Expressivity of Logics of Knowledge and Action

Kuijer, Bouke

Chapter 1

Introduction

1.1 An Introduction to Expressivity

This thesis is about logic and in particular about the *expressivity*[1] of different logics. Informally speaking, expressivity is a measure of what can be said (or *expressed*) in a logic. A logic $\mathcal{L}_2$ is at least as expressive as a logic $\mathcal{L}_1$, denoted $\mathcal{L}_1 \preceq \mathcal{L}_2$, if everything that can be said in $\mathcal{L}_1$ can also be said in $\mathcal{L}_2$, so if for every formula φ_1 of $\mathcal{L}_1$ there is a formula φ_2 of $\mathcal{L}_2$ with the same meaning. In other words, we can define expressivity in the following way.

Definition 1.1. A logic $\mathcal{L}_2$ is *at least as expressive* as a logic $\mathcal{L}_1$ if for every formula φ_1 of $\mathcal{L}_1$ there is an equivalent formula φ_2 of $\mathcal{L}_2$.

Definitions equivalent to the above one[2] can be found in many different publications, including for example [Emerson and Halpern, 1986, Rabinovich, 1992, Hirshfeld and Rabinovich, 2006, van Ditmarsch et al., 2007, Kooi, 2007, Renne, 2008, Ågotnes et al., 2010, Hunter, 2013, Wáng, 2013] among many others.

There are four things that are important to note about this definition. Firstly, the word "equivalent" in Definition 1.1 should be read as "having the same truth value in every (pointed) model." So we require both $\mathcal{L}_1$ and $\mathcal{L}_2$ to have models and semantics, and furthermore we require that $\mathcal{L}_1$ and $\mathcal{L}_2$ have the same class of models. The second thing is that $\mathcal{L}_1 \preceq \mathcal{L}_2$ is equivalent to there being a *translation* t from $\mathcal{L}_1$ to $\mathcal{L}_2$, namely the function t that maps the $\mathcal{L}_1$ formula φ_1 onto the equivalent $\mathcal{L}_2$ formula φ_2.[3] The third thing to note is that $\preceq$ is reflexive and transitive (and therefore a preorder) but not total. For many logics $\mathcal{L}_1$ and $\mathcal{L}_2$ there are $\mathcal{L}_1$ formulas for which there is no equivalent $\mathcal{L}_2$ formula as well as $\mathcal{L}_2$ formulas for which there is no equivalent $\mathcal{L}_1$ formula. In those cases we have $\mathcal{L}_1 \not\preceq \mathcal{L}_2$ and $\mathcal{L}_2 \not\preceq \mathcal{L}_1$. Finally we should note that we only consider *relative* expressivity; we have defined what it means for a logic $\mathcal{L}_2$ to be at least as expressive as a logic $\mathcal{L}_1$. We have not defined the *absolute* expressivity $e(\mathcal{L})$ of a single logic $\mathcal{L}$. The reason for only defining $\preceq$ but not $e(\mathcal{L})$ is quite simple; we know how to define $\preceq$ but not how to define $e(\mathcal{L})$.

[1] Or expressive power, or expressiveness. These three terms are used interchangeably.

[2] Or at least: equivalent to a restriction of Definition 1.1 to some domain. Sometimes expressivity is defined in such a way that only logics of a certain kind can be compared.

[3] This translation t need not be computable.

Because of our lack of an absolute definition we must make do with the relative definition. Fortunately, it turns out that relative expressivity is quite interesting. In order to support this claim that relative expressivity is interesting let us spend a few paragraphs considering what expressivity can do for us.

By being a measure of what can be expressed in a logic, expressivity is also to some extent a measure of the versatility or perhaps even usefulness of a logic. If $\mathcal{L}_2$ is at least as expressive than $\mathcal{L}_1$ then whenever we use $\mathcal{L}_1$ we could, in theory, use $\mathcal{L}_2$ instead. If we want to determine whether an $\mathcal{L}_1$ formula φ_1 is true in a particular model we can instead check whether the equivalent $\mathcal{L}_2$ formula φ_2 holds there. If we want to determine whether φ_1 is satisfiable we can instead check whether φ_2 is satisfiable. If some complicated (or simple) bit of reasoning can be represented in $\mathcal{L}_1$ then it can be represented in $\mathcal{L}_2$ as well.

We should stress that the ability of $\mathcal{L}_2$ to replace $\mathcal{L}_1$ might be only theoretical. In order for $\mathcal{L}_2$ to be at least as expressive as $\mathcal{L}_1$ it is sufficient that for every $\mathcal{L}_1$ formula φ_1 an equivalent $\mathcal{L}_2$ formula φ_2 *exists*, we do not have to be able to compute φ_2.[4] And even if we can compute φ_2 it is possible that φ_2 is much longer than φ_1 or harder to read than φ_1, and the satisfiability or model checking problems of $\mathcal{L}_2$ may have a much higher computational complexity than that of $\mathcal{L}_1$. Still, if we only care about what a logic can represent and not about how efficiently the logic does so, expressivity is a very useful tool that helps us compare logics. As an example let us consider the expressivity of three simple logics.

