



**HAL**  
open science

## Reasoning in attitudes

Franz Dietrich, Antonios Staras

► **To cite this version:**

| Franz Dietrich, Antonios Staras. Reasoning in attitudes. 2022. halshs-03023015v3

**HAL Id: halshs-03023015**

**<https://shs.hal.science/halshs-03023015v3>**

Submitted on 8 Jul 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

CES

Centre d'Économie de la Sorbonne  
UMR 8174

**Reasoning in attitudes**

Franz DIETRICH, Antonios STARAS

**2022.15RR**

*Version révisée*



# Reasoning in attitudes

This version: June 2022<sup>1</sup>

Franz Dietrich  
Paris School of Economics & CNRS

Antonios Staras  
University of Cardiff

## Abstract

People reason not just in beliefs, but also in intentions, preferences, and other attitudes. They form preferences from existing preferences, or intentions from existing beliefs and intentions, and so on, often facing choices between rival conclusions. Building on Broome (2013) and Dietrich et al. (2019), we present a philosophical and formal analysis of reasoning in attitudes with or without facing such choices. Reasoning in attitudes is a mental activity that differs fundamentally from reasoning *about* attitudes, a form of theoretical reasoning by which one discovers rather than forms attitudes. Reasoning in attitudes has standard formal features (such as monotonicity), but is indeterministic (reflecting choice in reasoning). Like theoretical reasoning, it need not follow logical entailment, but for different reasons related to indeterminism. This makes reasoning in attitudes harder to model logically than theoretical reasoning.

## 1 Introduction

A growing philosophical literature about rationality and reasoning teaches us that beliefs are not the only locus of reasoning. You also reason in intentions, preferences, and other attitudes (e.g., Broome 2006, 2013, Kolodny 2005, 2007, Boghossian 2014, Dietrich et al. 2019). By reasoning you can form preferences from preferences; this makes your preferences more transitive. You can form the intention to help a child cross a street from believing you ought to; this makes you less akratic. You can form the same intention from intending to make the child happy and believing that reaching this end requires your help; this makes you more instrumentally rational.

Two questions guide this paper:

- (1) What *is* reasoning in attitudes, philosophically and formally?
- (2) Can such reasoning follow entailment of some type, and hence be modelled logically?

---

<sup>1</sup>We are grateful for inspiring feedback from colleagues, especially John Broome, Christian List, Robert Sugden, and Frederik van de Putte, as well as from referees. The paper subsumes our earlier working paper with main title “Reasoning in versus about attitudes”. Franz Dietrich acknowledges support from the French Research Agency through the grants ANR-17-CE26-0003, ANR-16-FRAL-0010 and ANR-17-EURE-0001.

**On question (1):** For one, reasoning in attitudes – which John Broome calls reasoning ‘with’ attitudes<sup>2</sup> – differs fundamentally from reasoning about attitudes, in which you discover rather than form attitudes. Reasoning in attitudes is an internal, first-person process of forming attitudes. It happens routinely in daily life. It often is practical reasoning, by creating intentions that cause actions. It matters directly to practical philosophy, psychology, and even artificial intelligence.<sup>3</sup> By contrast, reasoning about attitudes is a third-person, meta-level process of discovering attitudes, i.e., either your own attitudes (in an act of reflection) or someone else’s attitudes (e.g., another player’s intentions in a game; cf. Perea 2012). It is a form of theoretical reasoning, i.e., reasoning in beliefs. Indeed, it is theoretical reasoning *about attitudes*.

For another, reasoning in attitudes differs from other processes that also change your attitudes, including processes driven by external causes (music can create desires) and internal psychological processes that are purely automatic and unconscious (desires can cancel intentions that stand in the way). We focus exclusively on *reasoning*.

Our philosophical account of reasoning in attitudes starts from Broome’s (2013) influential work, and presents a novel analysis of choice in reasoning that shows how choice is possible and is inextricably related to practical reasoning. Our Broome-inspired take on reasoning contrasts with more liberal approaches that count far more mental processes as reasoning.<sup>4</sup> Our formalisation of reasoning draws on Dietrich et al. (2019) and on the apparatus of abstract consequence operators.

**On question (2):** As we shall see, reasoning in attitudes departs fundamentally from entailment, even after twisting the notion of entailment. The main cause is choice in reasoning.

The difficulty to model reasoning in attitudes logically might at first surprise. After all, logic is so useful when modelling (deductive) *theoretical* reasoning;<sup>5</sup> and logic even provides powerful tools to model attitudes, namely modal operators such as belief operators, preference operators, or intention operators.<sup>6</sup> One might thus have conjectured that logic can ‘somehow’ model reasoning in attitudes, modulo standard idealisations or abstractions that come with any formal model. Truth of this conjecture is implicitly presupposed by the common conception of an ideal reasoner as someone possessing a deductively closed set of attitudes. If (ideal) reasoning does not follow entailment, then deductive closure is not the characteristic mark of an (ideal) reasoner.

Our analysis will reject the conjecture. This does not mean that reasoning in attitudes cannot be modelled: such reasoning can (and will) be modelled using the abstract machinery of consequence operators, simply because consequence operators can represent *any* transformation of attitudes. Our doubts about modelling reasoning in attitudes

---

<sup>2</sup>Our terminology aims to prevent any confusion with ‘reasoning about attitudes’.

<sup>3</sup>Sophisticated intelligent systems use (artificial) reasoning to form (artificial) attitudes, including intentions that cause actions.

<sup>4</sup>An example is Drucker’s (2021) broader account of reasoning, which he calls ‘generalism’.

<sup>5</sup>Non-deductive reasoning can behave very differently, for instance non-monotonically. See Harman’s (1984) seminal distinction between reasoning (inference) and entailment (implication).

<sup>6</sup>Examples of logics with attitude operators are logics of preferences (e.g., Liu 2011), of beliefs (e.g., Halpern 2017), or of beliefs, desires and intentions (‘BDI logics’).

refer to the use of concrete logics, not abstract logic.

## 2 Reasoning in attitudes – deterministic case

This section discusses and formalises reasoning in attitudes without yet addressing choice in reasoning (to which Section 3 turns). The philosophical account follows Broome (2013), and the formalism follows Dietrich et al. (2019).

### 2.1 Attitudes and constitutions

The agent – ‘you’ – holds various attitudes, also called (mental) states, such as: believing it snows, desiring to feel warm, intending to dress warm, preferring snow to rain, etc. The set of all possible attitudes is denoted  $M$ . Those attitudes which you possess form your *constitution*. Formally:

**Definition 1** A *constitution* is a set  $C \subseteq M$  of attitudes (‘your’ attitudes).

Think of attitudes in  $M$  as pairs of an attitude-content and an attitude-type. For many philosophers, contents are propositional: they are *single* propositions for monadic attitudes like intention, *pairs* of propositions for dyadic attitudes like preference, etc.<sup>7</sup> One could make this structure of states formally explicit.<sup>8</sup>

We use the term ‘attitude’ not only for mental states in  $M$  (such as: desiring to be warm), but also for attitude-types (such as: desire).

### 2.2 The Broomean account in a nutshell

Your constitution changes through reasoning. In reasoning, you form a (conclusion-)attitude from existing (premise-)attitudes: you form beliefs from beliefs, intentions from beliefs and desires, preferences from preferences, etc. The process is causal: the premise-attitudes cause the conclusion-attitude. It constitutes a conscious mental act. You bring the premise-attitudes to mind by saying their contents to yourself, normally using internal speech. This lets you construct a new attitude, again using (internal) speech. You might reason:

$$\textit{Paying taxes is legally required. So, I shall pay taxes.} \tag{1}$$

This is reasoning from a single premise-attitude, a belief, to an intention. The conclusion-attitude has the content *I pay taxes*. What you say however involves ‘shall’, a linguistic marker indicating that you entertain the content as an intention. In reasoning, you express to yourself the *marked contents* of your premise- and conclusion-attitudes, not

<sup>7</sup>On ‘propositionalism’, see Felappi (forthcoming). For a critical take, see Montague (2007).

<sup>8</sup>Let  $L$  be a set of *propositions*, and  $A$  a set of *attitude-types*, each carrying an *arity*  $n \in \{1, 2, \dots\}$ , usually 1 (monadic attitudes) or 2 (dyadic attitudes). Plausibly,  $A$  contains at least belief *bel* (monadic), desire *des* (monadic), intention *int* (monadic), preference  $\succ$  (dyadic), and indifference  $\sim$  (dyadic). Finally, define *attitudes* in  $M$  as tuples  $m = (p_1, \dots, p_n, a)$  where  $a$  is an attitude type in  $A$ ,  $n$  is its arity, and  $p_1, \dots, p_n$  are propositions in  $L$ . So,  $(p, \textit{bel})$  is believing  $p$ ,  $(p, \textit{int})$  is intending  $p$ ,  $(p, q, \succ)$  is preferring  $p$  to  $q$ , etc.

the contents simpliciter. Marked contents are contents marked by ‘how’ the content is entertained: as a belief, or intention, etc. The English language contains markers for various attitude types, allowing you to reason in those attitudes. Beliefs are special in that they need no linguistic marker. The same sentence – in the example: *Paying taxes is legally required* – expresses the belief’s content and its marked content.

It is debatable how exactly the English language expresses reasoning, i.e., which linguistic constructs serve to mark attitude-contents. Reasoning in preferences might at first seem obscure, as preferences are dyadic attitudes. Broome (2006) however points out (citing Jonathan Dancy for this insight) that English has a preference marker, namely a construction with ‘rather’. You can reason in preferences as follows:

*Rather bike than walk. Rather walk than drive. So, rather bike than drive.* (2)

You initially prefer biking to walking, and walking to driving. You come to prefer biking to driving.

You do not say to yourself *that you hold* the attitudes in question. For instance, in (1) you do not say:

*I believe paying taxes is legally required. So, I intend to pay taxes.*

This would be reasoning *about* your attitudes (cf. Section 4.2).

Reasoning is rule-governed: you draw the conclusion by following a *rule* that you endorse, although this endorsement is not an explicit act and requires no awareness of the rule, and indeed of the concept of rule. A rule allows forming a (conclusion-)attitude from existing (premise-)attitudes. But what is a rule more exactly?

