



DOMINIK KLEIN
ALESSANDRA MARRA

From Oughts to Goals: A Logic for Enkrasia

Abstract. This paper focuses on (an interpretation of) the Enkratic principle of rationality, according to which rationality requires that if an agent sincerely and with conviction believes she ought to X , then X -ing is a goal in her plan. We analyze the logical structure of Enkrasia and its implications for deontic logic. To do so, we elaborate on the distinction between basic and derived oughts, and provide a multi-modal neighborhood logic with three characteristic operators: a non-normal operator for basic oughts, a non-normal operator for goals in plans, and a normal operator for derived oughts. We prove two completeness theorems for the resulting logic, and provide a dynamic extension of the logic by means of product updates. We illustrate how this setting informs deontic logic by considering issues related to the filtering of inconsistent oughts, the restricted validity of deontic closure, and the stability of oughts and goals under dynamics.

Keywords: Enkrasia, Basic oughts, Derived oughts, Goals, Deontic logic, Neighborhood logic, Dynamic logic.

1. Introduction

Suppose I believe sincerely and with conviction that today I *ought* to repay my friend Ann the 10 euro that she lent me. But I do not make any *plan for* repaying my debt: Instead, I arrange to spend my entire day at the local spa enjoying aromatherapy treatments. This seems wrong.

Enkrasia is the principle of rationality that rules out the above situation. The principle plays a central role within the domain of practical rationality, and has recently been receiving considerable attention in practical philosophy.¹ In its most general formulation, Enkrasia is the principle according to which rationality requires that if an agent sincerely and with conviction believes she ought to X , then she intends to X . There might be several ways in which such an intention to X is to be understood. Inspired by Bratman [4], here we consider the agent's intention to X as indicating that the agent is committed to achieve X , and thus has, in some sense, a plan for

¹See the works of Broome [6], Kolodny [20], Shpall [27], Horty [18]. For a complementary account of the relation between oughts and plans, see Gibbard [13].

Special Issue: Permissions, Obligations, and Beyond
Edited by Piotr Kulicki and Olivier Roy

X-ing. When this is the case, we say that *X*-ing is a goal in the agent’s plan. Combining these aspects, we can understand Enkrasia as the principle of rationality requiring that if an agent sincerely and with conviction believes she ought to *X*, then *X*-ing is a goal in the agent’s plan. Such an interpretation of Enkrasia was first suggested by Horty [18], and constitutes the starting point of the present paper. Notably, this formulation does not refer to the agent’s intention. In fact, we drop the term “intention” altogether from our analysis to avoid confusion.

This paper pursues two aims. Firstly, we want to analyze the logical structure of Enkrasia in light of the interpretation just described. This is, to the best of our knowledge, a largely novel project within the literature. Much existing work in modal logic deals with various aspects of practical rationality starting from Cohen and Levesque’s seminal paper [10]. The framework presented here aims to complement this literature by explicitly addressing Enkrasia. The principle, in fact, bears some non-trivial conceptual and formal implications—which might be of interest to the practical philosopher as well as the modal logician. This leads to the second aim of the paper. We want to address the repercussions that Enkrasia has for deontic logic. To this end, we elaborate on the distinction between so-called “basic oughts” and “derived oughts”, and show how this distinction is especially meaningful in the context of Enkrasia. Moreover, we address issues related to the filtering of inconsistent oughts, the restricted validity of deontic closure, and the stability of oughts and goals under dynamics.

In pursuit of these two aims, we provide a multi-modal neighborhood logic for Enkrasia. The logic has three characteristic operators: A non-normal operator for basic oughts, a non-normal operator for goals in plans, and a normal operator for derived oughts. We prove two completeness theorems for the resulting logic, and provide a dynamic extension of the logic by means of product updates.

The paper proceeds along the following general lines. First, we clarify its philosophical foundations by introducing Enkrasia’s main characteristics and its connection with two principles of rationality requiring goals in plans to be consistent (Section 2). We then introduce three challenges that illustrate the relevance of Enkrasia for deontic logic (Section 3). After discussing some core features and design choices of our approach (Section 4), we present a static logic for Enkrasia (Sections 5–7). Finally, we provide a dynamic extension of the logic (Section 8). This paper fits in with a larger project aimed at investigating the logic of oughts in the context of practical rationality. We hence conclude by discussing some related logic literature, and considering possible future extensions.

2. Enkrasia and the Consistency of Goals

The starting point of this paper is the Enkratic principle of rationality, in the following interpretation:

ENKRASIA. If an agent believes she ought to X, then X-ing is a goal in the agent's plan.

Such an interpretation is inspired by Horty [18]. This section introduces ENKRASIA's main components and emphasizes its connection with two principles of rationality governing goals in plans. Let us stress before continuing that the aim pursued here is not to engage in a direct defense of ENKRASIA (for this, the interested reader can consult [6] and [18]). Rather, this section is meant to lay the groundwork for our formal analysis of ENKRASIA's structure and of its position within the domain of practical rationality.

Let us begin with the oughts to which ENKRASIA applies; where "ought" is used as a noun, roughly meaning "obligation". It should be stressed that ENKRASIA does not take as antecedents all possible oughts. For one, ENKRASIA applies only to those oughts that are **believed by the agent**—in fact, this straightforwardly follows from the above formulation of the principle. However, further constraints are in place. We take inspiration from Broome [6], and require oughts that fall within the scope of ENKRASIA to have at least two further properties: They are **normative** and **ascribed to the agent herself**. These constraints are better illustrated via examples, so let us briefly consider them in turn.

One constraint limits the scope of ENKRASIA to normative oughts. These are the oughts that have to do, for instance, with morality, law or prudence. "I ought to repay my friend (as morality demands me to)" is an illustrative example of a normative ought. Contrariwise, examples of non-normative oughts are often to be found where oughts are used to express what is typically expected to be the case (see [32]), as in "I ought to have heard from the landing module ten minutes ago" ([6], p. 9). It would make little sense to say that hearing from the landing module is something I plan for. Indeed, ENKRASIA does not apply there.

The other constraint demands that the agent ascribes the oughts to herself. We can put this point in various ways: We can say this constraint demands that the agent believes the ought is required of her, that she recognizes it as her job to bring about the ought, or that she believes she is the "owner" of the ought (cf. [6], p. 22). Examples of oughts ascribed to the agent herself are "I ought to get a sun hat" ([6], p. 12), and "I ought to see to it that the kids are alright". An ought that is *not* ascribed to the agent

herself is “I ought to get a punishment” ([6], p.12), in a (natural) context where it is not on me to ensure that I receive this punishment. As long as getting a punishment is not my job, it would be incorrect to say that I fall short of rationality if getting punished is not a goal in my plan. This is why we demand ENKRASIA to apply only to oughts that are ascribed to the agent herself.

We have just identified a way in which ENKRASIA is constrained: It applies *only* to the oughts that enjoy the three properties above, namely, that are believed by the agent, normative, and ascribed to the agent herself. In our formal framework, we will implicitly assume that the oughts of ENKRASIA are of that kind. This is not to mean, however, that *all* oughts with those properties will correspond, via ENKRASIA, to goals in the agent’s plans. In fact, in the next section, we will suggest that ENKRASIA needs to be further weakened.

So much for oughts. Let us now turn to another crucial component of ENKRASIA: Goals in plans. Drawing from Bratman [4], when saying that *X-ing* is a goal in the agent’s plan, we mean that the agent is *committed* to achieve *X*, which includes figuring out (to an appropriate degree) how to do so. To put it more succinctly, we mean that the agent has a plan *for X-ing*. For instance, repaying my friend is a goal in my plan only if I am committed to do so: I have a plan for repaying my friend which, minimally, for me rules out all the options (such as spending all my money, leaving the country, etc.) that I believe would make it impossible to achieve my goal. Those options become, given my commitment to repay my friend, no longer **admissible**. In this context, goals in plans differ from mere desires or wishes, which lack such a dimension of commitment [10,28]. Those notions should be kept apart here.

Furthermore, goals in plans are **future-directed**: The most natural reading of “*X-ing* is a goal” is the one in which *X* is something that still has to happen (see [4], p. 4). Indeed, when talking about having a goal, we generally refer to something we are committed to do in the future (by the end of today, tomorrow, next month, etc.). In line with these considerations, in this paper we will assume *X* to include an element of futurity.

The literature imposes constraints on goals in plans. For instance, Broome suggests a property that—paraphrased in our own terms—amounts to requiring that the agent has the ability, via forming the goal to *X*, to have an impact on *X-ing* ([6], pp. 162–163). Although we find such a suggestion worth further (formal) analysis, we do not pursue this direction here. Rather, we focus our attention on two minimal principles of rationality governing

goals in plans. These principles of rationality require goals in plans to be **consistent**, in the following two senses of the term:

INTERNAL CONSISTENCY. If X -ing, Y -ing, ... are goals in an agent's plans, then it is *logically* consistent to X and Y and

STRONG CONSISTENCY. If X -ing, Y -ing, ... are goals in an agent's plans, then the agent *believes* it is *possible* to X and Y and

Both principles reflect the idea that it should be possible for goals in plans to be successfully achieved: INTERNAL CONSISTENCY demands that goals in plans be jointly logically consistent, while STRONG CONSISTENCY requires that goals in plans be jointly consistent with respect to the agent's beliefs. Their motivation is ultimately rooted in the dimension of commitment that goals in plans have: I could not truly be committed to repaying my friend and, at the same time, be committed to spending all my money to see the movies, while believing that these two things are incompatible—let alone jointly logically impossible (see [4, 10, 18]).

Straightforward consequences of the above consistency principles are that if X -ing is a goal in a plan, then not X -ing is not a goal in a plan (from INTERNAL CONSISTENCY), and that if X -ing is a goal in a plan, then the agent believes it is possible to X (from STRONG CONSISTENCY). That is to say, goals in plans should neither be contradictory, nor believed to be impossible to achieve.

This is perhaps the right moment to mention some aspects of the current debate surrounding ENKRASIA—and the principles of rationality, more generally—that we will *not* address in this paper. The first has to do with the debate on whether principles of rationality are of wide or narrow scope. Consider ENKRASIA. Under the narrow scope, if the agent believes she ought to X , *rationality requires that* X -ing is a goal in the agent's plan. Under the wide scope, on the other hand, *rationality requires that* if the agent believes she ought to X , then X -ing is a goal in her plan. The two readings lead to different pictures of rationality. Under the narrow scope reading, rationality requires a particular *attitude* of the agent. Under the wide scope, rationality only requires a particular *relation* between the agent's attitudes, typically leaving the rational agent leeway to either adopt X -ing as a goal in her plan or to revise her belief that she ought to X .² Since the focus of the present paper is not on operators akin to “rationality requires that”, we take our

²See, among others, Broome [6], Kolodny [20] and Shpall [27]. Broome defends the wide scope reading, Kolodny the narrow scope one, while Shpall proposes a “conciliatory view” between the two camps of the debate.

contribution to be largely independent of the question whether ENKRASIA is a narrow or wide scope principle of rationality.

A second issue to which this paper does not contribute is whether principles of rationality are synchronic or diachronic. Consider again ENKRASIA, now enriched with time-indexes: If the agent believes at t that she ought to X , then X -ing is a goal in the agent's plan at t' . Diachronically, t precedes t' , while synchronically, t and t' refer to the same time. Under the diachronic reading, believed oughts can be thought of *generating* corresponding goals, while under the synchronic interpretation, believed oughts and goals *coexist* at the same time. For reasons of simplicity, we follow Broome [6] and focus on the synchronic interpretation of ENKRASIA. We hold, however, that both interpretations have a certain appeal, especially from a logical perspective.

3. Three Challenges

We now introduce three challenges surrounding ENKRASIA that are apt to illustrate the relevance such a principle holds for deontic logic.

3.1. Challenge I: From Inconsistent Oughts to Consistent Goals

There is a potential *tension* between ENKRASIA and the principles of INTERNAL and STRONG CONSISTENCY for goals in plans. Consider the following:

EXAMPLE 3.1. Suppose I believe I ought to repay 10 euro to my friend Ann. I also believe I ought to go to the movies with Barbara (I have promised her so). However, money is scarce, and I believe it is impossible to do both.

It is safe to suppose that the oughts in Example 3.1 are of the kind to which ENKRASIA may apply, i.e., they enjoy all three properties introduced in Section 2. Now if ENKRASIA were in fact applied to those oughts, I would need to plan for both repaying the money to Ann and for going to the movies with Barbara—ending up with two goals I believe to be *inconsistent*, and so violating STRONG CONSISTENCY in this specific case.

How to solve this tension? One way is to assume oughts are always consistent, both from a logical viewpoint and from the perspective of the agent's beliefs (see Broome, [6]). This assumption certainly solves the problem. But consider again the example above. Especially when oughts originate from different sources, it seems a viable possibility that these may end up being jointly inconsistent.

In what follows, we investigate another strategy to solve the tension between ENKRASIA, INTERNAL and STRONG CONSISTENCY. In a nutshell, this

strategy is not to rule out the possibility of inconsistent oughts, nor to abandon the consistency principles for goals in plans, but rather to weaken ENKRASIA. The rationale for maintaining both INTERNAL and STRONG CONSISTENCY is rather pragmatic: In the face of a normative conflict about how to act, the least I can do is to assure that whatever I commit to is achievable.

Allowing for oughts, but not goals, to be inconsistent has several major consequences. Firstly, since oughts are possibly inconsistent but goals are not, it straightforwardly follows that not *all* oughts can correspond to goals in plans. In fact, this makes ENKRASIA a logically *invalid* principle. Secondly, it is natural to ask if not all, then *which* oughts do correspond to goals in plans. The challenge consists then in formally determining how oughts can be filtered out, in order to move from inconsistent oughts to consistent goals.

3.2. Challenge II: Basic Oughts and Derived Oughts

The second challenge revolves around a family of logical principles and inference rules that goes under the name of “deontic closure under implication”—for short: *deontic closure*. A longstanding tradition in deontic logic rejects the validity of deontic closure, arguing that it leads to unacceptable conclusions. An example is given by Ross’ Paradox [15,26]: Suppose *I ought to mail the letter*; now, since mailing the letter logically implies mailing the letter or burning it, deontic closure would imply that *I ought to mail the letter or burn it*—which is intuitively implausible.

The issue is that *even if* we accept that deontic closure is in fact problematic and should not be generally valid, an outright rejection of deontic closure would not constitute an adequate solution. For one, it would lead us to miss out also on deontic inferences that *are* intuitively plausible.

To see this, consider the following example, which we owe to Horty [18]. For this example, let us forget about my promise to go to the movies with Barbara, and simply assume that going to the movies is something I like:

EXAMPLE 3.2. Suppose that I ought to repay Ann 10 euro. Now suppose that I would also like to go to the movies, but I do not have a lot of money. In fact, I believe that unless I refrain from going to the movies it is impossible to repay Ann. So, I conclude, I ought not go to the movies.

Such a conclusion strikes us as impeccable. Following von Wright [31], we call the above piece of reasoning *practical inference*, and schematically represent it as:

- (P1) I ought to repay Ann
 (P2) Necessarily, repaying Ann implies not going to the movies
 (C) Therefore, I ought not go to the movies

Practical inference is the cornerstone of instrumental reasoning.³ Yet, practical inference—just as Ross’ Paradox—is a variant of deontic closure (specifically, deontic closure under necessary implication). An outright rejection of deontic closure would have the effect of also blocking the above derivation.

The challenge then takes the following shape: Even assuming that deontic closure is not generally valid, a deontic logic should be “thick” enough to license crucial deontic inferences—including those instances of deontic closure that are valid. In the remainder of this section, we explore the boundaries between valid and invalid instances of deontic closure, and show that ENKRASIA provides us with the conceptual tools to do so.

All we need is to fix one *set* of oughts to start with. This set functions as input for the agent’s deliberation. We do not impose any requirements on this set other than demanding that all oughts enjoy the three properties described in Section 2, i.e., being believed by the agent, normative, and ascribed to the agent herself. It follows, hence, that these are oughts to which ENKRASIA *may* apply. We call the oughts in this set **basic oughts**. Apart from what we just said, there is nothing intrinsically special about these.⁴ We do not assume basic oughts to have any particular surface grammar, nor do we assume they share any further commonalities. In fact, we even admit the possibility that basic oughts are jointly inconsistent.⁵ Once the set of basic oughts is fixed, we call **derived oughts** those oughts that are *implied* by basic oughts.

