a specification of the the initial conversational state (the prior belief/informational states of the group of deliberators) and the opinion revision rules applied by deliberators, this would suffice to define a ‘deliberation game’ in which individuals choose speech acts and update their attitudes as quizpositions are brought into the common ground. An equilibrium of a deliberation game would be a point at which no participants have any incentive to make a speech act. The informal remarks of this section suggest that an analysis of this kind would produce a wider range of deliberational equilibria than is possible on the basis of information exchange alone. But evidently there is much work still to be done here.

Acknowledgements The authors are grateful for the feedback from the audience of the CoLAForm Seminar series, Sam Carter, Kyle Blumberg, Simon Goldstein and Nico Kirk-Gianinni, and thank Dane Stocks for his valuable research assistance. They are also grateful to Princeton University for hosting them both in the 2019-2020 academic year and thereby making this collaboration possible.

Funding Not applicable.

Availability of data and material Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Austen-Smith D, Feddersen TJ (2006) Deliberation, preference uncertainty, and voting rules. *Am Polit Sci Rev* 100(2):209–217. <https://doi.org/10.1017/S0003055406062113>
- Berto F (2019) Simple hyperintensional belief revision. *Erkenntnis* 84(3):559–575. <https://doi.org/10.1007/s10670-018-9971-1>
- Bledin J, Rawlins K (2020) Resistance and resolution: attentional dynamics in discourse. *J Semant* 37(1):43–82. <https://doi.org/10.1093/jos/ffz015>
- Bradley R (2009) Revising incomplete attitudes. *Synthese* 171(2):235–256. <https://doi.org/10.1007/s11229-009-9638-1>
- Carruthers P (2018) Basic questions. *Mind Lang* 33(2):130–147. <https://doi.org/10.1111/mila.12167>
- Ciardelli I, Groenendijk J, Roelofsen F (2019) *Inquisitive semantics*. Oxford University Press, Oxford
- Dryzek JS, List C (2003) Social choice theory and deliberative democracy: a reconciliation. *Br J Polit Sci* 33(1):1–28. <https://doi.org/10.1017/S0007123403000012>