Rules can be individuated more or less specifically. For now we focus on deterministic rules, i.e., rules in Broome’s sense. In its most specific individuation, the rule you follow in (1) is this: from believing that paying taxes is legally required, come to intend to pay taxes. In a broader individuation, the rule is a schema, such as: from believing that  $\phi$ -ing is legally required, come to intend to  $\phi$  (where  $\phi$  is any act). Many rules promote your rationality. Here are examples of rationality-promoting rules, stated informally:

- (a) *Modus-Ponens Rule*: From believing  $p$  and believing *if  $p$  then  $q$* , come to believe  $q$ . Parameters: propositions  $p, q$ .
- (b) *Enkratic Rule*: From believing *obligatorily  $p$* , come to intend  $p$ . Parameter: propositions  $p$ .
- (c) *Instrumental-Rationality Rule*: From intending  $p$  and believing  *$q$  is a means implied by  $p$* , come to intend  $q$ . Parameters: propositions  $p, q$ .
- (d) *Preference-Transitivity Rule*: from preferring  $p$  to  $q$  and preferring  $q$  to  $r$ , come to prefer  $p$  to  $r$ . Parameters: propositions  $p, q, r$ .

For instance, your reasoning in preferences (2) follows the Preference-Transitivity Rule, where  $p, q$  and  $r$  are *I bike, I walk* and *I drive*, respectively. The rules (a)–(d) are just examples; one could modify these rules (e.g., by reformulating or adding premises) or state entirely other rules. Exactly which rules you follow or should follow is not our topic.

### 2.3 Reasoning formalised – deterministic case

We now formalise Broomean reasoning following Dietrich et al. (2018), so far ignoring choice in reasoning. Our model individuates rules specifically rather than broadly. This modelling choice simplifies the formalism; nothing hinges on it. So, a (deterministic) reasoning rule is simply a combination  $(P, k)$  of a specific set of (premise-)attitudes  $P \subseteq M$  and a specific (conclusion-)attitude  $k \in M$ . The four rule schemas (a)–(d) in Section 2.2 can now be re-stated more formally:

- the rules  $(P, k) = (\{\textit{believing } p, \textit{believing if } p \textit{ then } q\}, \textit{believing } q)$  for propositions  $p, q$ ,
- the rules  $(P, k) = (\{\textit{believing obligatorily } p\}, \textit{intending } p)$  for propositions  $p$ ,
- etc. for (c) and (d).

These re-statements of rules stay semi-informal; but formal statements are possible.<sup>9</sup>

A rule  $(P, k)$  transforms your initial constitution  $C$ . If you possess all premise-attitudes in  $P$  but not the conclusion-attitude  $k$ , then you form the attitude  $k$ , acquiring the constitution  $C \cup \{k\}$ . Otherwise your constitution remains the same. Formally:

**Definition 2** A *deterministic reasoning rule* is a pair  $(P, k)$  of a set of premise-attitudes  $P \subseteq M$  and a conclusion-attitude  $k \in M$ . The *revision* of a constitution  $C$  by this rule is the constitution

$$C^+ = \begin{cases} C \cup \{k\} & \text{if } P \subseteq C \text{ \& } k \notin C \text{ (the rule 'applies')} \\ C & \text{otherwise.} \end{cases}$$

You can reason with certain rules – ‘your’ rules. The set of your rules is your *reasoning system*  $S$ . A constitution is *reachable* for you if you can acquire it by reasoning with some of your rules. Formally:

**Definition 3** A *deterministic reasoning system* is a set  $S$  of deterministic reasoning rules (‘your’ rules). Given an initial constitution  $C$ , a constitution  $C^+$  is **reachable from  $C$  by  $(S-)$ reasoning** if there are a number of steps  $n \geq 0$  and constitutions  $C_0, \dots, C_n$  such that  $C_0 = C$ ,  $C_n = C^+$ , and each  $C_t$  with  $t \neq 0$  is the revision of  $C_{t-1}$  by some rule in  $S$ .

## 3 Reasoning in attitudes – indeterministic case

You often face choices in reasoning. Broome (2013) discusses such choice, but develops no full account. We now present a philosophical analysis of indeterministic reasoning in

<sup>9</sup>Using the formalism in footnote 8, write  $(P, k) = (\{(p, \textit{bel}), (\textit{if } p \textit{ then } q, \textit{bel})\}, (q, \textit{bel}))$  ( $p, q \in L$ ) for (a); write  $(P, k) = (\{(\textit{obligatorily } p, \textit{bel})\}, (p, \textit{int}))$  ( $p \in L$ ) for (b); etc. This involves composite propositions. To given them formal meaning, assume that to any propositions  $p, q \in L$  is assigned a proposition *if p then q* in  $L$ ; that to any proposition  $p \in L$  is assigned a proposition *obligatorily p* in  $L$ ; etc. Technically, this defines a binary operator  $L \times L \rightarrow L$ ; a unary operator  $L \rightarrow L$ ; etc. The rules (a)–(d) are now formally specified. One could go further and model propositions in  $L$  syntactically (intensionally) as sentences in a formal language, or semantically (extensionally) as subsets of some set of possible worlds. This turns operators into syntactic or semantic operators, respectively (cf. Dietrich et al. 2022).

attitudes (Sections 3–3.4) and a formal model thereof (Sections 3.5–3.6). Philosophically oriented readers can focus mainly on Sections 3.1–3.4, technically oriented readers mainly on Sections 3.5–3.6. An appendix provides complementary discussion and formal material.

### 3.1 The problem of choice in reasoning

Consider a version of an example in Broome (2013: 163 ff.). You intend to visit Venice; you believe that taking a boat and taking a train are the two only possible means; and you believe that both means are equally good. From these three premise-attitudes you can reason either to a ‘boat’ intention or to a ‘train’ intention. You face choice in reasoning. Facing choice happens frequently in instrumental reasoning – not just when you believe that the available means are equally good, but also when you are unable to compare the means, and more generally when your attitudes create a ‘tie’.<sup>10</sup>

There are two fundamental types of choice in reasoning, to be called ‘choice between’ and ‘choice whether’. The former is a choice between rival conclusions, as in our example. The latter is a choice whether to derive any conclusion at all. You can face a ‘choice whether’ without a ‘choice between’ (as there can be just one conclusion, which is optional), and vice versa (as there can be many possible conclusions, abstention being impermissible). Broome’s account of reasoning (Section 2.2) already allows for ‘choice whether’, in that it does not imply a *requirement* to reason with a given rule – or so Broome would argue. We shall thus focus on ‘choice between’ (but return to this issue in the conclusion). ‘Choice’ and ‘indeterminism’ will therefore refer to ‘choice between’.

The question of how you reason when facing choice between rival conclusion-attitudes contains two subquestions:

- (i) How do you reason to any conclusion-attitude at all?
- (ii) How do you avoid reasoning to several conclusion-attitudes consecutively?

Both questions are non-trivial. In (i), the possibility to reach a conclusion *by reasoning* is prima facie threatened by the lack of reasons for favouring any conclusion over another. In (ii), the difficulty is that, prima facie, *if* you can reach some conclusion *then* you can by symmetry reach another conclusion too, assuming that your premise-attitudes are all you need to reach a conclusion.

Broome’s (2013) take on question (i) is strikingly simple: you can reason to a conclusion in the same way in which you do it when you face no choice, i.e., when there is no alternative conclusion. Formally speaking, if in the Venice example your initial constitution  $C$  includes the set  $P$  of premise-attitudes, and if your reasoning system contains the deterministic reasoning rules  $(P, k_1)$  and  $(P, k_2)$ , where  $k_1$  and  $k_2$  are the possible conclusion-attitudes, i.e., the ‘boat’ and ‘train’ intention, then you can form either conclusion by standard deterministic reasoning. We shall adopt this approach to question (i), but indicate an alternative in the concluding remarks.

We shall however grapple with question (ii). The trouble is that nothing in Broome’s

---

<sup>10</sup>In the Venice example we could have replaced your equal-goodness belief (the third premise-attitude) by another tie-creating attitude, e.g., a preferential indifference between the means, or a belief of being unable to compare the means in terms of betterness or equal goodness.

account stops you from forming ‘surplus conclusions’, as Dietrich et al. (2019) note. You can form both intentions by applying both rules  $(P, k_1)$  and  $(P, k_2)$  consecutively; and this *will* happen if you reason all the way, applying all your rules. What mechanism stops this counterintuitive ‘surplus reasoning’?

### 3.2 Four accounts of indeterministic reasoning

We now present four accounts of how you make *one and only one* choice in reasoning. Technically, a situation of choice in reasoning is characterised by a set of premise-attitudes  $P \subseteq M$  (in the Venice example: the intention to visit Venice and two beliefs) and a set of at least two possible conclusion-attitudes  $K \subseteq M$  (in the Venice example: the ‘boat’ and ‘train’ intention). Each account of choice in reasoning gives *some* explanation of how you succeed in reasoning from  $P$  to one and only one attitude in  $K$ .

We use the Venice example to illustrate each account. So, we must explain how you reach the ‘boat’ or ‘train’ intention, but not both. In principle, the four accounts generalise to other examples, whose conclusion-attitudes can be of another type than intentions and of another number than two. Still the Venice example is paradigmatic, as it addresses action-guiding reasoning and parallels the infamous problem of Buridan’s ass.

**First account.** Here, you reason to the broad intention to take either a boat or take a train, and later an automatic psychological process refines your intention, creating either specific intention, for instance the ‘boat’ intention. This might be your inner speech (with comments added in square brackets):

*I shall visit Venice. For this I must either take a boat or take a train.*  
*Both are equally good. So, I shall take a boat or a train.* (3)  
 [An automatic process intervenes.] *I shall take a boat.*

The final sentence expressing your ‘boat’ intention does not start with ‘So’, as it is not derived by reasoning. In a variant of (3), you do not verbalise your ‘boat’ intention; your inner speech act stops after deriving the broad intention. This makes sense because the ‘boat’ intention is not created actively, but automatically and perhaps subconsciously.

The first account avoids surplus intentions by the (plausible) assumption that your automatic psychology creates only one specific intention.

Clearly, such a two-stage process can happen. You routinely reason to broad intentions that get refined automatically. The problem with the account is not that it is implausible, but that it is not actually an account of choice *in reasoning*. It is an account of deterministic reasoning to a broad intention, followed by automatic ‘choice’, which is not a real choice, as it is not an action. The next three accounts are real accounts of choice *by reasoning*.

**Second account.** Here we presuppose a wider notion of reasoning: you can reason not just from attitudes, but also from absences of attitudes. You reason to the ‘boat’ intention or the ‘train’ intention based on your three attitudes *and* on the additional premise of not having either intention (or at least not having the other intention). This blocks

surplus reasoning because after reasoning to one intention – say the ‘boat’ intention – this intention is no longer absent, so that the fourth premise no longer holds. I never form the broad intention.