The distinction between basic and derived oughts is crucially meaningful in relation to ENKRASIA, and helps us to discern valid from invalid instances of deontic closure. Let us take practical inference as a case study. The central observation—originally noticed by Horty [18]—is that the oughts in (P1) and

³Typically (P2) expresses a *practical necessity*, which might vary with the circumstances or the agent’s beliefs thereof, cf. von Wright ([31], p. 161).

⁴Various interpretations can be imposed on the set of basic oughts. Horty [18] thinks of basic oughts as those oughts *directly* generated by normative requirements. Alternatively, one may think of basic oughts as those *explicitly* believed by the agent. These interpretations are compatible with our characterization of basic oughts. An alternative characterization of basic oughts is provided by Nair [24].

⁵This is why we have stressed that basic oughts are oughts to which ENKRASIA *may* apply. Since basic oughts are possibly inconsistent, while goals are not, it follows that not all basic oughts correspond to goals. See our discussion in Section 3.1.

in (C) interact differently with ENKRASIA. Suppose I deliberate about my day and I take as input that *I ought to repay Ann* (P1). In the absence of conflicts, this *basic ought* leads via ENKRASIA to the goal of repaying Ann, something that I plan for in itself. From there, via deontic closure, I do well in deriving that *I ought not go to the movies* (C). However this *derived ought* does not interact with ENKRASIA in the same way: Refraining from going to the movies is *not a goal in its own right*. Rather, it is something I necessarily have to do in order to fulfill my goal of repaying Ann. In other terms, the derived ought registers the *necessary* (though possibly not sufficient) *conditions* for the fulfillment of such a goal (see also [7]).

It is with respect to ENKRASIA that the different roles played by basic and derived oughts become evident. This motivates taking basic and derived oughts as *two* separate kinds of oughts in this context. Once these are understood as two separate oughts, having different logical meanings, it becomes non-trivial to say that there are instances of deontic closure that move from basic oughts to derived oughts. These instances will be valid in our logic. As elaborated above, this bears crucial implications for practical inference. Similar considerations apply to Ross' Paradox. Acknowledging the different roles of the oughts involved, I do well in deriving that *I ought to mail the letter or burn it* only to the extent that this expresses no more than the (logically) necessary—but not sufficient—conditions for the fulfillment of my goal of mailing the letter. In other terms, my inference is only valid to the extent that *I ought to mail the letter or burn it* is a derived ought.

3.3. Challenge III: Dynamic Conditions

The last challenge, finally, brings *dynamics* into the picture. Our initial reason for focusing on dynamics is mainly conceptual, as it is especially when seen through the lenses of dynamic change that the differences between basic oughts, goals and derived oughts become most prominent. The following two observations illustrate what we have in mind. Firstly, it is a well-known feature of goals in plans that they tend to be *stable* under various perturbations (cf. the seminal Bratman [4], pp. 16, 67). Reflecting the fact that goals are ultimately things the agent has committed to, there is a tendency for goals in plans to resist reconsideration, and in particular not to be discarded at every slight change that might occur in the agent's information.

EXAMPLE 3.3. Suppose that giving 10 euro back to Ann is a goal in my plan. My plan involves reaching Ann's house which I usually do by car. Now if I learned that my car is broken, I would not simply give up my goal to

repay Ann. Normally, I would rather maintain my goal and replan to get to her house by another means, for instance, by bus.

Let us call *practical dynamics* any changes in the agent’s information about what the world looks like. This is, in fact, the kind of dynamics at work in Example 3.3. Then, the basic idea is that goals in plans are not reconsidered whenever practical dynamics occur. Of course, goals are not irrevocable. They should, for instance, be dropped if they become inconsistent (cf. [4], p. 16). Yet, goals are stabler than other notions with respect to practical dynamics. This leads us to a second observation. Echoing Horty [18], we can appeal to a sort of “stability test” to illuminate the conceptual distinction between basic oughts, at least those that correspond to goals via ENKRASIA, and derived oughts. The following example shows that derived oughts are generally less stable with respect to practical dynamics:⁶

EXAMPLE 3.4. Suppose I ought to repay 10 euro to Ann (basic ought), and I hold the corresponding goal in my plan. Moreover, suppose that I do not have a lot of money, and hence conclude that I ought not go to the movies (derived ought), although I would really like to. Now if I learned that I have additional money at home, sufficient for both repaying Ann and buying a cinema ticket, I would give up that I ought not go to the movies. While I would maintain that I ought to repay Ann, and hence that doing so is a goal in my plan, I would *not* maintain that I ought to make sure not to go to the movies.

Hence, while the difference between basic oughts and derived oughts is not explicitly reflected in surface grammar, testing stability with respect to practical dynamics can help to make this conceptual distinction salient.⁷

The above observations show the relevance of investigating the effects of practical dynamics. Unanticipated obstacles or unexpected opportunities can diversely affect various notions at play in ENKRASIA, specifically goals in plans and derived oughts. The challenge then consists in precisely characterizing how and under which conditions these change dynamically.

⁶A version of Example 3.4 is discussed in Horty [18], p. 225.

⁷We thank an anonymous reviewer for drawing attention to the fact that the contrast between basic oughts and derived oughts is not marked in ordinary language. The distinction between the two oughts is predominantly conceptual.

4. Introducing the Framework

We can turn now to the first aim of this paper: Providing a logical framework for ENKRASIA. For the analysis, we posit a set of **minimal requirements** about basic oughts, goals in plans, and derived oughts. The reader may find these incomplete. However, our aim here is not to reveal the full logical principles governing oughts or goals. Rather, we aim for a minimal set of axioms strong enough to identify relations between basic oughts, goals and derived oughts that result from our analysis of ENKRASIA. In working towards a more complete logic of oughts and goals, additional axioms could be added in the future, validating further theorems. Of course, in such a stronger logic, the relationships identified here would continue to hold.

Despite the intended minimalism, devising a logic for ENKRASIA requires a variety of conceptual and formal choices. Some of these are core features of the framework developed. Others are mere design choices that could be altered easily. The following discussion details both.

4.1. Core Features

Basic Oughts, Goals and Derived Oughts The framework's first core component is three main logical operators: A modal operator for *basic oughts*, one for *goals* in plans, and finally one for *derived oughts*. We implicitly assume basic oughts to satisfy the three conditions identified in Section 2: They are believed by the agent, normative, and ascribed to the agent herself. Moreover, we take oughts and goals to be future-looking, referring to future states of affairs to be brought about.

Information States The framework focuses on a single moment in time, specifically, where the agent deliberates on what to do. The choice options represented in the logic are those believed possible by the agent at this moment; they form, in some sense, her *information state*.⁸ In fact, the framework with its various components is fully relative to the agent's beliefs, and so can do without any explicit doxastic operators.

A Thin Logic for Basic Oughts and Goals The starting point of the framework is a set of basic oughts. At present, the logic governing basic oughts remains thin. We do not assume basic oughts to have any logical structure such as being closed under implication or pairwise intersections, nor do we

⁸Unlike in most epistemic frameworks, this information state does not list epistemic possibilities the agent cannot distinguish between, but a set of possible options the agent can choose from.

require the content of a basic ought to be satisfiable, even in principle. The only requirement made is that basic oughts are independent of their exact description, i.e., the agent’s set of basic oughts is closed under replacement of logically equivalent formulas.⁹ The logic of basic oughts, hence, will turn out *weaker than normal* in the logical sense: It will be a neighborhood modal logic (cf. [25]). Similar considerations apply to goals. The set of goals in the agent’s plans will be a *consistent* subset of her basic oughts. Hence, also the goal modality will turn out to be a non-normal neighborhood operator.

A Thicker Logic for Derived Oughts While assuming the logic of basic oughts and goals to be thin, the resulting neighborhood logic is strong enough to license crucial deontic inferences. Derived oughts play a central role in such reasoning. To illustrate how these are represented in the framework, we first note that the agent may be committed to multiple goals in parallel. Following the principles of INTERNAL CONSISTENCY and STRONG CONSISTENCY, these goals are required to be jointly consistent. Put formally, this means that there must exist some possible course of events that satisfies all of the agent’s goals. We call such courses of events *admissible*. Derived oughts, then, denote those properties that all admissible courses of events have in common. In other words, derived oughts indicate the necessary (but possibly not sufficient) conditions for the fulfillment of *all* the agent’s goals. Derived oughts, unlike basic oughts, hence follow a *normal* modal logic.

4.2. Design Choices

Branching Temporal Trees Oughts and goals, we have said, are future-looking. Correspondingly, the agent’s relevant choices when deliberating on what to do are between possible future courses of events. In the present framework a fine-grained perspective on such future courses of events is assumed, representing the relevant temporal structure explicitly. To this end, all possible future unfoldings of the world are recorded in a temporally branching tree, where each maximal branch—each *history*—corresponds to a possible future course of events. For an illustration of a branching time setting, see Figure 1.

In accordance with this fine-grained perspective, oughts and goals need to be expressed in an adequate formal language rich enough to capture their

⁹We are arguably omitting certain structural properties of basic oughts. For instance, a plausible further requirement to impose on basic oughts could be what Cariani [8] calls “weakening”: $O\varphi \wedge O\psi \vDash O(\varphi \vee \psi)$. Nevertheless, we will show that the present framework contains sufficient structure for a logical analysis of ENKRASIA.




Figure 1. Left: The subtree compatible with the satisfaction of the agent’s basic ought and goal that φ (gray). Middle: The subtree compatible with the satisfaction of agent’s basic ought and goal that ψ (gray). Right: Interaction of both basic oughts and goals (dark gray). Bold arrows denote the admissible subtree, i.e., the courses of events compatible with the satisfaction of both goals φ and ψ

temporal structure. To this end, the framework involves a temporal logic that can express, for instance, that certain states of affairs should be always avoided, reached at least once or maintained throughout.

Notably, representing possible courses of events as temporally extended histories is not strictly necessary. For the static part of the logic (Section 7), it would suffice to treat each possible course of events as a single state, giving rise to a more classic neighborhood logic. It is only in the dynamic extension of Section 8 that the temporal structure becomes relevant.

From Basic Oughts to Goals An agent’s goals, we have said, form a subset of her basic oughts. The latter, however, are potentially inconsistent whereas goals are not. A central component for the transition from basic oughts to goals will hence be (maximally) consistent subsets of basic oughts, as these guarantee that the principles of INTERNAL CONSISTENCY and STRONG CONSISTENCY are satisfied.¹⁰ Note, however, that there can be multiple maximally consistent subsets of basic oughts. So how are goals related to maximally consistent sets of basic oughts? There exist at least two viable ways of approaching this:

- In a strict reading, a basic ought is adopted as goal if it is contained in *every* maximally consistent set of basic oughts.

¹⁰A competing notion of consistency, which we might call *free choice consistency*, is discussed in Veltman [29]. Veltman would consider $O(\neg\varphi)$ and $O(\varphi \vee \psi)$ inconsistent, as the former violates the free choice expressed by the latter. In this paper we do not deal with free choice, and hence we limit ourselves to a classic account of consistency. Free choice in the context of planning is considered in Marra and Klein [23].

- In a more tolerant approach, a basic ought is adopted as goal if it is contained in *some specific* maximally consistent set of basic oughts.

The tolerant approach will, in general, lead to more goals than the strict approach. In fact, by picking a single maximally consistent subset of basic oughts, it guarantees the agent to do the best she can in terms of adopting a multitude of goals without violating consistency. The following analysis follows the tolerant approach.¹¹ We are hence in need of a mechanism for selecting which maximally consistent set of basic oughts corresponds to goals.

Linear Priority on Basic Oughts For selecting a maximally consistent subset of basic oughts, we assume the latter to be ordered linearly.¹² By means of the lexicographic order (cf. Definition 6.6), this linear order extends to a priority ordering among sets of basic oughts. The agent then adopts the highest ranked maximally consistent subset of her basic oughts as goals. The current framework is, however, modular in this respect. Any other mechanism for picking out one element from any given set of maximally consistent set of oughts would function just as well. In fact, the choice of selection mechanism does not have *any* impact of the static analysis of Sections 6 and 7. In particular, the assumption of oughts being ordered *linearly* is non-substantial for the present purpose.

¹¹The strict approach is prominently pursued by Kratzer [21] in her seminal approach to the semantics of deontic operators. There, briefly, a possibly inconsistent set of normative requirements N creates an ideality ordering on a set of possible worlds W . To define the ordering, let $N(w)$ for a world w be the set of normative requirements from N satisfied at w . The ordering is then defined by $w > v$ (read “ w is more ideal than v ”) if $N(w) \supset N(v)$. A deontic necessity statement $\Box\varphi$, finally, holds true in the framework if φ is satisfied in *all* $>$ -maximal worlds. Notably, $>$ -maximality is tightly related to maximally consistent subsets. More specifically, world w is $>$ -maximal iff no M with $N(w) \subset M \subseteq N$ is satisfiable in any $v \in W$, i.e., iff $N(w)$ is maximally W -consistent. It follows that $\Box\varphi$ is true iff φ holds in *all* intersections of maximally consistent subsets of norms. This is exactly the above strict reading.

In fact, various aspects of Kratzer’s approach have counterparts in the present framework. To make these explicit: normative requirements N and possible worlds correspond to basic oughts and histories of tree \mathcal{T} respectively. The deontic necessity operator, finally, corresponds to our modality for derived oughts.

¹²Hence, although we do not rule out the possibility of having both $O\varphi$ and $O\neg\varphi$ as basic oughts, we exclude irresolvable dilemmas. One basic ought must take priority over the other.

4.3. Towards a Logic for Enkrasia

The construction of our formal framework proceeds in two steps. The first step (Sections 5–7) defines two static logics Λ_{Enkr} and $\Lambda_{Enkr, \square}$. Having modalities for basic oughts, goals and derived oughts, these already incorporate ENKRASIA through a number of axioms regulating the relationship between the three components. The second of these logics offers an additional global modality \square allowing the agent to reason about which options are available to her.

The second step (Section 8) adds dynamic operations to the logics defined. We focus on practical dynamics: The updating operations we consider can add or remove possible courses of events, but leave the agent’s basic oughts unaltered. Nevertheless, practical dynamics turn out to have complex effects on goals and derived oughts. Studying these—we hold—provides additional insights into the relationship between basic oughts, goals and derived oughts.

5. The Language

To begin, let us specify the logical language used. The construction proceeds in several steps. First, we define two languages \mathcal{L}_0 and \mathcal{L}_1 to talk about present and future states of affairs. This language will serve to express *the content* of oughts and goals. Afterwards, we introduce language \mathcal{L}_2 that allows to reason about basic and derived oughts, goals, and their interaction.

DEFINITION 5.1. Let At be a finite or countable set of atomic propositions. The **basic language** \mathcal{L}_0 is given by the standard language of propositional logic combined with a future-tensed operator F . It is defined by the following BNF:

$$\psi := p \mid \neg\psi \mid \psi \wedge \psi \mid F\psi$$

for $p \in \text{At}$. The intended reading of modal expressions $F\psi$ is “ ψ is true at least once in the future”. We denote the dual of F by G . $G\psi$ hence reads as “ ψ is always true in the future”. Operators \rightarrow and \vee , finally, are defined as usual.

It is convenient to consider the future looking fragment of \mathcal{L}_0 :

DEFINITION 5.2. The **language** \mathcal{L}_1 is the fragment of \mathcal{L}_0 containing only future-tensed formulas where every atomic proposition is in the scope of a temporal operator. Formally, \mathcal{L}_1 is defined as follows:

$$\varphi := F\psi | \neg\varphi | \varphi \wedge \varphi$$

for $\psi \in \mathcal{L}_0$. Building on \mathcal{L}_1 , the modal language for reasoning about basic oughts, goals in plans, and derived oughts can be defined.

DEFINITION 5.3. The modal **language** \mathcal{L}_2 is given by the following BNF:

$$\varphi := p | O\psi | Goal\psi | D\psi | \neg\varphi | \varphi \wedge \varphi$$

for $p \in \text{At}$ and $\psi \in \mathcal{L}_1$. The intended reading of the three modal operators is the following: $O\varphi$ reads as “it is a basic ought that φ ”, $Goal\varphi$ as “it is a goal in a plan that φ ”, and finally $D\varphi$ reads as “it is a derived ought that φ ”. Again, operators \rightarrow and \vee are defined as usual.