- Elster J (1986) The market and the forum. In: Elster J, Hylland A (eds) *Foundations of social choice theory*. CUP, Cambridge (Reprinted in: Bohman J, Rehg W (eds) *Deliberative democracy: essays on reason and politics*. MIT Press, Cambridge, pp 3–33, 1997)
- Estlund D, Landmore H (2018) The epistemic value of democratic deliberation. In: Bächtiger A, Dryzek J, Mansbridge J, Warren M (eds) *The Oxford handbook of deliberative democracy*. Oxford University Press, Oxford, pp 113–131
- Fagin R, Halpern JY (1988) Belief, awareness, and limited reasoning. *Artif Intell* 34:39–76. [https://doi.org/10.1016/0004-3702\(87\)90003-8](https://doi.org/10.1016/0004-3702(87)90003-8)
- Friedman J (2017) Why suspend judging? *Noûs* 53:296–315. <https://doi.org/10.1111/nous.12222>
- Friedman J (2018) Junk beliefs and interest-driven epistemology. *Philos Phenomenol Res* 97:568–583. <https://doi.org/10.1111/phpr.12381>
- Fritz P, Lederman H (2015) Standard state-space models of unawareness. In: *Proc. of TARK XV*
- Gemes K (1994) A new theory of content I: basic content. *J Philos Log* 23(6):595–620. <https://doi.org/10.1007/BF01052779>
- Ginzburg J (2012) *The interactive stance: meaning for conversation*. Oxford University Press, Oxford
- Groenendijk J, Stokhof M (1984) *Studies on the semantics of questions and the pragmatics of answers*. Dissertation, University of Amsterdam
- Heim I (1988) On the projection problem for presupposition. In: *Proc. of second annual west coast conference on formal linguistics*. Stanford University, pp 114–26 (Reprinted in: Portner S, Partee BH (eds) *Formal semantics*. Wiley-Blackwell, pp 249–260, 2002)
- Hintikka J (1962) *Knowledge and belief*. Cornell University Press
- Hoek D (2019) *The web of questions: inquisitive decision theory and the bounds of rationality*. PhD Dissertation, New York University
- Hoek D (forthcoming) Minimal rationality and the web of questions. In: Kindermann D, van Elswyk P, Egan A (eds) *Unstructured content*. Oxford University Press, New York
- Koralus P (2014) The erotetic theory of attention: questions, focus and distraction. *Mind Lang* 29(1):26–50. <https://doi.org/10.1111/mila.12040>
- Koralus P, Alfano M (2017) Reasons-based moral judgment and the erotetic theory. In: Bonnefon J-F, Trémolière B (eds) *Moral inferences*. Routledge/Taylor & Francis Group, Milton Park, pp 77–106
- Koralus P, Mascarenhas S (2013) The erotetic theory of reasoning. *Philos Perspect* 27(1):312–365. <https://doi.org/10.1111/phpe.12029>
- Landa D, Meirowitz A (2009) Game theory, information, and deliberative democracy. *Am J Polit Sci* 53(2):427–444. <https://doi.org/10.1111/j.1540-5907.2009.00379.x>
- Lehrer K, Wagner C (1981) Rational consensus in science and society. Reidel, Dordrecht
- List C (2017) Democratic deliberation and social choice: a review. In: Bächtiger A, Dryzek JS, Mansbridge JJ, Warren M (eds) *The Oxford handbook of deliberative democracy*. Oxford University Press, Oxford
- List C, Luskin RC, Fishkin JS, McLean I (2013) Deliberation, single-peakedness, and the possibility of meaningful democracy: evidence from deliberative polls. *J Polit* 75(1):80–95. <https://doi.org/10.1017/S0022381612000886>
- Luskin RC, Fishkin JS, Jowell R (2002) Considered opinions: deliberative polling in Britain. *Br J Polit Sci* 32(3):455–487. <https://doi.org/10.1017/S0007123402000194>
- Miller D (1992) Deliberative democracy and social choice. *Polit Stud* 40(1 Suppl):54–67. <https://doi.org/10.1111/j.1467-9248.1992.tb01812.x>
- Pérez Carballo A (2018) Good questions. In: Dunn J, Ahlstrom-Vij K (eds) *Epistemic consequentialism*. Oxford University Press, Oxford, pp 123–145
- Perote-Peña J, Piggins A (2015) A model of deliberative and aggregative democracy. *Econ Philos* 31(1):93–121. <https://doi.org/10.1017/S0266267114000418>
- Rad S, Roy O (2020) Deliberation, single-peakedness, and voting cycles
- Roberts C (2012) Information structure: towards an integrated formal theory of pragmatics. *Semant Pragmat* 5(6):1–69. <https://doi.org/10.3765/sp.5.6>
- Rothe A, Brenden ML, Gureckis TM (2018) Do people ask good questions? *Comput Brain Behav* 1:69–89. <https://doi.org/10.1007/s42113-018-0005-5>
- Schipper BC (2015) Awareness. In: Van Ditmarsch H et al (eds) *Handbook of epistemic logic*. College Publications, London, pp 77–146
- Stalnaker R (1974) Pragmatic presuppositions. In: Munitz MK, Unger P (eds) *Semantics and philosophy*. New York University Press, New York, pp 197–213
- Stalnaker R (2014) *Context*. Oxford University Press, Oxford

- Sunstein C (2006) Deliberating groups versus prediction markets (or Hayek's challenge to Habermas). *Episteme* 3(3):192–213. <https://doi.org/10.1353/epi.2007.0007>
- Yalcin S (2011) Nonfactualism about epistemic modality. In: Egan A, Weatherson B (eds) *Epistemic modality*. Oxford University Press, New York, pp 295–332
- Yalcin S (2018) Belief as question-sensitive. *Philos Phenomenol Res* 97(1):23–47. <https://doi.org/10.1111/phpr.12330>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.