Yet reasoning from absences is not explicit reasoning, violating the basics of Broome’s account. You can say to yourself *that* some attitudes are absent, but this expresses a belief in the absence, not the absence itself. It reports rather than expresses the absence. Your inner speech on this account might be this (note that the ‘premise of absence’ cannot be expressed explicitly):

*I shall visit Venice. For this I must either take a boat or take a train.*  
*Both are equally good.* [The absence of the ‘boat’ and the ‘train’ intention intervenes implicitly.] *So, I shall take a boat.* (4)

**Third account.** This account might come the closest to Broome’s (2013) short discussion of choice in reasoning. On the account, after reasoning to either intention, say the ‘boat’ intention, you *can but do not* reason to the ‘train’ intention. You *can* reason to the boat intention because your three premise-attitudes remain present and are all you need to derive a ‘train’ intention, unlike under the second account. But you *do not* derive the ‘train’ intention. Broomean reasoning is something you can, but need not do. Your premise-attitudes enable, but do not require forming the premise-attitudes.

This begs a question: how to you avoid reasoning again? In short, your ‘boat’ intention prevents the second reasoning from getting started or completed. From the three premise-attitudes *plus* your ‘boat’ intention you can no longer derive the ‘train’ intention, because the process of deriving the ‘train’ intention from the (now) *four* premises seems wrong to you, hence does not constitute Broomean reasoning.

One might object that, despite having the ‘boat’ intention, you might not bring it to mind. You might only bring to mind the three original premise-attitudes, from where you *can* derive the ‘train’ intention. Yet, on the account, this simply never happens: you always bring your ‘boat’ intention to mind, if it is not yet conscious. More precisely:

- (i) If your ‘boat’ intention is already conscious, perhaps because you have just acted on it by booking a boat ticket, then you usually do not even start reasoning to another intention. Even if you did start, i.e., brought to mind the three premise-attitudes, you could not complete your reasoning. You could not derive the ‘train’ intention from the *four* attitudes you now have in mind, including the ‘boat’ intention.
- (ii) If your ‘boat’ intention is initially subconscious (‘inactive’), perhaps because you formed it too long ago, then you might begin to reason. You bring to mind premises. But then, rather than rushing to the conclusion of a ‘train’ intention (by Broomean reasoning), you remember your ‘boat’ intention, i.e., bring it to mind (‘activate’ it). You remember it because your activity of reasoning makes it salient. You remember it after bringing all three premise-attitudes to mind, or already after bringing some of them to mind. Then you stop reasoning, again because you cannot derive a ‘train’ intention from the attitudes you have in mind, including your ‘boat’ intention.

Your inner speech act might be as follows, assuming case (ii) and assuming you remember your ‘boat’ intention after bringing to mind two premises:

*I shall visit Venice. For this I must either take a boat or take a train.  
Both are equally good. So, I shall take a boat.*  
[Break. Intentions and beliefs become unconscious. You start again:] (5)  
*I shall visit Venice. For this I must either take a boat or take a train.  
Oh, I already intend to take a boat.* [Reasoning stops.]

In another version of the third account, specifically of clause (ii), you remember your ‘boat’ intention because it *comes* to mind *automatically*, not because you *bring* it to mind *actively*. The effect is the same: you no longer derive the ‘train’ intention.

**Fourth account.** Perhaps the third account is satisfactory. But arguably a problem remains: reasoning from the three premises to intending a means is arguably *incorrect* reasoning, against what Broome thinks.<sup>11</sup> To derive the intention *correctly*, you need an extra premise-attitude, for instance the belief that you do not (yet) intend either means, or the belief that you do not (yet) intend the other means, or perhaps something else. These claims about correctness assume Broome’s own notion of correct reasoning. They are discussed in Appendix B; here we simply accept them.

The fourth account adds the missing fourth premise-attitude. For concreteness, let this attitude be the belief that you do not yet intend either means (some other versions would also work, *mutatis mutandis*). This additional belief is introspective or ‘second-order’: it is a belief about your attitudes. This premise-belief often remains implicit, and possibly subconscious. But even when it is implicit or subconscious, it needs to be present. Sometimes this belief is initially absent, and is formed while reasoning to an intention. It may be formed actively, as discussed in Section 3.4. The speech act by which you form the ‘boat’ intention might be this:

*I shall visit Venice. For this I must either take a boat or take a train.  
Both means are equally good.* [Your meta-belief is formed, actively  
or automatically.] *I do not yet intend a means. So, I shall take a boat.*

Appendix A presents a more complete version of your inner speech in case your meta-belief is formed actively.

Curiously, your reasoning makes one of its premises false. Indeed, the premise-belief of not intending either means becomes false by forming either intention, say the ‘boat’ intention. Being false, the belief might then disappear automatically. This prevents you from reasoning to the ‘train’ intention, as a premise has gone away. But, even if your introspective belief failed to disappear after becoming false, then something else saves you from starting or from finishing to reason to a ‘train’ intention: it is your ‘boat’ intention,

---

<sup>11</sup>More precisely, Broome regards such reasoning as correct provided one adds a fourth premise-attitude, namely the belief that both means are up to you. If the means are not up to you, say because you have delegated your travel planning to a travel agent, then your intention is unnecessary. Broome may be right that correctness requires his fourth premise. We ignore this premise for expositional simplicity.

which is either already conscious, or is brought to mind, or appears automatically. The precise mechanism resembles what we encountered in the third account.<sup>12</sup>

### 3.3 Indeterministic reasoning as a mental activity

Indeterministic reasoning is not standard Broomean reasoning, because it is an interplay between some process generating a conclusion-attitude and some mechanism blocking a second conclusion-attitude; details depend on the account. Under some accounts, indeterministic reasoning departs strongly from Broomean reasoning, by including automatic processes or reasoning from absences.

But on some versions of the third and fourth account, indeterministic reasoning is a broadly Broomean process: it is an activity, i.e., something you *do*, as opposed to something happening partly automatically (in fact, it even is an activity that can be explicit). We now explain why, using the Venice example. We must ask two questions:

- First, is a conclusion-attitude – the ‘boat’ or ‘train’ intention – formed *actively*? On the third account, you derive the intention by Broomean reasoning, hence by an activity. On the fourth account, the same is true, but with an extra caveat: one premise-attitude, namely the belief of not yet holding either intention, must first be formed (if you do not already have it). This belief appears either automatically, or actively by a mental act of *introspecting*. In the second case, all steps are mental acts, of introspecting or reasoning. We discuss introspection in Section 3.4.
- Second, is an additional conclusion-attitude avoided *actively*? On the third account, surplus reasoning gets blocked because you remember your existing conclusion. Remembering is a mental act of bringing the intention to mind, or an automatic process. In the first case, surplus reasoning is blocked by a mental activity. On the fourth account, surplus reasoning gets blocked either because the already derived intention lets you automatically lose your (by now false) introspective belief of holding neither intention; or because your ‘boat’ intention is conscious; or because you bring it to mind; or because it comes to mind automatically. In the second or third case, surplus reasoning is blocked by a mental activity.

### 3.4 The mental activity of introspecting

What is introspecting? This question matters to the fourth account of reasoning, but matters also in itself. Introspecting is a mental act that generates a belief about your attitudes, for instance the belief of (not) having some attitudes. Introspective beliefs often appear automatically; but here we focus on the mental act of introspecting, by which you create introspective beliefs actively. Introspecting resembles observing (and

---

<sup>12</sup>While you can derive the ‘train’ intention from the four premise-attitudes (including the introspective belief) you cannot derive it after adding the ‘boat’ intention as a fifth premise-attitude. Such reasoning would rest on a strange combination of premise-attitudes, as the ‘boat’ intention (the fifth premise-attitude) clashes with the belief of not intending either means (fourth premise-attitude). Like under (i)–(ii) in the third account, you either do not even start reasoning to the ‘train’ intention, because your ‘boat’ intention is conscious. Or you start reasoning, but then, after bringing to mind some or all of the premise-attitudes, you remember your ‘boat’ intention, which lets you stop reasoning as you cannot derive the ‘train’ intention with your ‘boat’ intention as an additional premise-attitude.

thereby coming to believe) a fact, except that the fact in question is one about your mind. Introspecting is thus a purely mind-internal form of observing, which requires no external perception or sensory data.

Introspecting is a different mental act from Broomean reasoning, as it is not inferential. You do not *infer* something about your attitudes; you *observe* it. For instance, introspectively discovering your ‘boat’ intention differs from deriving the belief of having this intention. The categorical difference between introspecting and Broomean reasoning becomes even clearer when what you observe is the absence of attitudes. A belief of not having certain attitudes cannot be inferred from attitudes, but only reached by introspecting. Accordingly, in the fourth account of choice in reasoning, you discover your lack of an intention by introspecting, not by reasoning.

While Broomean reasoning starts from premise-attitudes, introspecting arguably starts from a special attitude, namely one of ‘wondering whether’, which might be analysed as ‘intending to know whether’. You wonder whether  $p$  (where  $p$  is a given proposition about your attitudes), which makes you check whether  $p$ , and ultimately believe  $p$  or believe *not*  $p$ . ‘Wondering’ could have denoted an activity; but we understood it as an attitude. This attitude triggers an activity, namely introspecting.

Just as you can reason explicitly, you can introspect explicitly, using the interrogative mood to express your wondering, as explained in Appendix A. So your indeterministic reasoning under the fourth account can as a whole be an explicit activity, one component being explicit introspection.

### 3.5 Indeterministic reasoning formalised

A full psychological model of reasoning with choices would inevitably look different for the four accounts. We take a different strategy: we only model what is effectively achieved, namely the formation of exactly one conclusion-attitude. Our ‘shorthand model’ thus focuses on the result of reasoning, staying silent on the psychological process. This makes the model account-neutral: it encompasses all four accounts, rather than committing to one.

The model generalises our deterministic model of Section 2.3. It first generalises deterministic rules (Definition 2) into indeterministic rules:

**Definition 4** An *indeterministic reasoning rule* is a pair  $(P, K)$  of a set of premise-attitudes  $P \subseteq M$  and non-empty set of possible conclusion-attitudes  $K \subseteq M$ . A (**possible**) *revision* of a constitution  $C$  by this rule is a constitution given by

$$C^+ = \begin{cases} C \cup \{k\} \text{ for some } k \in K & \text{if } P \subseteq C \text{ \& } K \cap C = \emptyset \text{ (the rule ‘applies’)} \\ C & \text{otherwise.} \end{cases}$$

An indeterministic rule with a single conclusion-attitude  $(P, K) = (P, \{k\})$  is effectively deterministic, and can be identified with the corresponding deterministic rule  $(P, k)$ . For there is a single revision of a constitution  $C$  by  $(P, \{k\})$ , namely the revision of  $C$  by the deterministic rule  $(P, k)$ .