Two observations about \mathcal{L}_2 are in order. Firstly, the language does not allow for iterated modalities. This is a feature shared with several other systems of deontic logic. Secondly, being built over the temporal fragment \mathcal{L}_1 of \mathcal{L}_0 , the modal language \mathcal{L}_2 only allows for basic oughts, goals and derived oughts that scope over future-tensed formulas. Our oughts and goals are, as we have said, future-looking.

6. Semantics

Before introducing logical principles on the above languages, we specify the intended semantic structures for basic oughts, goals, and derived oughts. Section 7 then provides an axiomatization that is sound and complete with respect to the semantics introduced here. We begin our analysis by introducing trees, delineating how the agent envisages the possible unfoldings of future events.

DEFINITION 6.1. A **tree** is an ordered set $\mathcal{T} = \langle T, \prec_{\mathcal{T}} \rangle$ where T is a set of moments and $\prec_{\mathcal{T}}$ a tree-order on T . We make two additional assumptions about $\prec_{\mathcal{T}}$. First, the tree order is assumed to have a **root**, i.e., a minimal element t_0 satisfying $t_0 \prec_{\mathcal{T}} t$ for all $t \neq t_0$. Second, $\prec_{\mathcal{T}}$ is assumed serial, i.e., every moment must have at least one successor.¹³ A **history** h , finally, is a maximal linearly ordered subset of \mathcal{T} .

Intuitively, t_0 indicates the current time step, i.e, the moment at which the agent ponders what to do. Notably, a tree is the union of its histories, i.e., $\mathcal{T} = \bigcup \{h \subseteq \mathcal{T} \mid h \text{ history}\}$. We will make heavy use of this later. To ease

¹³This definition remains silent about the exact shape of a tree. It allows for finite as well as infinite branchings and also for discrete as well as dense orders.

terminology, we will use the term **subtree** for any tree \mathcal{T}' that is of the form $\bigcup_{h \in \text{Hist}} h$ with Hist a set of histories of \mathcal{T} . We will denote the set of subtrees of \mathcal{T} by $\mathcal{P}(\mathcal{T})$. Lastly, let \mathcal{T}' and \mathcal{T}'' be subtrees of \mathcal{T} given by $\mathcal{T}' = \bigcup_{h \in \text{Hist}' } h$ and $\mathcal{T}'' = \bigcup_{h \in \text{Hist}'' } h$ respectively. Then define the **intersection subtree** $\mathcal{T}' \pitchfork \mathcal{T}''$ of \mathcal{T} as the subtree generated by $\text{Hist}' \cap \text{Hist}''$, i.e.,¹⁴ $\mathcal{T}' \pitchfork \mathcal{T}'' := \bigcup_{h \in \text{Hist}' \cap \text{Hist}'' } h$.

Based on the definition of a tree, we can define a tree model for our temporal language \mathcal{L}_0 .

DEFINITION 6.2. A **pointed tree model** is a tuple $\mathcal{M} = \langle \mathcal{T}, t_0, v \rangle$ where $\mathcal{T} = \langle T, \prec_{\mathcal{T}} \rangle$ is a tree, t_0 the distinguished time, i.e., the root of \mathcal{T} and $v : \text{At} \rightarrow \mathcal{P}(T)$ is a valuation function that maps each atomic proposition of the background language into a set of moments of \mathcal{T} .

A pointed tree model provides a semantics for language \mathcal{L}_0 and hence also \mathcal{L}_1 .

DEFINITION 6.3. Let \mathcal{M} be a pointed tree model. The **evaluation** of formulas of \mathcal{L}_0 on time-history pairs t/h with $t \in h$ of \mathcal{M} is defined as follows:

- $\mathcal{M}, t/h \models p$ iff $t \in v(p)$ for p atomic
- $\mathcal{M}, t/h \models \neg\varphi$ iff $\mathcal{M}, t/h \not\models \varphi$
- $\mathcal{M}, t/h \models \varphi \wedge \psi$ iff $\mathcal{M}, t/h \models \varphi$ and $\mathcal{M}, t/h \models \psi$
- $\mathcal{M}, t/h \models F\varphi$ iff there is a $t' \in h$ such that $t \prec_{\mathcal{T}} t'$ and $\mathcal{M}, t'/h \models \varphi$

Finally, we say that a formula is true at t simpliciter iff it is true at t/h' for all histories h' passing through t .

DEFINITION 6.4. Let $\varphi \in \mathcal{L}_0$ and $t \in \mathcal{T}$. The proposition expressed by φ at t , i.e., the **truth subtree** $\llbracket \varphi \rrbracket^t$, is defined as follows:

$$\llbracket \varphi \rrbracket^t = \bigcup \{ h \mid t \in h \text{ and } \mathcal{M}, t/h \models \varphi \}$$

Towards developing a semantics for the language \mathcal{L}_2 , we finally extend tree models with neighborhoods representing the agent's basic oughts. A central component of these extended models will be sets of the form $\llbracket \varphi \rrbracket^{t_0}$, representing the truth set of φ as seen from the moment of deliberation t_0 .

DEFINITION 6.5. An **enkratic model** is a tuple $\mathcal{M} = \langle \mathcal{T}, t_0, v, N_O, \succ_O \rangle$ where $\langle \mathcal{T}, t_0, v \rangle$ is a pointed tree model, and $N_O \subseteq \mathcal{P}(\mathcal{T}) \times \mathcal{L}_1$ is a neighborhood with the additional condition that $(\mathcal{T}', \varphi) \in N_O$ implies that $\mathcal{T}' = \llbracket \varphi \rrbracket^{t_0}$.

¹⁴Note that $\mathcal{T}' \pitchfork \mathcal{T}'' \subseteq \mathcal{T} \cap \mathcal{T}'$. In general, however, $\mathcal{T}' \pitchfork \mathcal{T}''$ is a proper subset of $\mathcal{T} \cap \mathcal{T}'$.

Finally \succ_O is a conversely well-founded linear order on the set of all φ such that $(\llbracket\varphi\rrbracket^{t_0}, \varphi) \in N_O$.¹⁵

Presently, we are only interested in the agent's basic oughts at the time of reasoning t_0 . We can represent these with a set of treelike neighborhoods N_O listing all the basic oughts the agent is exposed to at t_0 .¹⁶ It might seem counterintuitive to represent a basic ought by a subset-formula pair $(\llbracket\varphi\rrbracket^{t_0}, \varphi)$ rather than simply a subtree $\llbracket\varphi\rrbracket^{t_0}$. The reason for this will become clear in Section 8 where dynamics enter the picture. Briefly, two propositions φ and ψ may be co-extensional in the current tree, but might cease to be so once new information about the world is acquired. For this case, it is necessary to keep track of whether the basic ought prescribes that φ or ψ .

On a given enkratic model, we can construct additional structures related to the semantics of goals and derived oughts. The first is the **goal-neighborhood** $N_G \subseteq \mathcal{P}(\mathcal{T})$. For the construction we recall the definition of a lexicographic order.

DEFINITION 6.6. Let \succ_O be a conversely well-founded linear order on a set of formulas $\Psi \subseteq \mathcal{L}_1$. Then the **lexicographic order** \succ_{Lex} on the power set $\mathcal{P}(\Psi)$ is defined by $X \succ_{Lex} Y$ iff there is some $x \in X$, $x \notin Y$ such that

$$\{z \in X \mid z \succ_O x\} = \{z \in Y \mid z \succ_O x\}.$$

In other words, x is the \succ_O -most important element on which X and Y disagree.

The goal neighborhood $N_G \subseteq \mathcal{P}(\mathcal{T})$ is determined by three conditions. First, the goals in an agent's plan must be derived from basic oughts. Second, the set of goals in a plan should be consistent. The third condition, finally, expresses that the set of goals is chosen optimally, given the agent's priority relation \succ_O between her basic oughts. Formally, the conditions on N_G are:

- (i) $N_G \subseteq \{\llbracket\varphi\rrbracket^{t_0} \mid (\llbracket\varphi\rrbracket^{t_0}, \varphi) \in N_O\}$.
- (ii) N_G is maximally consistent, i.e.,
 - (a) there is some history h of \mathcal{T} with $h \subseteq \llbracket\varphi\rrbracket^{t_0}$ for all $\llbracket\varphi\rrbracket^{t_0} \in N_G$ and

¹⁵Where \succ_O is a conversely well-founded linear order if and only if it is antisymmetric, transitive, total and every subset $B \subseteq \{\varphi \mid (\llbracket\varphi\rrbracket^{t_0}, \varphi) \in N_O\}$ has a \succ_O -maximal element.

¹⁶The approach could be extended to include the agent's basic oughts along all moments of a tree. Such an extension requires additional conceptual work, as basic oughts may, for instance, get discarded once they have been satisfied. Also, an extension will need to specify what happens to basic oughts in future moments where they have become unsatisfiable for pragmatic or principal reasons. Technically, such an extension would work by replacing the neighborhood N_O with a neighborhood function $n_O : \mathcal{T} \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{T}) \times \mathcal{L}_1)$.

(b) whenever $N_G \subset Y \subseteq \{\llbracket \varphi \rrbracket^{t_0} \mid (\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O\}$ there is no history h' with $h' \subseteq \llbracket \varphi \rrbracket^{t_0}$ for all $\llbracket \varphi \rrbracket^{t_0} \in Y$.

(iii) N_G is \succ_O -maximal, i.e., whenever Y satisfies (i) and (ii) then

$$\{\varphi \mid \llbracket \varphi \rrbracket^{t_0} \in N_G\} \succ_{Lex} \{\varphi \mid \llbracket \varphi \rrbracket^{t_0} \in Y\};$$

where \succ_{Lex} is the lexicographic order on $\mathcal{P}(\{\varphi \mid (\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O\})$ induced by \succ_O . (Cf. Definition 6.6).

Note that the three conditions uniquely determine the neighborhood N_G which is therefore well-defined.

From N_G the third central component of enkratic models—besides basic oughts and goals—can be defined. Let us begin by introducing what we call the **admissible subtree** \mathcal{T}_{adm} . The admissible subtree \mathcal{T}_{adm} , briefly, is the intersection of the various subtrees corresponding to the agent's goals. Hence, it consists of all those histories that guarantee *all* of the agent's goals to be satisfied. It is from this admissible subtree that the agent's derived oughts are determined. Derived oughts indicate what holds in \mathcal{T}_{adm} , and therefore can be thought of as expressing the necessary conditions for the fulfillment of all the agent's goals. To state things formally, the admissible subtree is defined as

$$\mathcal{T}_{adm} := \bigcap_{\llbracket \varphi \rrbracket^{t_0} \in N_G} \llbracket \varphi \rrbracket^{t_0}.$$

From the properties of N_G , it follows that \mathcal{T}_{adm} is non-empty. Having defined \mathcal{T}_{adm} , we can give the semantic conditions turning enkratic model into models for language \mathcal{L}_2 . Unlike \mathcal{L}_0 , the language \mathcal{L}_2 is evaluated on moments t rather than time-history pairs t/h . We take this to be a natural condition, as \mathcal{L}_2 represents the agent's oughts and goals at a moment in time t_0 where she has not yet acted on any particular course of events, i.e., any history h . The following definition builds on the evaluation of \mathcal{L}_0 (and hence \mathcal{L}_1) on pointed tree models (cf. Definition 6.3).

DEFINITION 6.7. The **evaluation** of \mathcal{L}_2 on an enkratic model \mathcal{M} is given by the following clauses:

- $\mathcal{M}, t \models p$ iff $t \in v(p)$ for p atomic
- $\mathcal{M}, t \models \neg\varphi$ iff $\mathcal{M}, t \not\models \varphi$
- $\mathcal{M}, t \models \varphi \wedge \psi$ iff $\mathcal{M}, t \models \varphi$ and $\mathcal{M}, t \models \psi$
- $\mathcal{M}, t \models O\varphi$ iff $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O$
- $\mathcal{M}, t \models Goal\varphi$ iff $\llbracket \varphi \rrbracket^{t_0} \in N_G$ and $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O$.

- $\mathcal{M}, t \models D\varphi$ iff $\mathcal{T}_{adm} \subseteq \llbracket \varphi \rrbracket^{t_0}$

Notably, the semantics of operators O , $Goal$ and D does not depend on the moment t of evaluation, but only on the initial time t_0 . These modalities, hence, are meant to represent the agent's basic oughts, goals and derived oughts *at the time of deliberation* t_0 .

In sum, the semantics of all three modalities supervenes on two components of the model: The neighborhood N_O and the priority ordering \succ_O . While the semantics of O , the basic ought modality, is directly given by N_O , the $Goal$ modality's neighborhood is derived by having \succ_O pick a maximally consistent subset of N_O . This goal neighborhood, in turn, defines the derived ought modality D 's admissible subtree by means of intersection.

7. Syntax: Axioms and Results

In this section, we provide an axiomatization for the various languages introduced in Section 5. We start with axioms for the temporal languages \mathcal{L}_0 and \mathcal{L}_1 .

K_G	$G(\varphi \rightarrow \psi) \rightarrow (G\varphi \rightarrow G\psi)$
4	$G\varphi \rightarrow GG\varphi$
L	$F\varphi \wedge F\psi \rightarrow (F(\varphi \wedge \psi) \vee F(\varphi \wedge F\psi) \vee F(\psi \wedge F\varphi))$
D_G	$\neg G\perp$

These are accompanied by the classic necessitation rule

$$\frac{\vdash \varphi}{\vdash G\varphi} \text{NEC}_G$$

The first two axioms are the standard K and 4 axioms, expressing that G is a normal modal operator and that the 'later' relation is transitive. The third axiom L reflects the fact that histories are linear, expressing that two future events φ and ψ will either be simultaneous, or that one comes after the other. Finally, the D -style axiom D_G expresses that time never ends, as there always is a future moment. We denote the temporal logic over language \mathcal{L}_0 generated by $K_G, 4, L, D_G$ and G -necessitation NEC_G by \mathbf{A}_{temp} .

Next, we turn to the extended language \mathcal{L}_2 . Operators O and $Goal$ only have a limited logical structure. Reflecting the actual content of oughts issued by a normative source, we do not presuppose any logical requirements on basic oughts other than being invariant under replacement with logical equivalents. This is the content of:

$$\frac{\vdash \varphi \leftrightarrow \psi}{\vdash O\varphi \leftrightarrow O\psi} \text{INT}_O$$

The corresponding intensionality condition for the *Goal* operator also holds, as is shown in Lemma 7.3. While goals and basic oughts are not closed under logical reasoning, derived oughts are. In particular, the *D*-operator is normal and non-trivial, as expressed by the following axioms

$$\begin{array}{ll} K_D & D(\varphi \rightarrow \psi) \rightarrow (D\varphi \rightarrow D\psi) \\ D_D & \neg D\perp \end{array} \quad \frac{\vdash \varphi}{\vdash D\varphi} \text{NEC}_D$$

Lastly, and most importantly, the logic is guided by three interaction axioms describing the interplay between goals, basic and derived oughts. It is these principles that embody the ENKRASIA principle in the logic.

$$\begin{array}{ll} \text{GO} & \textit{Goal}\varphi \rightarrow O\varphi \\ \text{GD} & \textit{Goal}\varphi \rightarrow D\varphi \\ \text{MAX} & O\varphi \wedge \neg \textit{Goal}\varphi \rightarrow D\neg\varphi \end{array}$$

The first of these expresses that basic oughts are the only admissible sources of goals in the agent's plan. Every *Goal* follows from a basic *Ought*. The second axiom, GD, is a weak converse, saying that every *Goal* gives rise to a corresponding *Derived* ought. The third axiom, MAX, finally embodies the *bounded validity* of ENKRASIA, as can best be seen from its counterpositive $\neg D\neg\varphi \rightarrow (O\varphi \rightarrow \textit{Goal}\varphi)$: If it is not the case that already $\neg\varphi$ is a derived ought, then a basic ought that φ triggers a goal that φ . Hence, in combination with K_D and D_D , MAX states that every basic ought has a corresponding goal *unless* this causes a violation of consistency. Put semantically, MAX expresses that the set of goals is a *Maximally* consistent subset of the agent's basic oughts.¹⁷

DEFINITION 7.1. The **Enkrasia logic** Λ_{Enkr} on language \mathcal{L}_2 is defined by all propositional tautologies together with the axioms $K_G, 4, L, D_G, K_D, D_D, GO, GD, MAX$ and the rules $\text{INT}_O, \text{NEC}_G$ and NEC_D (cf. Table 1).