As in the deterministic model, you have many reasoning rules at your disposal –

‘your’ rules – and you reason with them to reach new constitutions. This leads to an indeterministic generalisation of Definition 3:

**Definition 5** *An **indeterministic reasoning system** is a set  $S$  of indeterministic reasoning rules (‘your’ rules). Given an initial constitution  $C$ , a constitution  $C^+$  is **reachable from  $C$  by (S-)reasoning** if there are a number of steps  $n \geq 0$  and constitutions  $C_0, \dots, C_n$  such that  $C_0 = C$ ,  $C_n = C^+$ , and each  $C_t$  with  $t \neq 0$  is a revision of  $C_{t-1}$  by some rule in  $S$ .*

An indeterministic reasoning system  $S$  containing only deterministic rules  $(P, K) = (P, \{k\})$  can be identified with a deterministic reasoning system, by identifying each rule  $(P, \{k\})$  in  $S$  with  $(P, k)$ .

An indeterministic rule  $(P, K)$  is a metaphor or shorthand for a complex psychological process, whose nature depends on the account. None of the four accounts *literally* involves an indeterministic reasoning rule. The first, third and fourth account involve ordinary deterministic rules, and the second account involves non-ordinary deterministic rules (with absences of attitudes as additional premises). A full psychological model would therefore look rather different from our shorthand model. While the shorthand model ascribes to you an indeterministic rule  $(P, K)$ , a full psychological model on one of the accounts invokes for each possible conclusion-attitude  $k$  in  $K$  a deterministic rule that derives  $k$  based on certain premises.<sup>13</sup> Appendix C sketches a full psychological model on each account.

### 3.6 The structural properties of reasoning in attitudes

Reasoning in attitudes, using a given reasoning system  $S$ , defines a *consequence operator*, to be called a ‘Broomean’ operator. In general, a consequence operator over the set  $M$  transforms old into new constitutions (subsets of  $M$ ). Consequence operators are familiar objects in logic ever since Tarski (1956), except that they usually operate over sentences or propositions, not attitudes – but this interpretive difference does not affect the abstract analysis.

This section defines Broomean consequence operators formally, and investigates their structural properties. As we shall see, they are in one sense entirely standard: they are monotonic, inclusive, and idempotent. In another sense, they are non-standard: they are indeterministic, i.e., multi-valued, provided some rules in  $S$  are indeterministic.

A **deterministic consequence operator (over  $M$ )** is a function transforming each constitution  $C \subseteq M$  into a new constitution  $C^+ \subseteq M$ , the *consequence or revision* of  $C$ . More generally, an **indeterministic consequence operator (over  $M$ )** is a function transforming each constitution  $C \subseteq M$  into a set of constitutions, the (*possible*) *consequences or revisions* of  $C$ . If the operator transforms each constitution  $C$  into a singleton set  $\{C^+\}$ , i.e., if each  $C$  has a single revision, then we treat the operator as *deterministic*, by identifying an output set  $\{C^+\}$  with its single element  $C^+$ .

<sup>13</sup>On the first and third account the premise-attitudes are the attitudes in  $P$ ; on the second account they are the attitudes in  $P$  plus the absences of the attitudes in  $K$ ; on the fourth account, they are the attitudes in  $P$  plus the belief  $m$  of not yet possessing attitudes in  $K$ .

Your reasoning system  $S$  – whether deterministic or indeterministic – yields a consequence operator, called ‘Broomean’. How is it defined? Recall that your initial constitution  $C$  evolves as you reason with rules in  $S$ . A constitution  $C^+$  is reachable from  $C$  by ( $S$ -)reasoning if it emerges after applying any number  $n \geq 0$  of rules from  $S$  in any order (see Definitions 3 and 5). If moreover no rule from  $S$  applies to  $C^+$ , i.e., no further reasoning is possible, then we call  $C^+$  reachable by *maximal* ( $S$ -)reasoning from  $C$ . Maximal reasoning yields the Broomean operator:

**Definition 6** *The **Broomean consequence operator** given a reasoning system  $S$  is denoted  $Cn_S$  and transforms any constitution  $C$  into*

$$Cn_S(C) = \{C^+ : C^+ \text{ is reachable from } C \text{ by maximal } S\text{-reasoning}\}.$$

Defining the Broomean operator by *maximal* reasoning implicitly assumes that all your rules *must* be applied, i.e., that you face no ‘choice whether’ in reasoning, only a ‘choice between’ (see Section 3.1 for this distinction). ‘Choice whether’ is neglected in Definition 6 because it is not our primary focus. But one can generalise the notion of Broomean operators and the below theorem by permitting ‘choice whether’.<sup>14</sup>

The Broomean operator  $Cn_S$  is deterministic if all reasoning rules in  $S$  are deterministic, but usually indeterministic if some rules are indeterministic. Calling the operator ‘Broomean’ is a stretch if  $S$  contains indeterministic rules, as Broome never invokes indeterministic rules. But recall that indeterministic rules are a metaphor for any of the four accounts of choice in reasoning, the third of which seems compatible with Broome (2013). This can hopefully justify our term ‘Broomean’.

Deterministic consequence operators are far more common than indeterministic ones. Here are the three most standard properties of a deterministic operator  $C \mapsto C^+$ :

- *Inclusiveness*: For any initial constitution  $C \subseteq M$ ,  $C \subseteq C^+$  (‘revision never removes attitudes’).
- *Idempotence*: For any initial constitution  $C \subseteq M$ ,  $(C^+)^+ = C^+$  (‘revising again has no effect’).
- *Monotonicity*: For any initial constitutions  $C, D \subseteq M$ , if  $C \subseteq D$  then  $C^+ \subseteq D^+$  (‘the more attitudes you have ex-ante, the more you have ex-post’).

We generalise these properties to an indeterministic consequence operator:

- *Inclusiveness*: For any initial constitution  $C \subseteq M$  and any revision  $C^+$  of  $C$ ,  $C \subseteq C^+$ .
- *Idempotence*: For any initial constitution  $C \subseteq M$  and any revision  $C^+$  of  $C$ , the only revision of  $C^+$  is  $C^+$ .
- *Monotonicity*: For any initial constitutions  $C, D \subseteq M$ , if  $C \subseteq D$  then each revision  $D^+$  of  $D$  includes some revision  $C^+$  of  $C$ .

---

<sup>14</sup>How? Partition  $S$  into a set  $S_1$  of ‘mandatory’ rules (which you *must* apply, leaving no ‘choice whether’) and a set  $S_2$  of ‘optional’ rules (which you *can* apply, leaving a ‘choice whether’).  $S_1$  represents your reasoning obligations,  $S_2$  your reasoning permissions. The generalised Broomean operator  $Cn_{S_1, S_2}$  is definable like  $Cn_S$ , merely replacing ‘maximal  $S$ -reasoning’ with ‘ $S$ -reasoning that is maximal w.r.t.  $S_1$ ’ (suitably defined).  $Cn_{S_1, S_2}$  reduces to  $Cn_S$  if  $S_1 = S$  and  $S_2 = \emptyset$ , i.e., if all rules are mandatory. Theorem 1 generalises *partly* to  $Cn_{S_1, S_2}$ , which remains inclusive and monotonic, but can become non-idempotent.

The triple of indeterministic definitions reduces to the standard triple if the consequence operator is deterministic. The Broomean consequence operator has these properties, and is in this respect classical, albeit indeterministic.

**Theorem 1** *The Broomean consequence operator given by any (deterministic or indeterministic) reasoning system  $S$  is inclusive, idempotent, and monotonic.*

Part of this result is obvious. Broomean consequence is of course inclusive, as reasoning never removes attitudes; and it is idempotent, as no further change is possible after maximal reasoning (but see fn. 14 about adding ‘choice whether’). Monotonicity is less trivial; it is established in the appendix.

## 4 The difficulty to model reasoning in attitudes logically

Given that (deductive) theoretical reasoning follows logical entailment, one might expect that reasoning in attitudes also follows entailment of some sort, and hence could be modelled logically. By a logical model of reasoning we mean a model in terms of some kind of entailment relation.

We now go, largely unsuccessfully, through the three most natural attempts to model reasoning in attitudes logically (Section 4.1–4.3). To rule out trivial deviations from entailment, much of the discussion will implicitly assume that you are an ‘ideal’ reasoner, i.e., your reasoning system  $S$  contains the right kind of rules. We do not spell out what this means exactly.<sup>15</sup>

### 4.1 Content entailment: a model of reasoning in a *single* attitude

When theoretical reasoning is said to follow entailment, one refers to entailments between belief-*contents*. So, our first attempt must be to model reasoning in attitudes by entailments between attitude-*contents*.

Reasoning in attitudes would have no chance to follow entailment between attitude-contents if those contents were not propositional or proposition-like. We would commit a category error by invoking entailments between the contents of attitudes like *adoring Paris* and *imagining a fireball*, i.e., entailments between things like towns and fireballs. The ‘propositionalist’ view that all intentional attitudes have, or can be reduced to having propositional content is controversial (Montagne 2007). While not committing to full-blown propositionalism, we do assume that *reasoning* in attitudes is reasoning in *propositional* attitudes.

So again: does reasoning in (propositional) attitudes follow entailment between the propositions? In the first place, logic is about propositions, not about anyone’s attitudes. But insofar as propositions are the contents of attitudes, logicians can indirectly

---

<sup>15</sup>Intuitively,  $S$  contains only good rules (‘soundness’) and sufficiently many rules (‘completeness’). Examples of good rules might be the deterministic rules in (a)–(d) in Section 2.2 and the indeterministic rule in the Venice example. To be complete,  $S$  might need to contain these and many other rules; for instance, without the rule in the Venice example you are handicapped, i.e., unable to form a ‘boat’ or ‘train’ intention where needed. For deterministic rules, being good might mean being correct in Broome’s sense.

address attitudes. When logicians do so, they notoriously choose *beliefs*: they interpret propositions as *belief*-contents, which turns logical entailment into a model of reasoning in beliefs, not in desires, or in intentions, etc. Could logicians instead choose desires (or intentions, etc.), and take content entailment to model reasoning in desires (or intentions, etc.)? Such a model would support reasoning from desiring  $p$  into desiring  $p$  or  $q$  (or from intending  $p$  into intending  $p$  or  $q$ , etc.), as  $p$  entails  $p$  or  $q$ ; and it would support reasoning from nothing into desiring a tautology (or intending it, etc.), as the empty set entails the tautology. One might doubt such reasoning, and hence reject the idea that reasoning in desires (or in intentions, etc.) follows content entailment.