Before moving on to completeness, let us take a moment to derive a number of consequences of the above axioms. First, we note that whenever an agent has a goal to φ , her derived oughts contain all logical consequences of φ . This follows immediately from axioms K_D and GD together with NEC_D .

¹⁷If we had instead chosen the strict principle of translating basic oughts into goals, (cf. Section 4.2) i.e., only taking those basic oughts that are contained in all maximally consistent subsets instead, MAX would need to be replaced by the weaker $O\varphi \wedge \neg \textit{Goal}\varphi \rightarrow \neg D\varphi$.

Table 1. Axioms and rules of Λ_{temp} , Λ_{Enkr} and $\Lambda_{Enkr, \square}$

Axioms of Λ_{temp} over \mathcal{L}_0		
4	$G\varphi \rightarrow GG\varphi$	$\frac{\vdash \varphi}{\vdash G\varphi} \text{NEC}_G$
D_G	$\neg G\perp$	
K_G	$G(\varphi \rightarrow \psi) \rightarrow (G\varphi \rightarrow G\psi)$	
L	$F\varphi \wedge F\psi \rightarrow (F(\varphi \wedge \psi) \vee F(\varphi \wedge F\psi) \vee F(\psi \wedge F\varphi))$	
Axioms of Λ_{Enkr} over \mathcal{L}_2		
K_D	$D(\varphi \rightarrow \psi) \rightarrow (D\varphi \rightarrow D\psi)$	$\frac{\vdash \varphi}{\vdash D\varphi} \text{NEC}_D$
D_D	$\neg D\perp$	
GD	$Goal\varphi \rightarrow D\varphi$	
GO	$Goal\varphi \rightarrow O\varphi$	$\frac{\vdash \varphi \leftrightarrow \psi}{\vdash O\varphi \leftrightarrow O\psi} \text{INT}_O$
MAX	$O\varphi \wedge \neg Goal\varphi \rightarrow D\neg\varphi$	
Λ_{temp} for \mathcal{L}_1 formulas inside $O, Goal, D$		
Additional Axioms for $\Lambda_{Enkr, \square}$ over \mathcal{L}_\square		
ND	$\square\varphi \rightarrow D\varphi$	$\frac{\vdash \varphi}{\vdash \square\varphi} \text{NEC}_\square$
K_\square	$\square(\varphi \rightarrow \psi) \rightarrow (\square\varphi \rightarrow \square\psi)$	

FACT 7.2.

$$\frac{\vdash \varphi \rightarrow \psi}{\vdash Goal\varphi \rightarrow D\psi}$$

Second, we note that the *Goal* operator is closed under replacement with logical equivalents:

LEMMA 7.3.

$$\frac{\vdash \varphi \leftrightarrow \psi}{\vdash Goal\varphi \leftrightarrow Goal\psi}$$

PROOF. Assume $\vdash \varphi \leftrightarrow \psi$. For a contradiction, also assume that $Goal\varphi$ but $\neg Goal\psi$. By GO we have $O\varphi$ and hence by INT_O also $O\psi$. Hence we have $O\psi \wedge \neg Goal\psi$ which implies $D\neg\psi$ by MAX. On the other hand, GD implies $D\varphi$. By NEC_D and K_D this implies $D(\varphi \wedge \neg\psi)$ which, again by K_D , implies $D\perp$ contradicting D_D . ■

Next, note that the logic does not demand an agent's basic oughts to be jointly consistent. Our agent may, for instance, believe both $O\varphi$ and $O\neg\varphi$ simultaneously. The set of goals, however, is required to be internally consistent.

LEMMA 7.4. *Let $\Lambda \subseteq \mathcal{L}_2$ be a consistent set and let $S = \{\varphi \in \mathcal{L}_1 \mid \text{Goal}\varphi \in \Lambda\}$. Then $S \not\vdash_{\Lambda_{temp}} \perp$.*

PROOF. Assume for a contradiction that $S \vdash_{\Lambda_{temp}} \perp$. Since all its axioms correspond to first order expressible frame conditions, Λ_{temp} is compact (cf. [3], Chapter 2.4). Hence there is a finite $S_0 \subseteq S$ such that $S_0 \vdash_{\Lambda_{temp}} \perp$. By Fact 7.2, we have $\{\text{Goal}\varphi \mid \varphi \in S_0\} \vdash_{\Lambda_{Enkr}} \bigwedge_{\varphi \in S_0} D\varphi$. By K_D we then get $\{\text{Goal}\varphi \mid \varphi \in S\} \vdash_{\Lambda_{Enkr}} D \bigwedge_{\varphi \in S_0} \varphi$, i.e., $\{\text{Goal}\varphi \mid \varphi \in S\} \vdash_{\Lambda_{Enkr}} D\perp$ contradicting D_D . ■

An immediate consequence is that the *Goal* operator satisfies the *D*-axiom, i.e.,

$$\vdash \text{Goal}\varphi \rightarrow \neg \text{Goal}\neg\varphi$$

In fact, this consistency requirement is solely responsible for discrepancies between basic oughts and goals. By MAX, whenever $O\varphi \wedge \neg \text{Goal}\varphi$ hold at some state w , this is because *Goal* φ could not have been consistently added to the set of present goals, as it would require both $D\varphi$ and $D\neg\varphi$ to hold simultaneously.

Having specified our treatment of ENKRASIA, it is now time to present a general characterization result. However, before being able to do so, we need to make an extra assumption about enkratic models. For the rest of this paper, we assume the neighborhood N_O to be closed under logical equivalence. That is, if φ and ψ are logically equivalent in Λ_{temp} and $(\llbracket\varphi\rrbracket^{t_0}, \varphi) \in N_O$ then also $(\llbracket\psi\rrbracket^{t_0}, \psi) \in N_O$. It follows immediately that also N_G is closed under Λ_{temp} logical equivalence. With this assumption, we can show the following characterization result, which is proved in the appendix.

THEOREM 7.5. *The logic Λ_{Enkr} is sound and complete with respect to the class of enkratic models.*

7.1. Enriching the Language: A Global Modality

Note that language \mathcal{L}_2 suffers from what might be perceived as a lack of expressive power. So far, \mathcal{L}_2 can express whether the agent is under a certain basic ought that φ and whether this ought translates into a goal. What \mathcal{L}_2 cannot yet express is whether the agent considers φ possible in the first place, i.e., whether she believes her basic ought that φ to be satisfiable. To remedy this, we add a new modal operator \Box , where $\Box\varphi$ for some $\varphi \in \mathcal{L}_1$ is to express that φ holds in all histories considered possible by the agent. As usual, \Diamond stands for the dual of \Box . So $\Diamond\psi$ expresses that there is a possible ψ -history or, at least, one the agent considers possible. To incorporate \Box , we expand language \mathcal{L}_2 to \mathcal{L}_\Box given by the BNF:

$$\varphi := p|O\psi|Goal\psi|D\psi|\Box\psi|\neg\varphi|\varphi \wedge \varphi$$

for $p \in \text{At}$ and $\psi \in \mathcal{L}_1$. The semantics of \mathcal{L}_\Box on an enkratic model is given by the semantics of \mathcal{L}_2 extended with the clause

$$\mathcal{M}, t \models \Box\varphi \text{ iff } \mathcal{M}, t/h' \models \varphi \text{ for all branches } h' \text{ with } t \in h'$$

On the axiomatic side, the new modality \Box is governed by the axioms and rules below. The first, ND, expresses that the agent is under a derived ought to φ whenever φ is unavoidable to her. K_\Box is the K -axiom for \Box .

$$\begin{array}{l} \text{ND} \quad \Box\varphi \rightarrow D\varphi \\ \text{K}_\Box \quad \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi) \end{array} \qquad \frac{\vdash \varphi}{\vdash \Box\varphi} \text{NEC}_\Box$$

We denote the extension of Λ_{Enkr} with ND, K_\Box and NEC_\Box by $\Lambda_{\text{Enkr},\Box}$. That is, $\Lambda_{\text{Enkr},\Box}$ is the logic on \mathcal{L}_\Box defined by all propositional tautologies together with the axioms $K_G, 4, L, D_G, K_D, D_D, GO, GD, \text{MAX}, \text{ND}, K_\Box$ and the rules $\text{INT}_O, \text{NEC}_D, \text{NEC}_G$ and NEC_\Box . See Table 1 for an overview. As we will see below, ND covers the full interaction between \Box and the other operators.

Having the expressive resources of $\Lambda_{\text{Enkr},\Box}$, we are finally in a position to show that goals satisfy STRONG CONSISTENCY, as desired:

LEMMA 7.6. *Let $\Lambda \subseteq \mathcal{L}_2$ be consistent, and let $S \subseteq \{\varphi \in \mathcal{L}_1 \mid \text{Goal}\varphi \in \Lambda\}$ be finite. Then $\{\text{Goal}\varphi \in \Lambda\} \vdash_{\Lambda_{\text{Enkr},\Box}} \Diamond \bigwedge_{\varphi \in S} \varphi$.*

PROOF. Assume $\bigwedge_{\varphi \in S} \text{Goal}\varphi$. By iterated application of GD, we can derive $\bigwedge_{\text{Goal}\varphi \in S} D\varphi$. By K_D and NEC_D this implies $D \bigwedge_{\text{Goal}\varphi \in S} \varphi$. By K_D and D_D this implies $\neg D \neg \bigwedge_{\text{Goal}\varphi \in S} \varphi$. The counterpositive of K_D then allows us to derive $\Diamond \bigwedge_{\text{Goal}\varphi \in S} \varphi$. ■

An immediate consequence is:

$$\vdash \text{Goal}\varphi \rightarrow \Diamond\varphi$$

We turn now to a characterization result for $\Lambda_{\text{Enkr},\Box}$:

THEOREM 7.7. *Assume At is infinite. Then the logic $\Lambda_{\text{Enkr},\Box}$ is sound and weakly complete with respect to the class of enkratic models.*

The proof can be found in the appendix. Note that a strengthening of this result is not valid. Unlike Λ_{Enkr} , the extended logic $\Lambda_{\text{Enkr},\Box}$ is only weakly complete with respect to the class of enkratic trees, at least if At is infinite.¹⁸

¹⁸To see this take some $\varphi \in \mathcal{L}_1$ that is neither a tautology nor a contradiction. Then the set $\{\Diamond\varphi \wedge D\neg\varphi\} \cup \{\neg O\psi \mid \psi \in \mathcal{L}_1\}$ is $\Lambda_{\text{Enkr},\Box}$ consistent. In fact, every finite subset thereof

7.2. Back to Challenge I: From Inconsistent Oughts to Consistent Goals

This is the right moment to return to the first two of the three challenges posed in Section 3. We consider them in turn. The first challenge concerned the potential tension between ENKRASIA and the two principles of INTERNAL and STRONG CONSISTENCY of goals in plans. Within the current framework, the tension was solved by weakening ENKRASIA. Both INTERNAL and STRONG CONSISTENCY of goals are logically valid principles, while ENKRASIA is only valid *within bounds*. INTERNAL and STRONG CONSISTENCY are, in fact, the only bounds to ENKRASIA's validity. While basic oughts are possibly inconsistent, and hence not *all* basic oughts can correspond to goals, the agent's set of goals is guaranteed to be a maximally consistent subset of her basic oughts. That is, the agent validates as many instances of ENKRASIA as is possible without violating INTERNAL and STRONG CONSISTENCY.

Note that the bounds imposed on ENKRASIA do not, as such, fully determine for *which* basic oughts the principle is valid. There might be more than one maximally consistent subset of the agent's basic oughts, hence additional choices are necessary. To this end, the framework incorporates a selection mechanism, fueled by the agent's priority ordering \succ_O . This aspect is, however, less central for the resulting logic. Any alternative selection mechanism would validate the same logical principles.

To illustrate the choices made, let us return to **Example 3.1**: Suppose I believe I ought to repay 10 euro to my friend Ann. I also believe I ought to go to the movies with Barbara. However, money is scarce, and I believe it is impossible for me to do both.

Put formally, the basic oughts of Example 3.1 are $O(Fr)$ and $O(Fm)$ (we assume, for the sake of illustration, that no other basic oughts are in play). If ENKRASIA were applied unrestrictedly, it would follow that $Goal(Fr)$ and $Goal(Fm)$ which, given that $\neg\Diamond(Fr \wedge Fm)$, would violate STRONG CONSISTENCY. Hence, ENKRASIA can only be applied to a maximally consistent subset of those basic oughts, i.e., either to $\{OFr\}$ or to $\{OFm\}$. Which one of the two depends on the lexicographic order induced by \succ_O . Suppose I believe that settling my debt with Ann takes precedence over going to the movies with Barbara, i.e., $Fr \succ_O Fm$. It follows that ENKRASIA applies only to $\{OFr\}$. The only goal derived is $Goal(Fr)$, and since $\Diamond(Fr)$ holds, no violation of INTERNAL or STRONG CONSISTENCY occurs.

Footnote 18 continued

is realizable in an enkratic-model. However, for a enkratic-model \mathcal{M} to satisfy $\{\neg O\psi \mid \psi \in \mathcal{L}_1\}$ we need that $N_O = \emptyset$. This, however, implies that the admissible subtree is all of \mathcal{T} , which yields that $\mathcal{M}, t_0 \models D\neg\varphi$ iff $\mathcal{M}, t_0 \not\models \Diamond\varphi$, i.e., it is impossible that $\mathcal{M}, t_0 \not\models \Diamond\varphi \wedge D\neg\varphi$.

7.3. Back to Challenge II: Basic and Derived Oughts

The second challenge asked to distinguish valid from invalid logical inferences about oughts. Our focus was specifically on the notorious principle of deontic closure. Even if one accepts that deontic closure is not generally valid, we have argued that an outright rejection of the principle is a too strong, and ultimately unsatisfying, solution. The challenge, hence, is to provide a deontic logic that is thick enough to license those instances of deontic closure that are unproblematic.

A central step towards meeting this challenge was the distinction between two types of oughts, basic and derived. Building on this distinction, we can illustrate the main characteristics of those deontic inferences that are valid in our logics Λ_{Enkr} and $\Lambda_{Enkr,\square}$. Let us begin by limiting the possible conclusions derivable from valid deontic inferences. Leaving aside axiom GO and INT_O , the logics Λ_{Enkr} and $\Lambda_{Enkr,\square}$ can only produce derived oughts as conclusions. In valid instances of deontic closure, hence, the ought inferred as a conclusion is a derived ought D in our sense.

More positively, valid deontic inferences are of the following kinds. First, the axiom GD allows us to infer a derived ought $D\varphi$ from a corresponding basic ought, provided that $Goal\varphi$ also holds. This reflects the idea that derived oughts are necessary conditions for the fulfillment of goals. Second, since the derived ought operator D is a normal modality, classical reasoning within the scope of derived oughts is a valid mode of inference. Thus, new derived oughts can be derived by standard modal reasoning from old ones. Finally, within the extended logic $\Lambda_{Enkr,\square}$, axiom ND can be used to infer derived oughts also from global facts about the space of available options.

To illustrate the strength of this approach, let us recall the main lines of **Example 3.2**: Suppose that I ought to repay Ann 10 euro. Moreover, I believe that unless I refrain from going to the movies it is impossible to repay Ann. So, I conclude, I ought not go to the movies.

We have called the inference in the above example *practical inference*, and argued that valid practical inferences move from basic oughts to derived oughts. In the specific case of Example 3.2, practical inference moves from $O(Fp)$, indicating the basic ought to repay Ann (i.e., once), to $D(G\neg m)$, indicating the derived ought to refrain (i.e., always) from going to the movies. It is a characteristic of our approach that derived oughts indicate the necessary conditions for the fulfillment of the agent's goals. To derive that $D(G\neg m)$ we therefore need to require that repaying Ann is in fact a goal in the agent's plan. With this in place, we can apply the inference rules described in Section 7 to derive the desired conclusion:

(i)	$O(Fr)$	(P1)
(ii)	$\Box(Fr \rightarrow G\neg c)$	(P2)
(iii)	$Goal(Fr)$	(P3)
(iv)	$Goal(Fr) \rightarrow D(Fr)$	Axiom GD
(v)	$D(Fr)$	From (iii), (iv) and Modus Ponens
(vi)	$D(Fr \rightarrow G\neg c)$	From (ii) and ND
(vii)	$D(G\neg c)$	From (v)–(vi) and K_D and NEC_D

Let us conclude with some observations about Ross' Paradox. The distinction between basic and derived oughts allows us to disentangle different readings of Ross' Paradox. Some of these are problematic, others in fact are not. Suppose we start from a basic ought to mail the letter. From such a premise, the logics Λ_{Enkr} and $\Lambda_{Enkr, \Box}$ allow us to infer at best a *derived* ought to mail or burn the letter.¹⁹ Such an inference is not paradoxical. Being a derived ought, mailing or burning the letter is *not* an ought to which ENKRASIA may apply: It cannot become a goal in its own right. Such a derived ought merely describes a necessary condition for the fulfillment of the goal of mailing the letter, not a sufficient one. In fact, burning the letter is not an admissible option (i.e., states of affairs in which the letter is burnt lie outside the admissible subtree); hence, also a derived ought *not* to burn the letter can be inferred.

8. Dynamics

Finally, we turn towards the dynamics of basic oughts, goals and derived oughts. More precisely, we study how goals and oughts change when the agent receives new information that impacts her mental picture of the world and the options available to her. In general, informational changes may trigger significant replanning, as the agent's original options might not be admissible anymore or new, better options have become available.

Our focus here is exclusively on what we have called *practical dynamics*. Crucially, the dynamics studied here are not fueled by time progression and the according changes to the set of options available. For the present exposition, we assume the agent to rest in moment t_0 , i.e., she has not yet begun to put her plans into action. Even before beginning to act, the agent might receive epistemic updates that change her perception of available future courses of events [2, 12]. Let us stipulate that this information is purely descriptive: The agent does not receive any information that leads her to

¹⁹And even that only if also the goal to mail the letter is present.

adapt or discard any basic oughts. Rather, updates may only concern further available courses of events she had not yet considered, or that certain options she had considered are, in fact, not available. Crucially, leaving the set of basic oughts intact does not entail that the agent's set of goals remains unchanged. Which of the agent's basic oughts translate into goals precisely depends upon whether they are satisfiable in a given situation and, more generally, which sets of basic oughts are jointly satisfiable. In the following, we provide a dynamic account of how an agent's goals and derived oughts change when the available courses of events do. Let us illustrate this with the following example, which brings together our previous Examples 3.3 and 3.4:

EXAMPLE 8.1. Suppose that I ought to repay my friend Ann 10 euro today (basic ought), and that is a goal in my plan. I start to plan my day accordingly. I believe I can get to Ann's house only by bus or car, so it follows that I ought to take the bus or the car (derived ought). Moreover, since my money is scarce, it follows that I ought not go to the movies today (derived ought), although I would really like to. In fact, I have even promised my friend Barbara to go to the movies with her (basic ought), but repaying my debt takes precedence (priority relation). Consider the following updates.

Update 1 Suppose I learn that the car has a dead battery, and I cannot fix the problem on time to reach Ann's house to give her the 10 euro. So, I conclude I ought to take the bus and replan my day accordingly.

Update 2 Suppose that, later, I learn that I can simply walk to Ann's (her place is unexpectedly quite close). Hence, now there are again two ways in which I can reach her house: By walking or by bus. I conclude it is no longer true that I ought to take the bus.

Update 3 Finally, suppose that I find some extra money at home: It is no longer true that I do not have enough money to repay Ann and go to the movies. I can do both. Hence, it ceases to be true that I ought not go to the movies. In fact, as I had promised my friend Barbara to accompany her, it follows that I now entertain the goal of going to the movies.

Figure 2 illustrates this situation: The top left corner shows the initial tree, the top right corner the tree after learning that the car's battery is dead. The results of the following two updates are depicted in the bottom row.

Let us make things formal. In this section, all trees are assumed discrete. More precisely we assume—(in line with much work in computer science,

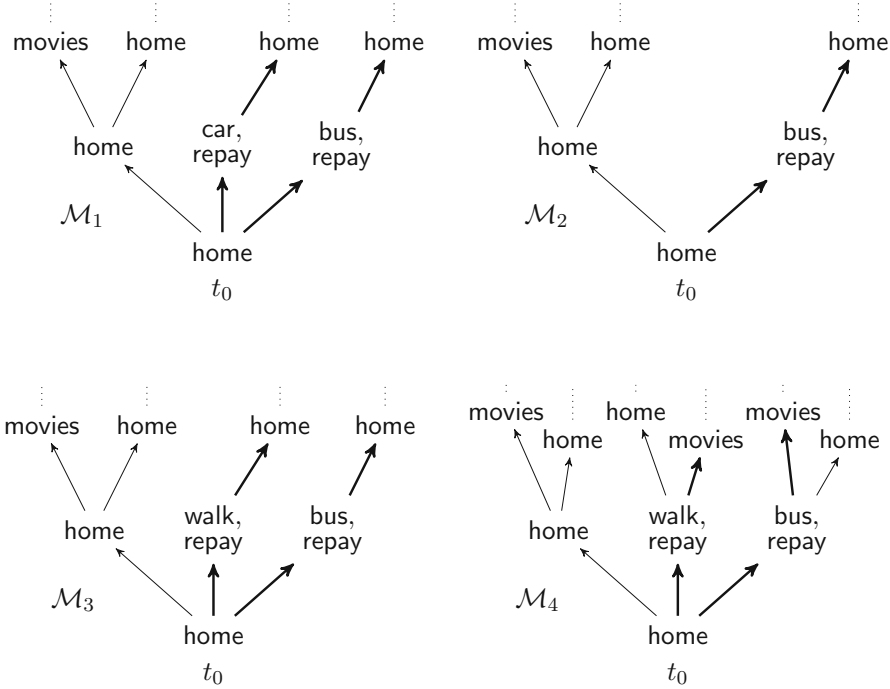


Figure 2. Four stages of planning about going to the movies and repaying money, indicated by the atoms *replay* and *movies*. In all models we have $F_{\text{replay}} \succ_O F_{\text{movies}}$ and $N_O = \{(\llbracket F_{\text{replay}} \rrbracket^{t_0}, F_{\text{replay}}), (\llbracket F_{\text{movies}} \rrbracket^{t_0}, F_{\text{movies}})\}$. Bold lines denote each model’s admissible subtree.

e.g., Ciuni and Zanardo [9])—that every history is isomorphic to the natural numbers, i.e., it can be written as $h = t_0 \prec_{\mathcal{T}} t_1 \prec_{\mathcal{T}} t_2 \prec_{\mathcal{T}} \dots$.

For a tree $\mathcal{T} = \langle T, \prec_{\mathcal{T}} \rangle$ let \prec_{im} be the **immediate predecessor relation**, e.g., $x \prec_{im} y$ iff $x \prec_{\mathcal{T}} y$ and there is no $z \in T$ with $x \prec_{\mathcal{T}} z \prec_{\mathcal{T}} y$. Note that $\prec_{\mathcal{T}}$ and \prec_{im} are in a tight relationship. As just shown, \prec_{im} is definable from $\prec_{\mathcal{T}}$. Under our assumption that every history is isomorphic to the natural numbers, the converse is also true: $\prec_{\mathcal{T}}$ is the transitive closure of \prec_{im} . Hence, providing an enkratic model $\mathcal{M} = \langle \mathbf{T}, t_0, v, N_o, \succ_O \rangle$ is equivalent to providing an **extended enkratic model** $\mathcal{M} = \langle \mathbf{T}, t_0, v, N_o, \succ_O, \prec_{im} \rangle$ that includes the relation \prec_{im} . We will make use of this property later. To technically define the dynamics of models, we refer to the concept of product updates with postconditions. See van Ditmarsch et al. [12] and Baltag et al. [2] for some technical background.

DEFINITION 8.2. A **(practical) update model** $\mathcal{E} = \langle S, s_0, R_{ims}, pre, post \rangle$ consists of a set of states S with $s_0 \in S$, a relation $R_{ims} \subseteq S \times S$, and maps $pre : S \rightarrow \mathcal{L}_0$ and $post : S \rightarrow \mathcal{P}(\text{At} \times \{\top, \perp\})$ such that if $(p, \top) \in post(s)$ then $(p, \perp) \notin post(s)$.

The attribute *practical* refers to the fact that update models do not introduce or revoke any basic oughts the agent is exposed to. Rather these models merely change the tree of possible future histories. Intuitively, S is a set of possible states with a temporal relation R_{ims} on it, similar to the relation \prec_{im} on T . Each possible state or event s in S can match and modify moments t in T . However, the matching might be subjected to additional conditions to be met by t . These conditions are recorded in $pre(s)$. Finally, a state of the update model might prescribe a change to atomic valuation at moment t . This change is represented by the postcondition $post(s)$, marking when a valuation should be forced true (i.e., $(p, \top) \in post(s)$) or false (i.e., $(p, \perp) \in post(s)$).

For the present purpose, we make two additional assumptions on the update model. First, we demand the transitive closure R of R_{ims} to be a discrete tree order on S with root s_0 such that R_{ims} is the immediate predecessor relation of R . Second, we require that for any $s \in S$ the set $\{\neg pre(t) \mid sR_{ims}t\}$ is Λ_{temp} inconsistent. In other words: For any formula $\psi \in \mathcal{L}_0$ that is not a Λ_{temp} contradiction, there is some successor t of s such that $pre(t)$ is logically compatible with ψ .

The second of the above assumptions, that the set $\{\neg pre(t) \mid sR_{ims}t\}$ is Λ_{temp} inconsistent for any $s \in S$, is non-standard. In fact, this assumption precludes classic approaches to public announcements or, more generally, the deletion of possible worlds. The condition is needed to ensure that the non-terminality axiom D_G continues to hold in the updated model. As will become clear in the formal treatment of Example 8.1, this restriction is far less severe than it may seem at first sight. Briefly, many cases of deletion can be mimicked by transforming superfluous worlds using postconditions.

DEFINITION 8.3. Let $\mathcal{M} = \langle \mathbf{T}, t_0, v, N_o, \succ_o, \prec_{im} \rangle$ be an extended enkratic model and $\mathcal{E} = \langle S, s_0, R_{ims}, pre, post \rangle$ be a practical update model. The **product update of \mathcal{M} with \mathcal{E}** , denoted by $\mathcal{M} \otimes \mathcal{E}$, is the extended enkratic model $\langle \langle T \otimes S, \prec'_T \rangle, (t_0, s_0), v', N'_o, \succ'_o, \prec'_{im} \rangle$ defined as follows

- $T \otimes S = \{(t, s) \in T \times S \mid \mathcal{M}, t \models pre(s)\}$
- Define a relation \prec'_{im} on $T \otimes S$ as $(t, s) \prec'_{im} (t', s')$ iff $t \prec_{im} t'$ and $sR_{ims}s'$.
The relation \prec'_T is then the transitive closure of \prec'_{im}

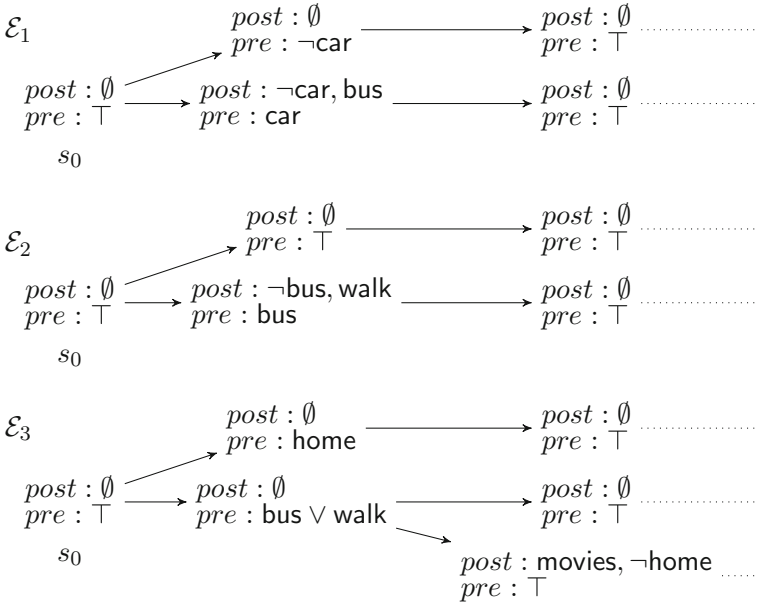


Figure 3. Practical update models corresponding to the three updating steps from Example 8.1

- The valuation $v' : At \rightarrow \mathcal{P}(T \otimes S)$ is defined by $(t, s) \in v'(p)$ if either
 - $(p, \top) \in post(s)$ or
 - $t \in v(p)$ and $(p, \perp) \notin post(s)$.
- $(\llbracket \varphi \rrbracket_{\mathcal{M} \otimes \mathcal{E}}^{(t_0, s_0)}, \varphi) \in N_O^{\mathcal{M} \otimes \mathcal{E}}$ iff $(\llbracket \varphi \rrbracket_{\mathcal{M}}^{t_0}, \varphi) \in N_O^{\mathcal{M}}$.
- $\succ_O^{\mathcal{M} \otimes \mathcal{E}} = \succ_O^{\mathcal{M}}$.

To begin with, we note that $\mathcal{M} \otimes \mathcal{E}$ is indeed an *enkratic* model. The proof of the following lemma can be found in the appendix, along with all other proofs of this section.

LEMMA 8.4. *Let \mathcal{M} be an extended enkratic model and \mathcal{E} a practical update model. Then $\langle (T \otimes S, \prec'_{\mathcal{T}}), (t_0, s_0), v', N'_O, \succ'_O, \prec'_{im} \rangle$ is an extended discrete enkratic model*

To demonstrate the versatility of this approach, we show how all three updating steps in Example 8.1 can be represented with update models. The first update model \mathcal{E}_1 in Figure 3 corresponds to learning that the car is not available. Note that by the second additional assumption on update models, the set $\{\neg pre(s') \mid sR_{ims}s'\}$ has to be inconsistent for any $s \in S$. We hence cannot simply delete the car worlds, but need to replace going by

car with something else, in this case going by bus. The next update model \mathcal{E}_2 corresponds to learning that I could walk to my friend's house. Here, a copy of the going-by-bus world is transformed into a walking-world. Finally, \mathcal{E}_3 corresponds to learning that I have sufficient money to see the movies even after repaying my friend.

The three update models displayed in Figure 3 generate the sequence of models $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4$ depicted in Figure 2. More precisely, we have that $\mathcal{M}_2 \otimes \mathcal{E}_2 = \mathcal{M}_3$ and $\mathcal{M}_3 \otimes \mathcal{E}_3 = \mathcal{M}_4$. For the transition from \mathcal{M}_1 to \mathcal{M}_2 this is not fully true: $\mathcal{M}_1 \otimes \mathcal{E}_1$ is not the same as \mathcal{M}_2 , since the former has two duplicate branches of going by bus. However, as \mathcal{M}_2 can be gained from $\mathcal{M}_1 \otimes \mathcal{E}_1$ by removing one of these duplicate branch, both models are logically equivalent.

8.1. Back to Challenge III: Dynamic Conditions

Generalizing from the previous examples, we show several general results that illustrate the complex relationship between updates, goals and derived oughts. Practical update models, despite not changing the set of basic oughts an agent is exposed to, can have intricate and non-monotonic effects on the agent's goals or derived oughts. We show that the *agent's set of goals need not necessarily grow when her available options grow, and it need not shrink if her options shrink*.

To formulate the following results, let $Hist(\mathcal{T})$ denote the set of histories of a tree \mathcal{T} . For a given situation \mathcal{M} , we call practical update model \mathcal{E} a **restriction** if $Hist(\mathcal{M} \otimes \mathcal{E}) \subset Hist(\mathcal{M})$ and an **expansion** if $Hist(\mathcal{M}) \subset Hist(\mathcal{M} \otimes \mathcal{E})$. Note that the first update in Example 8.1 is a restriction, while the second and third update exemplify expansions. We show two structural results about restrictions and expansions and how these impact the agent's goals. To show these results, we first introduce some notation. In the rest of this section, $n_G^{\mathcal{M}}$ denotes the goal formulas an agent pursues in model \mathcal{M} . Formally $n_G^{\mathcal{M}} = \{\varphi \mid \mathcal{M}, t_0 \models Goal\varphi\}$.