One reason for doubt is that entailment is deterministic, whereas reasoning in desires (or in intentions, etc.) might involve choices. We see no principled reason to exclude such choices, although convincing examples seem harder to find than when mixing attitude types, as done in the Venice example.

But even if reasoning in desires (or in intentions, etc.) followed content entailment, we would not have modelled general reasoning in attitudes. Reasoning in desires (or in intentions, etc.) is still mono-attitude reasoning. Once we mix attitude types, as practical reasoning routinely does, content entailment obviously cannot model reasoning: while  $p$  and *if  $p$  then  $q$*  entail  $q$ , you would not reason from desiring  $p$  and believing *if  $p$  then  $q$*  into intending  $q$ .

In sum, although content entailment can famously model theoretical reasoning (with qualifications discussed in Section 5), it is debatable as a model of reasoning in a given non-belief attitude such as desire or intention, and fails for reasoning in multi-attitudes.

For even simpler reasons, reasoning in *non-monadic* attitudes cannot follow content entailment, as such attitudes have complex contents. For instance, reasoning in preferences (Broome 2006) is reasoning in attitudes towards *pairs* of propositions. Entailments go between propositions, not between pairs.

## 4.2 Attitude entailment: a model of reasoning *about* attitudes

The mismatch between reasoning in attitudes and content entailment already shows that any relation between reasoning in attitudes and logic – if existent – is very different from that for theoretical reasoning. Such a different relation would exist if reasoning in attitudes followed *attitude* entailment. However, read literally, attitude entailment models theoretical reasoning *about* attitudes, not reasoning *in* attitudes. Why? Assume you reason in attitudes as follows:

$$I \text{ ought to pay taxes. So, } I \text{ shall pay taxes.} \tag{6}$$

Here you reason from a belief to an intention, following an instance of the Enkratic Rule in Section 2. Does this reasoning follow an attitude entailment, namely (in formal-logical terms) the entailment  $B(p) \models I(q)$ , where  $B$  is a belief operator,  $I$  is an intention operator, and  $p$  and  $q$  are sentences representing *I ought to pay taxes* and *I pay taxes*, respectively? Whenever we state attitude entailments in formal logic, we presuppose a suitable logic of attitudes, with modal operators for all relevant attitude-types, e.g., a belief operator, an intention operator, or a (dyadic) preference operator. (Logics of

attitudes exist in abundance. They can do many things.<sup>16</sup>)

Yet the literal reading of the entailment  $B(p) \models I(q)$  is not (6), but:

*I believe I ought to pay taxes. So, I intend to pay taxes.* (7)

Here you reason *about* your attitudes: you deduce you have an intention, from having a belief. In (6) you do not reason about your attitudes: you are not your own observer who recalls having a belief and deduces (‘discovers’) an intention, but you *form* an intention. Attitude entailment models attitude discovery, not attitude formation. It models reasoning about, not in attitudes. Reasoning about attitudes does not change *these* attitudes, but it creates beliefs about them (cf. Broome 2014 and Dietrich et al. 2019).

Worse, the different reasoning (7) which the entailment  $B(p) \models I(q)$  models is invalid as an inference about your attitudes: its premise can hold without its conclusion holding. Indeed, before starting your reasoning (6), you believed your ought to pay taxes without (yet) intending it; formally,  $B(p)$  was true and  $I(q)$  was false. Why, then, does the logic deem the inference  $B(p) \models I(q)$  valid? Nothing is wrong with the logic, but we have misapplied it. The logic is one of *rational* attitudes, unlike your imperfect attitudes. Attitude entailment represents reasoning about *rational* attitudes.

The whole point of reasoning *in* attitudes is to become more rational. Ironically, reasoning *about* your attitudes works if your attitudes are already rational, whereas reasoning *in* attitudes matters if your attitudes are not yet rational.

To be precise, improving rationality need not be the purpose of reasoning in attitudes. You could for instance reason as in Section 2:

*Paying taxes is legally required. So, I shall pay taxes.* (8)

This reasoning from a belief to an intention is not enkratic reasoning, because it starts from a belief about what is *legally* required, not what you ought to do. As Broome might say, you reason towards legality, not rationality. Modelling (8) by an entailment – namely by  $B(p') \models I(q)$  where  $p'$  represents the new premise content – is again problematic, still because the entailment represents a piece of reasoning *about* attitudes.<sup>17</sup> The novelty of the example is that the entailment  $B(p') \models I(q)$  is invalid in a logic of ‘merely’ rational attitudes, as the premise-belief does not *rationally* entail the conclusion-intention. The entailment holds in a logic of ideal attitudes in a suitably comprehensive sense of ‘ideal’ that captures the relevant norms, such as rational, legal, or moral norms.

In sum: attitude entailment in a suitable logic of rational (or otherwise ideal) attitudes does not model reasoning in attitudes, but reasoning about attitudes, more precisely about rational (or otherwise ideal) attitudes rather than your attitudes.

---

<sup>16</sup>Mono-modal logics address one attitude, e.g., belief in ‘doxastic logics’ (e.g., Halpern 2017) and preferences in ‘preference logics’ (e.g., Liu 2011). Multi-modal logics address more than one attitude, e.g., belief, desire and intention in ‘BDI logics’. Logics of attitudes capture rationality of attitudes by axioms (e.g., axioms requiring that tautologies are believed).

<sup>17</sup>The entailment reads: *I believe paying taxes is legally required. So, I intend to pay taxes.*

### 4.3 Attitude entailment: an as-if model of reasoning in attitudes?

Could attitude entailment at least serve as an as-if model of reasoning in attitudes, instead of a literal model? That is, could reasoning in attitudes *effectively* change your attitudes in a way that mimics attitude entailment, so that attitude entailment *predicts* the effect of reasoning? This as-if approach is psychologically silent: it does not model the psychological process (the passage from premises to conclusions), but only the resulting attitudinal changes. Reasoning would then behave as if following attitude entailment, hence be extensionally equivalent to attitude entailment. As-if models are popular in rational-choice theory.<sup>18</sup> They remain the orthodoxy among economists, whilst being increasingly criticised (e.g., Dietrich and List 2016, Guala 2019).

Such an as-if model must fail, because reasoning is an indeterministic process whereas attitude entailment is deterministic. We now spell this fact out precisely, by formulating and later rejecting the hypothesis of extensional equivalence between reasoning and attitude entailment. This is the equivalence hypothesis on which the as-if model rests:

**Extensional Equivalence Hypothesis (EE) – informal statement:** You reason to an attitude if and only if your initial attitudes entail that attitude.

The hypothesis can be re-stated formally, after formalising reasoning and entailment:

- Your reasoning is based on your reasoning system  $S$ , and formally captured by the Broomean reasoning operator, which transforms each constitution  $C \subseteq M$  into a (usually non-unique) revision, i.e., into any constitution  $C^+$  reachable from  $C$  by maximal  $S$ -reasoning (cf. Section 3.6).
- Attitude entailments are formally given by the entailment relation of a suitable logic of (rational or otherwise ideal) attitudes. Each attitude  $m \in M$  is represented by a sentence saying that you have this attitude; it is denoted  $m^*$  and takes the form  $O(\phi)$  where  $O$  is the relevant attitude operator and  $\phi$  is the relevant sentence. For instance, if  $m$  is *intending to swim*, then  $m^*$  is  $O(\phi)$  where  $O$  is the intention operator and  $\phi$  reads ‘you swim’.<sup>19</sup> An entailment from your initial attitudes (in  $C$ ) towards a new attitude  $k$  is thus formalised as  $\{m^* : m \in C\} \models k^*$ . In words: the attitudes in  $C$  (rationally or otherwise ideally) entail the attitude  $k$ .

**Extensional Equivalence Hypothesis (EE) – formal statement:** For any initial constitution  $C \subseteq M$ , revisions  $C^+$  of  $C$ , and attitude  $k \in M$ ,

$$k \in C^+ \Leftrightarrow \{m^* : m \in C\} \models k^*. \quad (9)$$

The left side of the equivalence (9) says that after reasoning you have the attitude  $k$ . The right side says that  $k$  is entailed by your initial attitudes. In the example (7),  $k$  is

<sup>18</sup>Under an as-if interpretation, a standard rational agent behaves *as if* maximising expected utility, so that utilities and probabilities carry no meaning beyond representing behaviour. Under more literal or mentalist interpretations, utilities and probabilities are psychological constructs capturing values and beliefs.

<sup>19</sup>Presumably, the assignment  $m \mapsto m^*$  defines a bijective correspondence between  $M$  and the set of logical sentences of type  $O(\phi)$  for some attitude operator  $O$ .

*intending to pay taxes*, logically represented as  $I(q)$  ( $= k^*$ ), and your initial constitution  $C$  contains *believing you ought to pay taxes*, logically represented as  $B(p)$ . Here EE says: this intention is among the attitudes formed by reasoning if and only if it is entailed by your initial attitudes, formally,  $k \in C^+ \Leftrightarrow \{m^* : m \in C\} \models I(q)$ .

Such an as-if model of reasoning is psychologically silent. It captures reasoning to  $k$  by an abstract condition ( $'k \in C^+'$ ) that tells only that you effectively form attitude  $k$ , without revealing the procedural fact of how many reasoning steps are needed to derive  $k$  from  $C$  and which initial attitudes from  $C$  enter your reasoning as premise-attitudes; you might reach  $k$  in just one step from just one premise (as in your reasoning in (7)) or in many steps drawing on many initial attitudes. Since the premises of reasoning are left out, your reasoning is not modelled by an entailment from your premise-attitudes, but by an entailment from the totality  $C$  of initial attitudes (formally, by  $\{m^* : m \in C\} \models k^*$ ).

In sum, the as-if model based on EE treats reasoning as a black box that generates new attitudes by entailment from initial attitudes, regardless of the psychological mechanism at work. This procedural blindness reflects the reduced ambition of the as-if approach, which aims to model what reasoning achieves *in effect*, not *how* it achieves it – an approach we took reluctantly after the more substantive and mentalistic attempts had failed.

But EE is not tenable, and with it the as-if model. We first discuss the central objection, based on choice in reasoning and directed against the necessity of entailment for reasoning (direction  $'\Rightarrow'$ ). We then mention an objection against sufficiency (direction  $'\Leftarrow'$ ), and an ad-hoc-ness concern.

*Against necessity.* You can face choice in reasoning – either ‘choice whether’ (reason or not reason?) or ‘choice between’ (reason to which conclusion?). If you then reason to an attitude, that attitude is not entailed by your initial attitudes, as you could have abstained from reasoning or reasoned otherwise.

As usual, let us focus on ‘choice between’. In our Venice example, you reason to a ‘boat’ intention, which is not entailed by your initial attitudes, as you could form the ‘train’ intention instead. (Recall: ‘entails’ means ‘rationally entails’, or more generally ‘ideally entails’ under some sense of ideal attitudes.)

If anything, a modified version of EE holds, whereby entailment governs which attitudes you *necessarily* reason to, i.e., belong to *all* revisions  $C^+$  of your initial constitution  $C$ .<sup>20</sup>

One might defend EE against the ‘choice’ objection by rejecting the very idea of choice in reasoning, instead making two claims: (i) you never reason to a ‘boat’ intention, but to the broad intention to take a boat *or* a train, this intention being later refined by an automatic process; (ii) the broad intention *is* rationally entailed by your initial attitudes. We doubt both claims. Claim (i) follows the first account of choice

---

<sup>20</sup>The modified hypothesis is: You necessarily reason to an attitude if and only if your initial attitudes entail that attitude. Formally, for any initial constitution  $C \subseteq M$  and attitude  $k \in M$ ,

$$k \in \bigcap_{C^+ \in C_{n_S}(C)} C^+ \Leftrightarrow \{m^* : m \in C\} \models k^*.$$

in reasoning, which is actually an account of *deterministic* reasoning followed by automatic conclusion-refinement (cf. Section 3.2). No doubt, you *sometimes* reason to broad intentions that get refined automatically. Yet sometimes you reason directly to a more specific intention, facing a choice. So the ‘never’ in Claim 1 is false.

Claim (ii) initially seems correct. It seems that your intention to visit Venice and your beliefs about possible means rationally entail intending to take a boat or train, following the requirement (\*) ‘intend what you believe is a necessary means to your intended end’. Perhaps Claim (ii) is true. We however see two counter-arguments:

- On one counter-argument, taking a boat or a train is not a means, but a disjunction of means. Means, on this argument, are special things, presumably certain actions or certain causes (in a ‘production’ sense rather than ‘difference-making’ sense). Taking a boat and taking a train are two actions or causes of relevant kind, but their disjunction is not, because it is somehow non-unified. What you need to intend according to (\*) is a necessary *means*, not disjunction of means. So your initial attitudes do not rationally entail intending the disjunction of means. Indeed, you can perfectly reach Venice without ever forming this broad intention, as long as you intend a means, e.g., to take a ‘boat’. You may form the broad intention as an intermediate step towards intending a means, but this step is not required.<sup>21</sup> This argument rests on a sophisticated notion of a means that is tricky to spell out precisely. A difficulty is that any action can be refined. For instance, taking a boat is the disjunction of taking a large or small boat. Why then should taking a boat qualify as a means, while taking a boat or train does not? Presumably because taking a boat is something unified or connected (albeit refinable), while taking a boat or train is something disunified or disconnected, hence not an action or cause of relevant type.<sup>22</sup> The next counter-argument avoids this difficulty, by being liberal about the notion of means.
- On another counter-argument, due essentially to Broome (2013), even if taking a boat or train *is* a means for visiting Venice, more precisely is believed to be a necessary means to that intended end, you are not required to intend it, because it can come about without being intended. Indeed, intending a specific means (e.g., to take a boat) ensures not just this specific means, but also the broad means. This argument effectively amounts to rejecting the rationality requirement (\*). One could replace it with a weaker requirement: ‘if you believe something is a necessary means to your intended end, then intend something that you believe implies this means’. Or, following Broome (2013: 159), one could refine (\*) by adding the condition that (you believe that) intending the means is necessary for the means. Both refinements of (\*) block the undesired conclusion of having to intend to take a boat or train in the Venice example.

Under either counter-argument, your initial attitudes do not rationally entail the

---

<sup>21</sup>In this context, note that you may intend something without intending its implications; e.g., you may intend to take a boat without intending to take a boat or a train.

<sup>22</sup>To draw an analogy with mathematical topology: in the space of logically possible actions, the set of actions of taking some boat is topologically connected, unlike the set of actions of taking some boat or some train.

broad intention. They also do not rationally entail any given specific intention. They rationally require holding *some* specific intention, as a matter of instrumental rationality; cf. Broome’s (2013: 170) ‘generalised instrumental requirement’.

*Against sufficiency.* Sometimes you cannot form an attitude although your attitudes entail it. Perhaps you are akratic, and unable to form an intention  $k$  which rationally follows from your beliefs about what you ought to do. Here your constitution entails  $k$  (formally,  $\{m^* : m \in C\} \models k^*$ ), but you cannot reason to  $k$  (formally,  $k \notin C^+$  for every revision  $C^+$  of  $C$ ). Perhaps you believe that attitude  $k$  makes you happy; this belief (let us assume) rationally entails forming  $k$ , to which you are unable. Perhaps you intend to become wise and believe wisdom requires studying, but you are psychologically unable to intend to study, although this intention is (rationally) entailed.

However, these counterexamples apply to an imperfect reasoner, who is unable to perform some correct reasoning. One can rehabilitate sufficiency by assuming a perfect reasoner who does not suffer from psychological ‘reasoning barriers’. This reasoner’s reasoning system  $S$  not only contains only correct rules, but also contains sufficiently many rules, enabling barrier-free reasoning towards rationality. This perfect reasoner might satisfy the sufficiency claim, and the modified hypothesis in fn. 20.

*Ad-hoc-ness charge:* Attitude entailments are entailments between attitude propositions of the simplest type: propositions saying that you possess a certain attitude, e.g., that you desire  $p$ . Call them *atomic* attitude propositions. There exist many *non-atomic* attitude propositions: that you do *not* desire  $p$ , that you desire  $p$  *and* believe  $q$ , etc. Entailments between non-atomic attitude propositions do not correspond to (Broomean) reasoning in attitudes. For instance, the entailment  $\{B(p) \vee I(q), \neg D(r)\} \models \neg D(s)$  (for operators of belief  $B$ , intention  $I$ , and desire  $D$ ) does not correspond to any reasoning in attitudes, because Broomean reasoning cannot start from or generate disjunctions or absences of attitudes. You can reason *about* absences or disjunctions, but not *in* them (cf. Dietrich et al. 2019). It might seem ad hoc to pick out particular entailments – those between *atomic* attitude propositions – and grant them a perfect correspondence to reasoning, while denying such a correspondence for all other entailments.

## 5 Concluding remarks: reasoning in attitudes vs. theoretical reasoning

Where do we stand? Reasoning in attitudes differs fundamentally from reasoning about attitudes, and more generally from theoretical reasoning. Conceptual differences aside, reasoning in attitudes does not follow entailment between attitude-contents – unlike (deductive) theoretical reasoning. It does not even follow entailment between attitudes – unlike (deductive) reasoning about your attitudes. The main driving force is choice in reasoning. There may be choice between rival conclusions (‘choice between’) and choice whether to derive a conclusion at all (‘choice whether’). We have mainly addressed ‘choice between’. Our focus lied on how you avoid surplus conclusions (question (ii) in Section 3.1), not how you reach any conclusion at all (question (i) in Section 3.1). This

was possible by adopting Broome’s position that standard (Broomean) reasoning already solves the second problem, i.e., allows you to derive a conclusion even when you face a choice. This position is interesting. It is bound to be controversial, as it essentially implies that Buridan’s ass faces no real dilemma, i.e., can reach one of the stacks of hay by the same simple reasoning by which it could reach it in the absence of choice. If one disagrees, then one could argue that reasoning with choice becomes possible by using a *tie-breaking rule* that selects a conclusion from the given set of possible conclusions. There are many possible tie-breaking rules, such as to choose like you did last time you faced an analogous choice. How and whether tie-breaking can enter your reasoning is a problem of its own.<sup>23</sup>

The special status of theoretical reasoning might surprise. After all, such reasoning is simply a special kind of reasoning in attitudes, namely reasoning in beliefs. Why then can (deductive) theoretical reasoning follow entailment? The central point is the absence of choice in theoretical reasoning: you do not choose between different possible conclusions when forming beliefs rather than, say, intentions. For beliefs track an external truth. They aim to match the external world. As the external world obeys logic, so does (deductive) theoretical reasoning.

There are well-known exceptions, where theoretical reasoning arguably departs from content entailment. You might derive *more* beliefs than are entailed, by reasoning inductively. You might derive *fewer* beliefs than are entailed, because subjectively probable (believed) propositions sometimes jointly entail subjectively improbable (disbelieved) propositions, as in the Lottery Paradox (Kyberg 1961, Dietrich and List 2021). We say ‘might’ because our Broomean account of reasoning might escape at least the second phenomenon, as explicit theoretical reasoning might exclude implicit probabilistic considerations.<sup>24</sup> There might be a third (controversial) reason for departure from entailment: non-epistemic motives, such as the pursuit of happiness. It is certainly questionable whether you can form a belief in order to become happy, even if you wanted to. But if it did work, then it would strongly disconnect theoretical reasoning from content entailment, perhaps to the point of introducing choice into theoretical reasoning.

Still, content entailment is a first-order approximation of theoretical reasoning. By contrast, reasoning in general attitudes goes beyond logic. No doubt, it can be studied formally, namely by using indeterministic consequence operators. But such tools are ‘logical’ only in an abstract sense.

---

<sup>23</sup>Under one approach, tie-breaking is an explicit part of reasoning that requires additional premise-beliefs, for instance the belief that such-and-such possible conclusion is selected by such-and-such tie-breaking rule, *plus* the belief that you ought to use this rule. But is this ought-belief true? Were it true, you would not face a choice in the first place, as one might argue. Under another approach, tie-breaking happens automatically. But might such an automatic process undermine the idea of reasoning as a mental act? We cannot settle these and other issues here.

<sup>24</sup>You may reason in *partial beliefs*; and you can do so explicitly, using markers such as ‘probably’. Such reasoning does of course not follow entailment between the contents of partial beliefs. But by ‘reasoning in beliefs’ we mean ‘reasoning in straight beliefs’.