LEMMA 8.5. *Let \mathcal{M} be an enkratic model where for each $(\llbracket\varphi\rrbracket^{t_0}, \varphi)$ and $(\llbracket\psi\rrbracket^{t_0}, \psi)$ in N_O with $\psi \succ_O \varphi$ it holds that $\psi \in n_G^{\mathcal{M}}$ whenever $\varphi \in n_G^{\mathcal{M}}$. Let \mathcal{E} be an expansion of \mathcal{M} . Then $n_G^{\mathcal{M}} \subseteq n_G^{\mathcal{M} \otimes \mathcal{E}}$.*

Lemma 8.5 identifies a condition under which the agent's set of goals increases if new options become available to her. This additional condition is crucial. In general, an agent might drop some of her goals when new options become available. The following example provides an enkratic model \mathcal{M} and an expansion \mathcal{E} of \mathcal{M} such that $n_G^{\mathcal{M}} \not\subseteq n_G^{\mathcal{M} \otimes \mathcal{E}}$.




Figure 4. An expansion \mathcal{M}' of \mathcal{M} that does not \subseteq -increase the set of goals. Bold lines indicate the admissible subtree.

EXAMPLE 8.6. Consider the enkratic models \mathcal{M} , \mathcal{M}' displayed in Figure 4. In both models we set $N_O = \{(\llbracket Fp \rrbracket^{t_0}, Fp), (\llbracket Fq \rrbracket^{t_0}, Fq), (\llbracket Gr \rrbracket^{t_0}, Gr)\}$ and $Fp \succ_O Fq \succ_O Gr$. Clearly, $Hist(\mathcal{M}) \subseteq Hist(\mathcal{M}')$ and there is an update model \mathcal{E} such that $\mathcal{M}' = \mathcal{M} \otimes \mathcal{E}$. Hence \mathcal{M}' is an expansion of \mathcal{M} . In \mathcal{M} we have $\mathcal{M}, t_0 \models GoalFp$ and $\mathcal{M}, t_0 \models GoalGr$ but $\mathcal{M}, t_0 \not\models GoalFq$. In the expansion \mathcal{M}' , on the other hand, we have $\mathcal{M}', t'_0 \models GoalFp \wedge GoalFq$ but $\mathcal{M}', t'_0 \not\models GoalGr$. Hence $n_G^{\mathcal{M}} \not\subseteq n_G^{\mathcal{M} \otimes \mathcal{E}}$.

Next, we turn to restrictions. Here, we show that an agent's goal set cannot grow as some of her options are removed.

LEMMA 8.7. *Let \mathcal{M} be an enkratic model and let $\mathcal{M} \otimes \mathcal{E}$ be a restriction of \mathcal{M} . Then $n_G^{\mathcal{M}} \not\subseteq n_G^{\mathcal{M} \otimes \mathcal{E}}$.*

This does not imply that a restriction cannot give rise to new goals. When certain goals become unreachable, other basic oughts the agent had discarded before may move into focus. In formal terms: It is, in general, not true that $n_G^{\mathcal{M} \otimes \mathcal{E}} \subseteq n_G^{\mathcal{M}}$. This is the subject of the following example.

EXAMPLE 8.8. We use the same construction as in Example 8.6. Again, consider the enkratic models \mathcal{M} , \mathcal{M}' displayed in Figure 4. Again, we set $Fp \succ_O Fq \succ_O Gr$ and $N_O = \{(\llbracket Fp \rrbracket^{t_0}, Fp), (\llbracket Fq \rrbracket^{t_0}, Fq), (\llbracket Gr \rrbracket^{t_0}, Gr)\}$ in both models. Evidently, $Hist(\mathcal{M}) \subset Hist(\mathcal{M}')$ and there is a practical update model \mathcal{E} such that \mathcal{M} is equivalent to $\mathcal{M}' \otimes \mathcal{E}$. Hence \mathcal{M} is a restriction of \mathcal{M}' . Then the same argument as above shows that $n_G^{\mathcal{M}' \otimes \mathcal{E}} \not\subseteq n_G^{\mathcal{M}'}$.

The two examples above illustrate that purely practical updates can have complex and non-monotonic effects on the agent's goals or derived oughts.

In particular, even practical updates merely expanding the set of available histories may trigger the agent to drop certain goals of hers. Partially, such phenomena are due to the exact choice of updating rule. In the present framework, the agent calculates her set of goals from scratch after each update, picking as the new goal set the maximally consistent subset of basic oughts that is maximal in the lexicographic order induced by \succ_O .

An alternative updating policy might opt for minimal changes instead, adopting as the updated set of goals some maximally consistent subset that differs minimally from the goal set the agent pursued before the update. While immune to the non-monotonicity described above, such minimal change rule may trigger a different type of non-conservativeness. Take for instance a two step update where an agent is first informed that some history h of a given enkratic tree model \mathcal{M} is not available, followed by a second update indicating that the first information was wrong and h is, in fact available. After executing both updates, the tree of available options is exactly as it was in the starting model \mathcal{M} . With the original updating policy described above, the set of goals after both updates is also the same as the initial goal set. However, this would, in general, cease to hold if we followed a minimal change rule instead.²⁰ We take this to illustrate that the existence of complex interaction patterns between changes of available histories and the set of goals pursued does not hinge on the exact updating policy, but is a general fact of planning when exposed to possibly incompatible sets of basic oughts.

9. Conclusion and Open Ends

In this paper we pursued two aims: Analyzing the logical structure of ENKRASIA, and addressing some of the implications ENKRASIA has, in combination with certain other principles of practical rationality, for deontic logic.

As for the first aim, we have argued that ENKRASIA is a principle of bounded validity. Goals are subjected to two requirements of INTERNAL and

²⁰To see that this holds true for *any* minimal change updating rule, consider an enkratic model \mathcal{M} consisting of three branches f, g and h . Moreover, assume N_O to contain three basic oughts O_1 – O_3 which are satisfied in the subtrees $\{f, g\}$, $\{f, h\}$ and $\{g, h\}$ respectively. Maximally consistent subsets hence are $\{O_1, O_2\}$, $\{O_1, O_3\}$ and $\{O_2, O_3\}$. Assume wlog that $\{O_1, O_2\}$ is adopted as set of goals in \mathcal{M} . After removing branch f , this set is not consistent anymore, now only $\{O_1, O_3\}$ and $\{O_2, O_3\}$ are maximally consistent. One of these is selected as new set of goals, wlog $\{O_1, O_3\}$. Since $\{O_1, O_3\}$ remains maximally consistent after adding f again, any minimal change rule must retain it as set of goals. In particular, the set of goals after removing and re-adding f is different from before.

STRONG CONSISTENCY which basic oughts are not. Both of these conditions set boundaries for the translation of basic oughts into goals. ENKRASIA, then, is only valid as far as it does not conflict with either requirement of consistency.

In relation to the second aim in this paper, we have elaborated on the distinction between basic and derived oughts. This distinction allows us to represent and indeed validate, e.g., a plausible reading of practical inference without generating an unrestricted validity of deontic closure. In fact, by restricting the conclusions of deontic closure to derived oughts, many of the paradoxical implications usually associated with deontic closure no longer obtain. Differences between basic and derived oughts surface, as shown, in their interaction with goals and their kinematics within a dynamical logical framework.

Related Approaches To the best of our knowledge, ENKRASIA has not been previously investigated explicitly from a logical perspective. There are, however, a variety of logical frameworks that deal with notions pertaining to practical rationality. We point to some of these. The analysis presented here is linked to the logical tradition of interpreting intentions according to Bratman's [4] planning theory. Related works—mainly focusing on normal modal logic—include Cohen and Levesque [10] and Lorini and Herzig [22]. The latter paper also discusses practical inference, though in the context of forming instrumental intentions rather than derived oughts.

For what concerns the dynamics of intentions and plans, related work includes van der Hoek et al. [16], with a focus on the operation of deleting plans. A second reference is Icard et al. [19], who provide an account of intention revision based on AGM theory. Further complementary work includes Craven and Sergot's [11] account of permitted and obligatory actions within transition systems, and Broersen et al.'s [5] syntactic, default-based approach on conflicts between beliefs, obligations, intentions and desires.

While various frameworks address the relationship between obligations, plans and intentions, combinability across approaches is sometimes hampered by the fact that these may refer to different notions. Consider, for instance, Broersen et al.'s [5] BOID framework on beliefs, obligations, intentions and desires mentioned above. BOID distinguishes, *inter alia*, between obligations (be they believed by the agent or not) and intentions, denoting actions the agent plans on doing. In the present framework, in contrast, basic oughts are assumed believed by the agent. Goals in a plan, moreover, correspond to basic oughts the agent decided to pursue. These form a category between BOID's obligations and intentions.

Future Directions The present approach fits in with a larger project on the logic of oughts in the context of practical rationality. The analysis and the framework presented here can be extended in different directions. We mention two.

One possible extension of our framework pertains to the relation between ENKRASIA and permissions. We have seen that the relation between oughts and goals is not unidirectional. Oughts translate into goals, and further, derived oughts can be generated from a given set of goals. However, it seems that not only oughts, but also permissions can be derived from what the agent is committed to bring about. The admissible subtree, denoting the intersection of all goals pursued by the agent, can be thought of as a weakest permission, i.e., as the largest subtree the agent is permitted to arrive in, given her commitments. Naturally, stronger permissions may also hold, permitting the agent to arrive in *any* subtree of the admissible tree. A classic reference on this is Anglberger et al. [1].

A second possible extension relates to the way in which possible future courses of events are represented. In the current framework we use a tree-structure to represent the agent's choice options among future unfoldings of events. This picture can be refined in various ways. We may, for instance, restrict the agent's ability to fully select the future course of actions. That is, agents may no longer be able to pick a specific branch, but merely to choose some subtree to stay within. Conceptually, this might require a refinement of the principle of STRONG CONSISTENCY, e.g., by demanding that if $Goal\varphi$ then the agent believes she has an available choice option that *guarantees* φ . Formally, this would amount to having equivalence classes of histories (representing choice uncertainty) paired with a quantification over such classes. Similar topics have been investigated by Horty [17] and Ciuni and Zanardo [9]. The results presented in these works could form a fruitful starting point for expanding the logic of ENKRASIA.

Acknowledgements. Authors are ordered alphabetically. We would like to thank John Horty, Olivier Roy, Igor Sedlár, Frederik Van De Putte, workshop audiences at DEON 2018 Utrecht and in Lublin, Bayreuth and Prague, as well as three anonymous reviewers for valuable feedback and suggestions. The work of DK was partially supported by the Deutsche Forschungsgemeinschaft (DFG) and Agence Nationale de la Recherche (ANR) as part of the joint project Collective Attitude Formation [RO 4548/8-1], by DFG and Grantová Agentura České Republiky (GAČR) through the joint project

From Shared Evidence to Group Attitudes [RO 4548/6-1] and by the National Science Foundation of China as part of the project Logics of Information Flow in Social Networks [17ZDA026].

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

Appendix: Proofs

Before we can prove Theorem 7.5, we need the following auxiliary lemma:

LEMMA A.1. *Assume $\Lambda \subseteq \mathcal{L}_2$ is maximally Λ_{Enkr} consistent. Let $\Lambda^D = \{\varphi \in \mathcal{L}_1 \mid D\varphi \in \Lambda\}$, let $\Lambda^{-D} = \{\varphi \in \mathcal{L}_1 \mid \neg D\varphi \in \Lambda\}$ and let $\Lambda^{lit} \subseteq \Lambda$ be the set of all literals, i.e., atoms and negated atoms, occurring in Λ . Then for every $\psi \in \Lambda^{-D}$, there is a linear²¹ tree $\mathcal{H}_{\neg\psi} = \langle h_{\neg\psi}, \prec_{\mathcal{H}} \rangle$ with root $t_{\neg\psi}$ such that $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \neg\psi$ and $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \varphi$ for every $\varphi \in \Lambda^D \cup \Lambda^{lit}$.*

PROOF. Let $\chi \in \Lambda^{-D}$, we will construct the desired linear tree $\mathcal{H}_{\neg\chi} = \langle h_{\neg\chi}, \prec_{\mathcal{H}} \rangle$. By axioms K_D and D_D , the set $\Lambda^D \cup \{\neg\chi\}$ is Λ_{temp} consistent. Since all $\varphi \in \Lambda^D \cup \{\neg\chi\}$ are future looking, i.e., every atom is in the scope of a modal operator, also $\Lambda^D \cup \{\neg\chi\} \cup \Lambda^{lit}$ is Λ_{temp} consistent. We can hence expand it to a maximally Λ_{temp} consistent subset $\Gamma \subseteq \mathcal{L}_0$. The rest of the proof proceeds by a classic bulldozing argument. For the sake of completeness, we hint at the details.

Let $\mathcal{M} = \langle W, R \rangle$ be the Λ_{temp} canonical model over \mathcal{L}_0 . That is, W is the set of all maximally Λ_{temp} consistent subsets of \mathcal{L}_0 and $\Theta R \Sigma$ iff $F\varphi \in \Theta$ for all $\varphi \in \mathcal{L}_0$ with $\varphi \in \Sigma$. Let $H = \{\Sigma \in \mathcal{M} \mid \Gamma R \Sigma\}$. We show that R is transitive and complete on H . Transitivity follows from the 4 axiom. For completeness let $\Theta \neq \Sigma \in H$, wlog $\Theta, \Sigma \neq \Gamma$. We have to show $\Theta R \Sigma$ or $\Sigma R \Theta$.

Pick enumerations $\varphi_0^\Theta, \varphi_1^\Theta \dots$ of Θ and $\varphi_0^\Sigma, \varphi_1^\Sigma \dots$ of Σ . For $i \in \omega$ let $\psi_i^* = \bigwedge_{j=1}^i \varphi_j^*$ for $*$ $\in \{\Theta, \Sigma\}$. Note that since $\Theta \neq \Sigma$, there are $j, k \geq 0$ such that $\varphi_j^\Theta = \neg\varphi_k^\Sigma$. Letting $i_0 = \max(j, k)$ we have that $\vdash_{\Lambda_{temp}} \neg(\psi_i^\Theta \wedge \psi_i^\Sigma)$ for all $i > i_0$. As Θ, Σ are maximally consistent we get $\psi_i^\Theta \in \Theta, \psi_i^\Sigma \in \Sigma$ for all $i \in \omega$.

By construction, we have $\Gamma R \Theta$ and $\Gamma R \Sigma$. The Truth Lemma then implies that $F\psi_i^* \in \Gamma$ for $*$ $\in \{\Theta, \Sigma\}$ and all $i \in \omega$. Hence, we have for all i that $F\psi_i^\Theta \wedge F\psi_i^\Sigma \in \Gamma$. By L, this implies that $F(\psi_i^\Theta \wedge \psi_i^\Sigma) \vee F(\psi_i^\Theta \wedge F\psi_i^\Sigma) \vee F(\psi_i^\Sigma \wedge F\psi_i^\Theta) \in \Gamma$. In particular Γ contains either $F(\psi_i^\Theta \wedge \psi_i^\Sigma)$ for infinitely many i , $F(\psi_i^\Theta \wedge F\psi_i^\Sigma)$ for infinitely many indices i or $F(\psi_i^\Sigma \wedge \psi_i^\Theta)$ for infinitely many i . Since $\vdash_{\Lambda_{temp}} \neg(\psi_i^\Theta \wedge F\psi_i^\Theta)$ for all but finitely many i , the first case is impossible. We treat the case where Γ contains $F(\psi_i^\Theta \wedge F\psi_i^\Sigma)$ for infinitely many i , the other case being similar. We will show that $F\psi_i^\Sigma \in \Theta$ for all $i \in \omega$. By K_G and the construction of the ψ_i^Σ this entails that $F\varphi \in \Theta$ for all $\varphi \in \Sigma$, which, together with the definition of R implies that $\Theta R \Sigma$.

²¹I.e., a tree where $\prec_{\mathcal{H}}$ is a linear order.