# Appendix

## A Explicit introspection

While Broomean reasoning is a process of bringing to mind premise-attitudes and then creating a new attitude, introspecting is arguably a process of bringing to mind your wondering and then creating a new (meta-)belief. Like Broomean reasoning, introspecting can in principle be done explicitly. How? By saying to yourself the marked content of your wondering and your resulting belief. Using the interrogative mood as a linguistic marker of wondering, you may introspect explicitly as follows in the fourth account of indeterministic reasoning:

*Do I already intend to take a boat or train? I do not yet intend either.*

You say no ‘So’, as you draw no inference, unlike in Broomean reasoning.

This suggests an objection against the fourth account: if indeterministic reasoning is indeed premised on a meta-belief, whose formation by introspection requires wondering (as just claimed), then you cannot reason without wondering in the first place. In the Venice example: if reasoning to a ‘boat’ or ‘train’ intention requires finding out introspectively that you have neither intention, which requires wondering whether you do, then you cannot reason without initially wondering about this. But normally you do not initially wonder about this (why should you?). Technically, if your initial constitution  $C$  does not contain the relevant attitude of wondering, you cannot introspect, hence cannot reason. What could help? You might be lucky and start wondering automatically when needed. But if the fourth account relied on ‘automatic wondering’, we could not uphold our claim that indeterministic reasoning on this account can be a mental activity (cf. Section 3.3).

Surprisingly, however, you can come to wonder about something by an act of Broomean reasoning. From what premises do you derive a wondering? Intuitively, the premises are the justifications or grounds for wondering. In the Venice example, your wondering whether you already have a ‘boat’ or ‘train’ intention might be derived from your intention to visit Venice and your beliefs about possible means. Then the complete process by which you form a ‘boat’ intention (on the fourth account) is made explicit by the following inner speech:

*(1) I shall visit Venice. (2) For this I must either take a boat or take a train. (3) Both means are equally good. (4) Do I already intend a means? (10)  
(5) I do not intend a means. (6) So, I shall take a boat.*

In (1)–(3) you bring to mind attitudes; from (1)–(3) you derive the wondering (4) by reasoning; from (4) you reach (5) by introspecting; from (1)–(3) and (5) you derive (6) by reasoning.<sup>25</sup>

---

<sup>25</sup>On an alternative view, your wondering is derived from *no* premises, as wondering needs no grounds. Then your complete inner speech might again take the form (10), but this time your wondering in (4) is derived from no premises rather than from (1)–(3).

## B Correct indeterministic reasoning?

According to Broome (2013), you can reason *correctly* to one of several possible conclusions. That is, given suitable sets  $P$  of premise-attitudes and  $K$  of possible conclusion-attitudes, you can reason correctly from  $P$  to any attitude in  $K$ . Broome focuses on instrumental reasoning, in which  $K$  contains intentions of a means to an end; for instance, in our Venice example  $K$  contains the ‘boat’ and ‘train’ intention, and  $P$  contains the intention to visit Venice and certain beliefs (about which Broome is more sophisticated than us, as mentioned in fn. 11). We have questioned Broome’s correctness claim, suggesting that correctness requires adding a premise-attitude to  $P$ , namely the belief of not yet having any attitude from  $K$ .

We here discuss the issue. For argument’s sake, we adopt Broome’s general characterisation of correctness: reasoning from a set of premise-attitudes  $P$  to a conclusion-attitude  $k$  is *correct* if  $P$  is a rationally permissible basis of  $k$ , more precisely, if it rationality permits you to (i) hold the attitudes  $p$  in  $P$  at some moments and (ii) hold the attitude  $k$  at some moment *based on* the former attitudes. This so-called ‘basing permission’ is diachronic: it relates your attitudes at different moments. Indeed, reasoning takes time. By the time you reach the conclusion, you may have lost some premise-attitude(s), despite the conclusion-attitude being based on the premise-attitudes. Broome does not analyse what ‘based on’ means; nor shall we.

Broome applies this correctness criterion to certain examples of reasoning with different possible conclusions, similar to the Venice example. He claims that reasoning to any possible conclusion passes the correctness test. Technically, if  $P$  and  $K$  are the relevant attitude sets, say those in the Venice example, then for any  $k$  in  $K$  the rule  $(P, k)$  passes the correctness test, i.e., in short,  $P$  is a permissible basis of  $k$ . We doubt this claim. Were it true, you could reason correctly to all attitudes in  $K$  one by one, thereby forming conflicting intentions in the Venice example. Arguably, Broome has underspecified the basis of a  $k$  in  $K$ : a permissible basis is not  $P$ , but  $P \cup \{m\}$ , where  $m$  is the belief of not (yet) having any attitude from  $K$ . One could replace  $m$  by other introspective beliefs.<sup>26</sup> It is hard to say whether one could replace  $m$  by a non-introspective (‘first-order’) attitude.<sup>27</sup>

After deriving the conclusion-attitude  $k$  from  $P \cup \{m\}$ , your introspective belief  $m$  is false. So rationality does not permit holding the conclusion-attitude and all premise-attitudes in  $P \cup \{m\}$  *simultaneously*. Yet rationality permits holding these attitudes at different times. This is why the reasoning rule  $(P \cup \{m\}, k)$  can meet the correctness test.

---

<sup>26</sup>An alternative basis of  $k$  might be the set  $P \cup \{m_l : l \in K \setminus \{k\}\}$ , where  $m_l$  is the belief of not possessing attitude  $l$ .

<sup>27</sup>Is  $P \cup \{m'\}$  a permissible basis of  $k$  if  $m'$  is, say, the intention not to take a boat or train to Venice, or the intention to travel to India (which prevents taking a train or boat to Venice)? This hinges on the notion of basis, which Broome and we find hard to spell out. Perhaps  $P \cup \{m'\}$  is a permissible basis. If not, perhaps it becomes one after adding the belief that the attitude  $m'$  implies not intending to take a boat or train to Venice. But a so-increased basis would again contain an introspective belief.

## C Full psychological models of the four accounts of choice in reasoning

The main text worked with a shorthand model of choice in reasoning, which is account-neutral thanks to focusing on the ultimate effect of reasoning on attitudes rather than the psychological process. This appendix sketches how a full psychological model might roughly look under each account. In fact, we only discuss how to model a particular instance of choice in reasoning: reasoning from a set of premise-attitudes  $P \subseteq M$  to any attitude from a set of possible conclusion-attitudes  $K \subseteq M$  (in our Venice example,  $P$  contains the intention to visit Venice and two beliefs, and  $K$  contains the ‘boat’ intention and the ‘train’ intention). The shorthand model represents this instance by the indeterministic rule  $(P, K)$ . A full model might instead take the following form.<sup>28</sup>

**First account.** Here, a full model of reasoning from  $P$  to any attitude in  $K$  has two ingredients. One is a deterministic rule  $(P, k^\vee)$  where  $k^\vee$  is a suitable ‘disjunctive’ attitude (in the Venice example: the broad ‘boat or train’ intention). The second ingredient represents the automatic process that refines your disjunctive attitude  $k^\vee$  into an arbitrary attitude  $k$  in  $K$  (in the Venice example: into the ‘boat’ intention or ‘train’ intention). The composition of your reasoning rule  $(P, k^\vee)$  and the automatic process is effectively equivalent to the indeterministic rule  $(P, K)$  in our shorthand model.

**Second account.** Here, a full model of reasoning from  $P$  to any attitude in  $K$  involves, for each possible conclusion attitude  $k \in K$ , a rule that derives  $k$  if all attitudes in  $P$  are present *and* all attitudes in  $K$  are absent; denote this rule by  $(P, K, k)$ . These rules are generalised deterministic rules, premised on presences *and absences* of attitudes.<sup>29</sup> Reasoning with the rules  $(P, K, k)$  ( $k \in K$ ) is effectively equivalent to reasoning with the single indeterministic rule  $(P, K)$ . Why? For any initial constitution  $C$ , either  $K \cap C \neq \emptyset$ , in which case none of these rules applies and the constitution stays  $C$ ; or  $K \cap C = \emptyset$ , in which case any rule  $(P, K, k)$  applies and leads to the new constitution  $C \cup \{k\}$ , after which none of the rules applies anymore, so that the constitution stays  $C \cup \{k\}$ . The result is the same as for reasoning with the indeterministic rule  $(P, K)$ .

**Third account.** Here, a full model of reasoning from  $P$  to any attitude in  $K$  involves, for each possible conclusion attitude  $k \in K$ , the standard deterministic rule  $(P, k)$ , which forms attitude  $k$  if you have all attitudes in  $P$ . The model also contains a precondition for applying these rules: each of these rules can only be applied to constitutions not yet containing any attitude from  $K$ . This precondition operationalises the assumption that when you start reasoning from the premise-attitudes in  $P$  *but already possess* an

<sup>28</sup>After knowing how to model an instance of indeterministic reasoning, one can easily model indeterministic reasoning in general. For the shorthand model, this step was taken by introducing a reasoning system  $S$ , and defining how constitutions can change by reasoning with  $S$  (Definition 5). For our four full models, one could proceed analogously, by introducing a construct analogous to a reasoning system and defining how constitutions can change by reasoning with this construct.

<sup>29</sup>A rule of this sort is definable as a triple  $(P_1, P_2, k)$ , and adds the attitude  $k$  to a constitution  $C$  if  $P_1 \subseteq C$  and  $P_2 \cap C = \emptyset$ .

attitude from  $K$ , then (on the account) your reasoning stops, caused by your having in mind or bringing to mind a preexisting attitude from  $K$ . The precondition prevents the rules from operating the usual way: they effectively operate like the generalised deterministic rules from our model of the second account (i.e., rules premised on the absence of attitudes from  $K$ ). This is why reasoning in the current model is effectively equivalent to reasoning in the second model, and hence to reasoning in the shorthand model based on the indeterministic rule  $(P, K)$ .

**Fourth account.** Here, a full model of reasoning from  $P$  to any attitude in  $K$  involves, for each possible conclusion attitude  $k \in K$ , the standard deterministic rule  $(P \cup \{m\}, k)$ , which forms attitude  $k$  from the attitudes in  $P \cup \{m\}$ , where the additional premise-attitude  $m$  is the belief of having no attitude from  $K$  (in the Venice example: the belief of having no ‘boat’ intention and no ‘train’ intention). We must also model the mechanism that prevents repeated reasoning to different attitudes from  $K$  (we shall only model the first such mechanism envisaged by the fourth account). To do this, we prescribe that each rule  $(P \cup \{m\}, k)$  ( $k \in K$ ) applies in a non-standard way to a constitution  $C$ : systematically, *before application* the attitude  $m$  is added to your constitution, and *after application*  $m$  is removed again. Interpretation: while reasoning you get aware of not having any attitude from  $K$ , i.e., form the belief  $m$ , and after reasoning you lose that belief as it has become false through forming an attitude from  $K$ . More precisely,  $m$  is only added (and later removed) if  $K \cap C = \emptyset$  and  $P \subseteq C$ . Why only then? If  $K \cap C \neq \emptyset$  then the belief  $m$  is false, while if  $P \not\subseteq C$  then you do not have all attitudes in  $P$ , hence stop reasoning prematurely. In both cases, you never get to the point of forming (and later losing) the belief  $m$ . In sum, the rule  $(P \cup \{m\}, k)$  applies in the following non-standard way to any constitution  $C$ . If  $K \cap C = \emptyset$  and  $P \subseteq C$ , then  $C$  is first transformed into  $C \cup \{m\}$ , which is then transformed by the rule into  $C \cup \{m, k\}$ , which is then transformed into  $(C \cup \{k\}) \setminus \{m\}$ .<sup>30</sup> Otherwise,  $C$  is not transformed.

This model is hardly parsimonious – a drawback of modelling the full psychological process postulated by the fourth account. Reasoning on the fourth model is effectively equivalent to reasoning on the other models or the shorthand model, because the non-standard rules of the fourth model produce the same result as the non-standard rules of the second account, and hence as the indeterministic rule of the shorthand model. To be precise, this effective equivalence holds with respect to all attitudes *except* the introspective belief  $m$ . Indeed, reasoning in the other models never affects the presence of  $m$ , whereas reasoning in the fourth model can have a (final) effect on the presence of  $m$ .<sup>31</sup>

<sup>30</sup>Note that  $(C \cup \{k\}) \setminus \{m\} = C \cup \{k\}$  if  $m \notin C$ .

<sup>31</sup>One might at first think that even in the fourth model reasoning has no *final* effect on  $m$ , since  $m$  is removed after having been added. This conclusion is however hasty, because the initial constitution  $C$  might already contain  $m$ , so that ‘adding’  $m$  means ‘keeping’  $m$ . Here reasoning does have a final effect on  $m$ :  $m$  was initially present, but is finally absent.

## D Proof of Theorem 1

For any constitution  $C \subseteq M$ , reasoning system  $S$ , and number  $n \in \{0, 1, \dots\}$ , let  $Cn_{S,n}(C)$  denote the set of constitutions reachable from  $C$  in  $n$  steps of  $S$ -reasoning. Now fix a reasoning system  $S$ . The corresponding Broomean operator  $Cn_S$  is obviously inclusive and idempotent, for reasons already indicated. To show monotonicity, consider constitutions  $C, D \subseteq M$  such that  $C \subseteq D$ , and fix a  $D^+ \in Cn_S(D)$ . Let  $S'$  be the reasoning system arising from  $S$  by replacing each rule  $(P, K) \in S$  satisfying  $K \cap D^+ \neq \emptyset$  with the rule  $(P, K \cap D^+)$ . So,

$$S' = \{(P, K) \in S : K \cap D^+ = \emptyset\} \cup \{(P, K \cap D^+) \in S : K \cap D^+ \neq \emptyset\}.$$

*Claim 1:* For every number of steps  $n \in \{0, 1, 2, \dots\}$ , all constitutions  $C^+ \in Cn_{S',n}(C)$  satisfy  $C^+ \subseteq D^+$  and  $C^+ \in Cn_{S,n}(C)$ .

We prove this claim by an induction on  $n$ . For  $n = 0$ , then the claim is obvious because  $Cn_{S',0}(C) = Cn_{S,0} = \{C\}$  and because  $C \subseteq D \subseteq D^+$ . Now assume  $n > 0$  and suppose the claim holds for smaller numbers than  $n$ . Fix any  $C^+ \in Cn_{S',n}(C)$ . Pick a  $\tilde{C} \in Cn_{S',n-1}(C)$  such that  $C^+ \in Cn_{S',1}(\tilde{C})$ . The fact that  $C^+ \in Cn_{S',1}(\tilde{C})$  and the definition of  $S'$  imply that  $C^+ \subseteq \tilde{C} \cup D^+$ . So, as by induction hypothesis  $\tilde{C} \subseteq D^+$ , we have  $C^+ \subseteq D^+$ .

It remains to show that  $C^+ \in Cn_{S,n}(C)$ . As  $\tilde{C} \in Cn_{S,n-1}(C)$ , it suffices to prove  $C^+ \in Cn_{S,1}(\tilde{C})$ . As  $C^+ \in Cn_{S',1}(\tilde{C})$ , we can pick a rule  $(P, K') \in S'$  by which  $C^+$  arises from  $\tilde{C}$ . By definition of  $S'$ , there exists a  $K \subseteq M$  such that  $(P, K) \in S$  and *either*  $[K \cap D^+ = \emptyset$  and  $K' = K]$  *or*  $[K \cap D^+ \neq \emptyset$  and  $K' = K \cap D^+]$ . The first case is impossible: it would imply that  $D^+$  is not closed under  $S$ -reasoning, because  $S$  contains the rule  $(P, K)$  which modifies  $D^+$  since  $P \subseteq D^+$  (as  $P \subseteq \tilde{C}$  and  $\tilde{C} \subseteq D^+$ ) and since  $K \cap D^+ = \emptyset$ . So the second case holds. Note that

$$\tilde{C} \cap K = \tilde{C} \cap [K' \cup (K \setminus D^+)] = (\tilde{C} \cap K') \cup (\tilde{C} \cap (K \setminus D^+)).$$

In the last expression,  $\tilde{C} \cap K' = \emptyset$  (as otherwise the rule  $(P, K')$  could not change  $\tilde{C}$ ) and  $\tilde{C} \cap (K \setminus D^+) = \emptyset$  (as  $\tilde{C} \subseteq D^+$  by induction hypothesis). So  $\tilde{C} \cap K = \emptyset$ . Since  $P \subseteq \tilde{C}$  and  $\tilde{C} \cap K = \emptyset$ , the rule  $(P, K)$  applies to  $\tilde{C}$ , just as the rule  $(P, K')$ . So, one can reason from  $\tilde{C}$  to  $C^+$  not just using  $(P, K')$ , but also using  $(P, K)$ . In other words,  $C^+$  belongs not just to  $Cn_{S',1}(\tilde{C})$ , but also to  $Cn_{S,1}(\tilde{C})$ . Q.e.d.

*Claim 2:*  $Cn_{S'}(C) \subseteq Cn_S(C)$ .

Let  $C^+ \in Cn_{S'}(C)$ . Then  $C^+ \in Cn_{S',n}(C)$  for some  $n \in \{0, 1, \dots\}$ . So, by Claim 1,  $C^+ \in Cn_{S,n}(C)$ . It remains to show that  $C^+$  is stable under  $S$ -reasoning. To show this, consider any rule  $(P, K) \in S$  such that  $P \subseteq C^+$ . We must show that  $K \cap C^+ \neq \emptyset$ . Form the rule  $(P, K') \in S'$ , where  $K'$  is  $K$  if  $K \cap D^+ = \emptyset$  and is  $K \cap D^+$  otherwise. Since  $C^+$  is closed under  $S'$ -reasoning, the rule  $(P, K')$  does not change  $C^+$ , i.e.,  $K' \cap C^+ \neq \emptyset$ . So,  $K \cap C^+ \neq \emptyset$ . Q.e.d.

*Claim 3:* Some  $C^+ \in Cn_S(C)$  satisfies  $C^+ \subseteq D^+$ .

All  $C^+ \in Cn_{S'}(C)$  satisfy  $C^+ \in Cn_S(C)$  by Claim 2, whence  $C^+ \subseteq D^+$  by Claim 1. Q.e.d.

## References

- Boghossian, P. (2014) What is Reasoning? *Philos Stud* 169: 1–18
- Broome, J. (2006) Reasoning with preferences? In: *Preferences and Well-Being*, S. Olsaretti ed., Cambridge University Press, 2006, pp. 183–208
- Broome, J. (2013) *Rationality Through Reasoning*, Hoboken: Wiley
- Dietrich, F., List, C. (2016) Mentalism versus Behaviourism in Economics: A Philosophy-of-science Perspective, *Economics & Philosophy* 32(2): 249–281
- Dietrich, F., List, C. (2021) The relation between degrees of belief and binary beliefs: a general impossibility theorem. In: I. Douven (ed.) *Lotteries, Knowledge, and Rational Belief. Essays on the Lottery Paradox*, Cambridge University Press, 2021, pp. 223–54
- Dietrich, F., Staras, A., Sugden, R. (2019) A Broomean model of rationality and reasoning, *Journal of Philosophy* 116: 585–614
- Dietrich, F., Staras, A., Sugden, R. (2022) Abstract rationality: the ‘logical’ structure of attitudes, working paper
- Drucker, D. (2021) Reasoning beyond belief acquisition, *Noûs*, see <https://doi.org/10.1111/nous.12363>
- Felappi, G. (forthcoming) Propositionalism and Questions that do not have Correct Answers, *Erkenntnis*, forthcoming. See <https://doi.org/10.1007/s10670-021-00442-5>
- Guala, F. (2019) Preferences: neither behavioural nor mental, *Economics & Philosophy* 35(3): 383–401
- Halpern, J. Y. (2017) *Reasoning About Uncertainty*, Cambridge, Massachusetts & London: MIT Press
- Harman, Gilbert (1984), Logic and reasoning, *Synthese* 60(1):107–127
- Kolodny, N. (2005) Why be rational? *Mind* 114: 509–563
- Kolodny, N. (2007) State or process requirements? *Mind* 116: 371–385
- Kyburg, H. E., Jr. (1961) *Probability and the Logic of Rational Belief*, Middletown, CT: Wesleyan University Press
- Liu, F. (2011) *Reasoning About Preference Dynamics*, Dordrecht: Springer
- Montague, M. (2007) Against Propositionalism, *Noûs* 41: 503–518
- Perea, A. (2012) *Epistemic Game Theory: Reasoning and Choice*, Cambridge University Press
- Staffel, J. (2013) Can there be reasoning with degrees of belief? *Synthese* 190(16): 3535–3551
- Tarski, A. (1956) *Logic, semantics and metamathematics*, Oxford University Press