To see that $F\psi_i^\Sigma \in \Theta$ for all $i \in \omega$, assume for a contradiction that this is false, i.e., that there is some i_m with $F\psi_{i_m}^\Sigma \notin \Theta$. By maximal consistency, this implies that $\neg F\psi_{i_m}^\Sigma \in \Theta$. Hence, there is some $\psi_j^\Theta \in \Theta$ with $\vdash_{\Lambda_{temp}} \neg(\psi_j^\Theta \wedge F\psi_{i_m}^\Sigma)$. By construction of the ψ_j , this implies $\vdash_{\Lambda_{temp}} \neg(\psi_{j'}^\Theta \wedge F\psi_{j'}^\Sigma)$ for all $j' > \max(i_m, j)$. In particular, since Γ contains $F(\psi_i^\Theta \wedge F\psi_i^\Sigma)$ for infinitely many i , there is some $j_0 > \max(i_m, j)$ with $F(\psi_{j_0}^\Theta \wedge F\psi_{j_0}^\Sigma) \in \Gamma$ but $\vdash_{\Lambda_{temp}} \neg(\psi_{j_0}^\Theta \wedge F\psi_{j_0}^\Sigma)$. This, in connection with D_G contradicts the consistency of Γ . Hence the assumption was false and we obtain that $F\psi_i^\Sigma \in \Theta$ for all $i \in \omega$.

To finish the proof, we need to ensure that R is a linear order on H . This is, in general, not true. However, by a classic bulldozing argument (cf. [30]), we can transform H into a linearly ordered tree $\mathcal{H} = \langle h, \prec_{\mathcal{H}} \rangle$ with $\prec_{\mathcal{H}}$ minimal element Γ such that $\mathcal{H}, \Gamma \models \varphi \Leftrightarrow \varphi \in \Gamma$. Renaming Γ to $t_{-\chi}$, h to $h_{-\chi}$ and \mathcal{H} to $\mathcal{H}_{-\chi}$ finishes the proof. ■

Now, we can finally show Theorem 7.5. The completeness direction proceeds by constructing a special tree model \mathcal{T} , with the property that for each $(\llbracket \varphi \rrbracket^{t_0} \varphi) \in N_O$ either $\llbracket \varphi \rrbracket^{t_0} = \mathcal{T}$ or $\llbracket \varphi \rrbracket^{t_0} = \emptyset$. Consequently, it will hold that $\mathcal{T}_{adm} = \mathcal{T}$. The main task of the construction, hence is to ensure that the histories of \mathcal{T} are such that $\mathcal{T} \subseteq \llbracket \varphi \rrbracket^{t_0} \Leftrightarrow D\varphi \in \Lambda$. The corresponding construction bears some resemblance to the completeness proof for ATL [14].

PROOF OF THEOREM 7.5 Soundness is trivial. For completeness, we show that every maximally Λ_{Enkr} consistent subset Λ of \mathcal{L}_2 is satisfiable in an enkratic model. Let $\Lambda^D = \{\varphi \in \mathcal{L}_1 \mid D\varphi \in \Lambda\}$, let $\Lambda^{-D} = \{\varphi \in \mathcal{L}_1 \mid \neg D\varphi \in \Lambda\}$ and let $\Lambda^{lit} \subseteq \Lambda$ be the set of all literals, i.e., atoms and negated atoms, occurring in Λ . Using Lemma A.1, we pick linear trees $\mathcal{H}_{-\psi} = \langle h_{-\psi}, \prec_{\mathcal{H}} \rangle$ with root $t_{-\psi}$ for each $\psi \in \Lambda^{-D}$ as above. Note that all $t_{-\psi}$ share the same atomic valuation, as this is completely determined by Λ^{lit} . Moreover, note that $\neg D \perp \in \Lambda$ by D_D . Hence $\perp \in \Lambda^{-D}$ and thus the set of linear trees picked is non-empty.

We have to construct an enkratic model $\mathcal{M} = \langle \mathcal{T}, t_0, v, N_o, \succ_o \rangle$. As tree \mathcal{T} we take the union of the $\mathcal{H}_{-\psi}$ where we identify all $t_{-\psi}$. Formally, for a linear tree $\mathcal{H} = \langle h, \prec_{\mathcal{H}} \rangle$, let $T_{\mathcal{H}}^>$ be the set of all moments but the first of \mathcal{H} . Let

$$T = \{t_0\} \cup \bigcup_{\psi \in B} T_{\mathcal{H}_{-\psi}}^>$$

and $\prec_{\mathcal{T}}$ the inherited tree-order, making t_0 the root. Finally, the valuation v is defined by

$$t \in v(p) \quad \text{iff} \quad \begin{cases} t = t_0 \text{ and } p \in \Lambda \\ \text{or } t \in T_{\mathcal{H}_{-\psi}}^> \text{ and } t \in v_{\mathcal{H}_{-\psi}}(p). \end{cases}$$

Finally, we define N_O by $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O$ iff $O\varphi \in \Lambda$. Moreover, we pick an arbitrary well-founded ordering \succ_o on $\{\varphi \mid (\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O\}$.

For the completeness argument, we begin with some observations about the root t_0 of this tree. First we note that, for $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O(t_0)$, we have $\llbracket \varphi \rrbracket^{t_0} = T$ if $Goal\varphi \in \Lambda$ and $\llbracket \varphi \rrbracket^{t_0} = \emptyset$ else. In the former case, we have $D\varphi \in \Lambda$ by GD.

By construction, this implies that every $\mathcal{H}_{\neg\psi}$ and hence every branch h/t_0 of \mathcal{T} satisfies $h/t_0 \models \varphi$, i.e., $\llbracket \varphi \rrbracket^{t_0} = T$, which is what had to be shown. In the other case, $Goal\varphi \notin \Lambda$, we have, by MAX, that $D\neg\varphi \in \Lambda$. Again by construction, every branch h/t_0 of \mathcal{T} satisfies $h/t_0 \models \neg\varphi$, i.e., $\llbracket \varphi \rrbracket^{t_0} = \emptyset$. The last two observations imply that

$$N_G = \{\llbracket \varphi \rrbracket^{t_0} \mid O\varphi \in \Lambda, Goal\varphi \in \Lambda\}. \quad (1)$$

In particular, by GD, $N_G(t_0) = \{\mathcal{T}\}$ if there is some $Goal\psi \in \Lambda$ and $N_G(t_0) = \emptyset$ else. In either case we have $\bigcap_{X \in N_G(t_0)} X = \mathcal{T}$.

Now we can show that our model is as desired, i.e., $\mathcal{M}, t_0 \models \varphi$ iff $\varphi \in \Lambda$. The argument is an induction over the complexity of φ . As induction base, we show the claim for φ an atom or of the form $O\psi, Goal\psi$ or $D\psi$ for some $\psi \in \mathcal{L}_1$. In the induction step we then show that if the claim holds for $\varphi_1, \varphi_2 \in \mathcal{L}_2$ then also for $\neg\varphi_1$ and $\varphi_1 \wedge \varphi_2$. This induction step is trivial. We only need to show the claim for the induction base.

If φ is atomic, the claim holds by definition of the valuation on t_0 . If φ is $O\psi$ for some $\psi \in \mathcal{L}_1$, this follows immediately from the construction of N_O . If φ is $Goal\psi$ for $\psi \in \mathcal{L}_1$ the claim follows immediately from Eq. 1.

The only non-trivial case is when φ is of the form $D\psi$ for $\psi \in \mathcal{L}_1$. For the left to right direction, assume $\mathcal{M}, t_0 \models D\psi$. We have to show $D\psi \in \Lambda$. First, note that $\mathcal{M}, t_0 \models D\psi$ implies that $\bigcap_{X \in N_G} X \subseteq \llbracket \psi \rrbracket^{t_0}$. Since for each $X \in N_G$ holds that $X = \{\mathcal{T}\}$ or $X = \emptyset$, this implies that $\mathcal{T} \subseteq \llbracket \psi \rrbracket^{t_0}$. Assume for a contradiction that $D\psi \notin \Lambda$. By maximality, this implies $\neg D\psi \in \Lambda$. By construction, there is a branch $h_{\neg\psi}$ of \mathcal{T} with $\mathcal{T}, t_0/h_{\neg\psi} \models \neg\psi$. In particular, $\mathcal{T} \not\subseteq \llbracket \psi \rrbracket^{t_0}$ which is a contradiction. Hence the assumption was false and $D\psi \in \Lambda$. For the right to left direction assume $D\psi \in \Lambda$. Recall that by construction, every branch h of \mathcal{M} satisfies $\mathcal{M}, t_0/h \models \psi$. In particular, $\bigcap_{X \in N_G} X \subseteq \llbracket \psi \rrbracket^{t_0}$ which implies that $\mathcal{M}, t_0 \models D\psi$. ■

PROOF OF THEOREM 7.7 The proof borrows heavily from the proof of Theorem 7.5. We only highlight the relevant differences. Soundness, again, is trivial. For weak completeness, we have to show that whenever there is some $\tilde{\varphi} \in \mathcal{L}_{\square}$ with $\vdash_{\neg\Lambda_{Enkr, \square}} \neg\tilde{\varphi}$, there is some enkratic model \mathcal{M}, w with $\mathcal{M}, w \models \tilde{\varphi}$. Let such $\tilde{\varphi}$ be given. Without loss of generality, we can assume that $\tilde{\varphi}$ is in disjunctive normal form, i.e.,

$$\tilde{\varphi} = \bigvee_{i=1}^n \bigwedge_{j=1}^{k_i} \chi_{i,j}$$

where each χ_i either is an atom, a negated atom, or of the form $X\varphi$ or $\neg X\varphi$ for $X \in \{O, Goal, D, \square\}$ and $\varphi \in \mathcal{L}_1$. By MAX, we have $\vdash_{\neg\Lambda_{Enkr, \square}} O\varphi \leftrightarrow (O\varphi \wedge Goal\varphi) \vee (O\varphi \wedge D\neg\varphi)$. We can hence assume, without loss of generality, that every disjunct $\bigwedge_{j=1}^{k_i} \chi_{i,j}$ of $\tilde{\varphi}$ has the property that for each $\chi_{i,j}$ of the form $O\psi$, either $Goal\psi$ or $D\neg\psi$ appears as some of the $\chi_{i,j'}$ with $j' \leq k_i$. Likewise, as $\vdash_{\neg\Lambda_{Enkr, \square}} \neg\square\varphi \leftrightarrow (\neg\square\varphi \wedge \neg D\varphi) \vee (\neg\square\varphi \wedge D\varphi)$, we can assume that for each $\chi_{i,j}$ of the form $\neg\square\psi$, either $\neg D\psi$ or $D\psi$ appears as some of the $\chi_{i,j'}$ with $j' \leq k_i$. Moreover, by GO, we can assume that for each $\chi_{i,j}$ of the form $Goal\psi$, also $O\psi$ appears as some of the $\chi_{i,j'}$ with $j' \leq k_i$. Finally, by D_D , we can assume that for each $i \leq n$ $\neg D\perp$,

appears as some $\chi_{i,j'}$ with $j' \leq k_i$. To show that $\tilde{\varphi} = \bigvee_{i=1}^n \bigwedge_{j=j'}^{k_i} \chi_i$ is satisfiable in an enkratic model, it suffices to show that one of its disjuncts is satisfiable in an enkratic model. Hence, it suffices to show the claim for $i = 1$, that is for $\tilde{\varphi}$ of the form $\bigwedge_{j=1}^k \chi_j$. Let such a φ be given and let $X = \{\chi_1, \dots, \chi_k\}$.

In a similar fashion as in Theorem 7.5, let $\Lambda^D = \{\rho \in \mathcal{L}_1 \mid D\rho \in X\}$, let $\Lambda^\square = \{\rho \in \mathcal{L}_1 \mid \square\rho \in X\}$, let $\Lambda^{-D} = \{\rho \in \mathcal{L}_1 \mid \neg D\rho \in X\}$ and let $\Lambda^{\text{At}} = \{\chi \in X \mid \chi = p \text{ or } \chi = \neg p\}$. Note that $\Lambda^{-D} \neq \emptyset$ as $\perp \in \Lambda^{-D}$. As At is infinite, we can pick an atom p_0 that does not occur in any of the formulas in X . Pick a valuation Λ^{lit} extending Λ^{At} , i.e., some maximally consistent $\Lambda^{\text{lit}} \subseteq \{p, \neg p \mid p \in \text{At}\}$ with $\Lambda^{\text{lit}} \supseteq \Lambda^{\text{At}}$ such that $\neg p_0 \in \Lambda^{\text{lit}}$. By a slight adaptation of Lemma A.1, there are linear trees $\mathcal{H}_{\neg\psi} = \langle h_{\neg\psi}, \prec_{\mathcal{H}} \rangle$, with root $t_{\neg\psi}$ for each $\psi \in \Lambda^{-D}$, such that $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \neg\psi$ and $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \chi$ for every $\chi \in \Lambda^D \cup \Lambda^\square \cup \Lambda^{\text{lit}}$. Moreover let $\Lambda^{\diamond/D} = \{\neg\rho \in \mathcal{L}_1 \mid \neg\square\rho \in X \text{ and } D\rho \in X\}$ the set of formulas that are possible without being a derived goal. By another slight adaptation of Lemma A.1, there are linear trees $\mathcal{H}'_\rho = \langle h'_\rho, \prec'_{\mathcal{H}} \rangle$ with root t'_ρ for each $\rho \in \Lambda^{\diamond/D}$ such that $\mathcal{H}'_\rho, t'_\rho \models \rho$ and $\mathcal{H}'_\rho, t'_\rho \models \chi$ for every $\chi \in \Lambda^\square \cup \Lambda^{\text{lit}}$.

Before constructing the desired enkratic model, we show the following claim: There is a formula ρ_0 such that $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \rho_0$ for all $\psi \in \Lambda^{-D}$, but $\mathcal{H}'_\psi, t'_\psi \models \neg\rho_0$ for all $\psi \in \Lambda^{\diamond/D}$. To see this, note that by definition of sets Λ^D and $\Lambda^{\diamond/D}$, every member of $\Lambda^{\diamond/D}$ is of the form $\neg\chi$ with $\chi \in \Lambda^D$. Let $\rho_0 = \bigwedge_{\neg\chi \in \Lambda^{\diamond/D}} \chi$. We hence have that $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \rho_0$ for all $\psi \in \Lambda^{-D}$, but $\mathcal{H}'_\psi, t'_\psi \models \neg\rho_0$ for all $\psi \in \Lambda^{\diamond/D}$ i.e., ρ_0 has the desired property.

Now, we construct an enkratic model $\mathcal{M} = \langle \mathcal{T}, t_0, v, N_O, \succ_O \rangle$ satisfying $\tilde{\varphi}$ as follows. We pick linear trees $\mathcal{H}_{\neg\psi}$ for each $\psi \in \Lambda^{-D}$ and \mathcal{H}'_ψ for all $\psi \in \Lambda^{\diamond/D}$ as above. As p_0 does not occur in any of the formulas in X and $\neg p_0 \in \Lambda^{\text{lit}}$ we can assume wlog that p_0 is globally false on all $\mathcal{H}_{\neg\psi}$ and \mathcal{H}'_ψ . With these, we define a tree \mathcal{T} as in the proof of Theorem 7.5. By the previous claim, there is formula ρ_0 such that $\mathcal{H}_{\neg\psi}, t_{\neg\psi} \models \rho_0$ for all $\psi \in \Lambda^{-D}$, but $\mathcal{H}'_\psi, t'_\psi \models \neg\rho_0$ for all $\psi \in \Lambda^{\diamond/D}$. Since p_0 is false everywhere, $\rho = \rho_0 \vee Gp_0$ has the same property. Moreover, since p_0 does not occur in X , we also have that $O\rho, \neg O\rho, \text{Goal}\rho, \neg\text{Goal}\rho$ are all not contained in X . Define the neighborhood N_O by $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \in N_O$ iff $O\varphi \in X$ or $\varphi = \rho$. Finally, as priority relation \succ_O , we pick any well-founded ordering on $\{\varphi \mid O\varphi \in X\} \cup \{\rho\}$ that has ρ as maximal element.

We can now show that $\mathcal{M}, t_0 \models \tilde{\varphi}$. Since $\tilde{\varphi} = \bigwedge_{\psi \in X} \psi$, it suffices to show that $\mathcal{M}, t_0 \models \psi$ for all $\psi \in X$. If ψ is an atom or negated atom, this follows immediately from the construction. The same holds if ψ is of the form $O\varphi, \neg O\varphi$ or $\square\varphi$. If ψ is of the form $\neg\square\varphi$ we have by our assumption that either $D\varphi \in X$ or $\neg D\varphi \in X$. In the first case $\neg\varphi \in \Lambda^{\diamond/D}$ and $\mathcal{H}'_{\neg\varphi}, t'_{\neg\varphi} \models \neg\varphi$ and hence $\mathcal{M}, t_0 \not\models \square\varphi$. In the second case $\varphi \in \Lambda^{-D}$ and $\mathcal{H}_{\neg\varphi}, t_{\neg\varphi} \models \neg\varphi$ witnessing again that $\mathcal{M}, t_0 \not\models \square\varphi$.

For the remaining cases, we define subtrees \mathcal{S} and \mathcal{S}' of \mathcal{T} by $\mathcal{S} = \bigcup\{\mathcal{H}_{\neg\psi} \mid \psi \in \Lambda^{-D}\}$ and $\mathcal{S}' = \bigcup\{\mathcal{H}'_\psi \mid \psi \in \Lambda^{\diamond/D}\}$. Since for each $O\psi \in X$ also $\text{Goal}\psi \in X$ or $D\neg\psi \in X$, we can use the same argument as in the previous theorem to show that for $O\psi \in X$, we have that $\mathcal{S} \subseteq \llbracket \psi \rrbracket^{t_0}$ if $\text{Goal}\psi \in X$, and $\llbracket \psi \rrbracket^{t_0} \subseteq \mathcal{S}'$ if $D\neg\psi \in X$.

Since $\llbracket \rho \rrbracket^{t_0} = S$, and ρ is the \succ_O maximal element of $\{\varphi \mid O\varphi \in X\} \cup \{\rho\}$, we get that $\llbracket \varphi \rrbracket^{t_0}$ with $O\varphi \in X$ can only be in N_G if $\llbracket \varphi \rrbracket^{t_0} \not\subseteq S'$, i.e., if $Goal\varphi \in X$. Moreover, since $S \subseteq \llbracket \varphi \rrbracket^{t_0}$ whenever $Goal\varphi \in X$, we get that $N_G = \{\llbracket \varphi \rrbracket^{t_0} \mid O\varphi, Goal\varphi \in X\}$. By our assumption that $O\varphi \in X$ whenever $Goal\varphi \in X$, this simplifies to $N_G = \{\llbracket \varphi \rrbracket^{t_0} \mid Goal\varphi \in X\}$ and hence $\bigcap_{\llbracket \varphi \rrbracket^{t_0} \in N_G} \llbracket \varphi \rrbracket^{t_0} = S$. From there, the same argument as in the previous theorem implies that $\mathcal{M}, t_0 \models \psi$ if $\psi \in X$ is of the form $Goal\varphi$ or $D\varphi$. If ψ is of the form $\neg D\varphi$, note that there is, by construction, a branch $h_{\neg\varphi}$ with $\mathcal{M}, t_0/h_{\neg\varphi} \models \neg\varphi$. Since $h_{\neg\varphi} \subseteq S$ and $\bigcap_{\llbracket \varphi \rrbracket^{t_0} \in N_G} \llbracket \varphi \rrbracket^{t_0} = S$, this implies $\mathcal{M}, t_0 \models \neg D\varphi$. Finally, if ψ is of the form $\neg Goal\varphi$, we need to distinguish whether $O\varphi \in X$ or not. If not, we have, by construction, that $(\llbracket \varphi \rrbracket^{t_0}, \varphi) \notin N_O$, which immediately implies that $\mathcal{M}, t_0 \models \neg Goal\varphi$. For the case that $O\varphi \in X$, we have by construction that also $D\neg\varphi \in X$. For this case, we have shown above that $\llbracket \varphi \rrbracket^{t_0} \subseteq S'$, which implies $\llbracket \varphi \rrbracket^{t_0} \notin N_G$. Again, we get $\mathcal{M}, t_0 \models \neg Goal\varphi$ as desired. ■

PROOF OF LEMMA 8.4 We begin with showing that \prec'_T is a tree-order on $T \otimes S$. Transitivity is immediate, as \prec'_T is transitively closed. For irreflexivity assume the contrary, i.e., assume that there is some (t', s') with $(t', s') \prec'_T (t', s')$. Hence there are $(t', s') \prec'_{im} (t_2, s_2) \prec'_{im} \dots \prec'_{im} (t_n, s_n) = (t', s')$. By definition of $\mathcal{M} \otimes \mathcal{E}$, this entails that $t' = t_1 \prec_{im} t_2 \dots \prec_{im} t_n = t'$. In particular, we get $t' \prec_T t'$, as \prec_T is the transitive closure of \prec_{im} . But this contradicts the fact that \prec_T is irreflexive.

Finally for inverse linearity, let $(t, s) \in T \otimes S$. We need to show that

$$P = \{(t', s') \in T \otimes S \mid (t', s') \prec'_T (t, s)\}$$

is linearly ordered by \prec'_T . Let $(t', s'), (\tilde{t}, \tilde{s}) \in P$ be given. We have to show that $(t', s') \prec'_T (\tilde{t}, \tilde{s})$ or $(\tilde{t}, \tilde{s}) \prec'_T (t', s')$. By assumption, we have that $t', \tilde{t} \prec_T t$. By construction and the assumption that \prec_T has the order type of the natural numbers, there is $n' \in \omega$ such that $(t', s') = (t'_0, s'_0) \prec'_T \dots \prec'_T (t'_{n'}, s'_{n'}) = (t, s)$ with $t'_0 \prec_{im} \dots \prec_{im} t'_{n'}$ and $s'_0 R_{ims} \dots R_{ims} s'_{n'}$. Likewise, there is $\tilde{n} \in \omega$ such that $(\tilde{t}, \tilde{s}) = (\tilde{t}_0, \tilde{s}_0) \prec'_T \dots \prec'_T (\tilde{t}_{\tilde{n}}, \tilde{s}_{\tilde{n}}) = (t, s)$ with $\tilde{t}_0 \prec_{im} \dots \prec_{im} \tilde{t}_{\tilde{n}}$ and $\tilde{s}_0 R_{ims} \dots R_{ims} \tilde{s}_{\tilde{n}}$. By definition of a practical update model, R_{ims} is a unique predecessor relation, i.e., $x R_{ims} z$ and $y R_{ims} z$ implies $x = y$. We hence have that $s'_{n'-1} = \tilde{s}_{\tilde{n}-1}$, $s'_{n'-2} = \tilde{s}_{\tilde{n}-2} \dots$. The same reasoning yields that $t'_{n'-1} = \tilde{t}_{\tilde{n}-1}$, $t'_{n'-2} = \tilde{t}_{\tilde{n}-2} \dots$. Since (T, \prec_T) is a tree, we have that $t' \prec_T \tilde{t}$, $\tilde{t} \prec_T t'$ or $t' = \tilde{t}$. Without loss of generality, we assume the first, the other cases being similar. Since $t' \prec_T \tilde{t}$, we have $n' > \tilde{n}$ and hence $\tilde{t} = t'_{n'-\tilde{n}}$. By the above, this implies that $(\tilde{t}, \tilde{s}) = (t'_{n'-\tilde{n}}, s'_{n'-\tilde{n}})$, and hence $(t', s') \prec'_T (\tilde{t}, \tilde{s})$.

Finally, we have to show that the tree order \prec'_T is serial. To this end, let $(t, s) \in T \otimes S$. We have to show that there is some $(t', s') \in T \otimes S$ with $(t, s) \prec'_T (t', s')$. Note that by seriality of T , there is some t' with $t \prec_T t'$. By our discreteness assumption, we can pick t' such that $t \prec_{im} t'$. Next, we claim that there is some $s' \in S$ with $s R_{ims} s'$ and $\mathcal{M}, t' \models pre(s')$. Note that this claim implies that $(t, s) \prec'_T (t', s')$, finishing the proof. So let us show the claim. Assume not. That is, assume that $\mathcal{M}, t' \not\models pre(s')$ for all $s' \in S$ with $s R_{ims} s'$. Hence, $\mathcal{M}, t' \models \neg pre(s')$ for all s' with $s R_{ims} s'$. This shows that $\{\neg pre(t) \mid s R_{ims} t\}$ is consistent, contradicting the assumption that S is a practical update model. ■

PROOF OF LEMMA 8.5 Since $\mathcal{M} \otimes \mathcal{E}$ is an expansion of \mathcal{E} , the set $S = \{[\psi]_{\mathcal{M} \otimes \mathcal{E}}^{(t_0, s_0)} \mid [\psi]_{\mathcal{M}}^{t_0} \in N_G^{\mathcal{M}}\}$ in $\mathcal{M} \otimes \mathcal{E}$ is consistent. By assumption, $n_G^{\mathcal{M}}$ is \succ_O -upward closed in $\{\varphi \mid ([\varphi]_{\mathcal{M}}^{t_0}, \varphi) \in N_O^{\mathcal{M}}\}$ and hence also in $\{\varphi \mid ([\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{t_0}, \varphi) \in N_O^{\mathcal{M} \otimes \mathcal{E}}\}$. In particular, the \prec_{Lex} maximally consistent subset of $\{[\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{(t_0, s_0)} \mid ([\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{t_0}, \varphi) \in N_O^{\mathcal{M} \otimes \mathcal{E}}\}$ contains S . This shows that $n_G^{\mathcal{M}} \subseteq n_G^{\mathcal{M} \otimes \mathcal{E}}$. ■

PROOF OF LEMMA 8.7 Assume for a contradiction that $n_G^{\mathcal{M}} \subset n_G^{\mathcal{M} \otimes \mathcal{E}}$. This implies that $n_G^{\mathcal{M}} \prec_{Lex}^{\mathcal{M}} n_G^{\mathcal{M} \otimes \mathcal{E}}$. Hence, by construction of the neighborhood $N_G^{\mathcal{M}}$, the set $\{[\varphi]_{\mathcal{M}}^{t_0} \mid [\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{t_0} \in N_G^{\mathcal{M} \otimes \mathcal{E}}\}$ must be inconsistent, i.e., there is no history h of \mathcal{T} such that $h \in [\varphi]_{\mathcal{M}}^{t_0}$ for all φ with $[\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{t_0} \in N_G^{\mathcal{M} \otimes \mathcal{E}}$. This, however, is impossible: $N_G^{\mathcal{M} \otimes \mathcal{E}}$ is by definition consistent in $\mathcal{M} \otimes \mathcal{E}$, i.e., there is some history h of $\mathcal{M} \otimes \mathcal{E}$ with $h \subseteq [\varphi]^{t_0}$ whenever $[\varphi]^{t_0} \in N_G^{\mathcal{M} \otimes \mathcal{E}}$. Since $\mathcal{M} \otimes \mathcal{E}$ is a restriction of \mathcal{M} , h is also a history of \mathcal{M} , showing that $\{[\varphi]_{\mathcal{M}}^{t_0} \mid [\varphi]_{\mathcal{M} \otimes \mathcal{E}}^{t_0} \in N_G^{\mathcal{M} \otimes \mathcal{E}}\}$ is consistent. ■

References

- [1] ANGLBERGER, A., N. GRATZL, and O. ROY, Obligation, free choice, and the logic of weakest permissions, *The Review of Symbolic Logic* 8(4):807–827, 2015.
- [2] BALTAG, A., L. S. MOSS, and S. SOLECKI, The Logic of Public Announcements, Common Knowledge, and Private Suspicions (extended abstract), in *Proc. of the Intl. Conf. TARK 1998*, Morgan Kaufmann Publishers, 1998, pp. 43–56.
- [3] BLACKBURN, P., M. DE RIJKE, and Y. VENEMA, *Modal Logic*. Cambridge University Press, 2001.
- [4] BRATMAN, M., *Intention, plans, and practical reason*, Harvard University Press, 1987.
- [5] BROERSEN, J., M. DASTANI, and L. VAN DER TORRE, Resolving conflicts between beliefs, obligations, intentions, and desires, in *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer, 2001, pp. 568–579.
- [6] BROOME, J., *Rationality Through Reasoning*, Wiley-Blackwell, 2013.
- [7] BROWN, M. A., Rich deontic logic: a preliminary study, *Journal of Applied Logic* 2(1):19–37, 2004.
- [8] CARIANI, F., Consequence and contrast in deontic semantics, *The Journal of Philosophy* 113(8):396–416, 2016.
- [9] CIUNI, R., and A. ZANARDO, Completeness of a branching-time logic with possible choices, *Studia Logica* 96(3):393–420, 2010.
- [10] COHEN, P. R., and H. J. LEVESQUE, Intention is choice with commitment, *Artificial Intelligence* 42(2):213–261, 1990.
- [11] CRAVEN, R., and M. SERGOT, Agent strands in the action language nc+, *Journal of Applied Logic* 6(2):172–191, 2008.
- [12] VAN DITMARSCH, H., W. VAN DER HOEK, and B. KOOI, *Dynamic Epistemic Logic*, Springer, 2008.
- [13] GIBBARD, A., *Reconciling Our Aims: In Search of Bases for Ethics*, Oxford University Press, 2008.
- [14] GORANKO, V. and S. VAN DRIMMELEN, Complete axiomatization and decidability of alternating-time temporal logic, *Theoretical Computer Science* 353(1–3):93–117, 2006.

- [15] HILPINEN, R., and P. MCNAMARA, Deontic logic: A historical survey and introduction, in *Handbook of deontic logic and normative systems*, Volume 1, College Publications, 2013, pp. 3–136.
- [16] VAN DER HOEK, W., W. JAMROGA, and M. WOOLDRIDGE, Towards a theory of intention revision, *Synthese* 155(2):265–290, 2007.
- [17] HORTY, J., *Agency and Deontic Logic*, Oxford University Press, 2001.
- [18] HORTY, J., Requirements, oughts, intentions, *Philosophy and Phenomenological Research* 91(1):220–229, 2015.
- [19] ICARD, T., E. PACUIT, and Y. SHOHAM, Joint revision of belief and intention, in *Proceedings of the 12th International Conference on Knowledge Representation*, 2010, pp. 572–574.
- [20] KOLODNY, N., Why be rational? *Mind* 114(455):509–563, 2005.
- [21] KRATZER, A., *The notional category of modality*, de Gruyter, 1981, pp. 38–74.
- [22] LORINI, E., and A. HERZIG, A logic of intention and attempt, *Synthese* 163(1):45–77, 2008.
- [23] MARRA, A., and D. KLEIN, Logic and ethics: An integrated model for norms, intentions and actions, in W. van der Hoek, W. H. Holliday, and W. Wang (Eds.), *Lecture Notes in Computer Science (LNCS) 9394*, Springer, 2015, pp. 268–281.
- [24] NAIR, S., Consequences of reasoning with conflicting obligations, *Mind* 123(491):753–790, 2014.
- [25] PACUIT, E., *Neighborhood semantics for modal logic*, *Short Textbooks in Logic*, Springer, 2017.
- [26] ROSS, A., Imperatives and logic, *Theoria* 7:53–71, 1941.
- [27] SHPALL, S., Wide and narrow scope, *Philosophical Studies* 163(3):717–736, 2013.
- [28] THOMASON, R. H., Desires and defaults: A framework for planning with inferred goals, in *KR*, Citeseer, 2000, pp. 702–713.
- [29] VELTMAN, F., Or else, what? imperatives on the borderline of semantics and pragmatics. Unpublished notes, 2011.
- [30] VENEMA, Y., Temporal logic, in L. Goble (Ed.), *The Blackwell guide to philosophical logic*, Blackwell, 2001, pp. 203–223.
- [31] VON WRIGHT, G. H., Practical inference, *The philosophical review* 72(2):159–179, 1963.
- [32] YALCIN, S., Modalities of normality, in N. Charlow and M. Chrisman (Eds.), *Deontic Modality*, Oxford University Press, 2016, pp. 230–255.

D. KLEIN, A. MARRA
Philosophy Department
University of Bayreuth
Universitätsstraße 30
95447 Bayreuth
Germany
dominik.klein@uni-bayreuth.de

A. MARRA
alessandra.marra@uni-bayreuth.de

D. KLEIN
Political Science Department
University of Bamberg
Feldkirchenstraße 21
96052 Bamberg
Germany