Example 1.1. Suppose we want to use logic to formalize the sentences "Alice is either hungry or tired" and "If Alice is tired then, necessarily, she is not hungry." There are many logics we could choose to formalize these sentences but suppose that for some reason we have narrowed down our choice to the following three logics.

- $\mathcal{L}_\text{p}$, a classical propositional logic with the connectives $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$.
- $\mathcal{L}_{\{\neg,\wedge\}}$, a classical propositional logic with the connectives $\neg$ and $\wedge$.
- $\mathcal{L}_\text{m}$, a basic modal logic with the connectives $\neg, \vee, \wedge, \rightarrow, \leftrightarrow$ and $\Box$.

For technical reasons we assume that all three logics are evaluated on Kripke models, even though the relational structure of such models is only relevant for $\mathcal{L}_\text{m}$. We can now start to formalize the first sentence. This is most straightforwardly done in $\mathcal{L}_\text{p}$, where we can formalize it as $h \vee t$. But we could formalize it in either of the other logics as well. In $\mathcal{L}_{\{\neg,\wedge\}}$ we could, slightly awkwardly, formalize it as $\neg(\neg h \wedge \neg t)$ and in $\mathcal{L}_\text{m}$ we could formalize it as $h \vee t$ just like in $\mathcal{L}_\text{p}$. It is of course no coincidence that $\mathcal{L}_{\{\neg,\wedge\}}$ and $\mathcal{L}_\text{m}$ can formalize the sentence as well. After all, for every formula in $\mathcal{L}_\text{p}$ there are equivalent $\mathcal{L}_{\{\neg,\wedge\}}$ and $\mathcal{L}_\text{m}$ formulas. If an $\mathcal{L}_\text{p}$ formula φ formalizes a given sentence S then the equivalent $\mathcal{L}_{\{\neg,\wedge\}}$ and $\mathcal{L}_\text{m}$ formulas also formalize S.

But now let us look at the second sentence, "If Alice is tired then, necessarily, she is not hungry." This sentence can be formalized in $\mathcal{L}_\text{m}$ as $t \rightarrow \Box\neg h$.[5] But there is no way to formalize this sentence in either of the two other logics.

Everything that can be expressed in $\mathcal{L}_\text{p}$ can also be expressed in $\mathcal{L}_{\{\neg,\wedge\}}$ and $\mathcal{L}_\text{m}$. So $\mathcal{L}_{\{\neg,\wedge\}}$ and $\mathcal{L}_\text{m}$ are at least as expressive as $\mathcal{L}_\text{p}$. It is easy to see that, likewise,

[4] Although usually we can in fact compute φ_2.

[5] One could also argue that the correct formalization is $\Box(t \rightarrow \neg h)$ instead of $t \rightarrow \Box\neg h$. Which formalization we choose is unimportant for the matter at hand; with either choice the sentence can be formalized in $\mathcal{L}_\text{m}$ but not in the two other logics.

$\mathcal{L}_\text{p}$ is at least as expressive as $\mathcal{L}_{\{\neg,\wedge\}}$. This means that $\mathcal{L}_\text{p}$ and $\mathcal{L}_{\{\neg,\wedge\}}$ are equally expressive. On the other hand, there are things that can be expressed in $\mathcal{L}_\text{m}$ that cannot be expressed in $\mathcal{L}_\text{p}$. So $\mathcal{L}_\text{m}$ is (strictly) more expressive than $\mathcal{L}_\text{p}$. In the end, if we wish to formalize both sentences we have no choice but to use $\mathcal{L}_\text{m}$; neither of the other logics is expressive enough.

One important thing to note about this example is that we used expressivity to determine which logic is most appropriate for our needs: if we want to formalize "If Alice is tired then, necessarily, she is not hungry" then a basic propositional logic will not suffice but a modal logic will. So expressivity can be a reason to prefer one logic over another.

Expressivity is of course not the only reason why we could prefer one logic over another. We could prefer $\mathcal{L}_2$ over $\mathcal{L}_1$ because it has more intuitive or natural semantics, because its decision problems have a better computational complexity or because it is more succinct. But expressivity does to some extent have a privileged position among these reasons. Suppose the semantics of $\mathcal{L}_1$ are rather unintuitive, perhaps because the only connective it has is the Sheffer stroke. Then we can still use $\mathcal{L}_1$, although it will take some effort. Suppose then that the decision problems of $\mathcal{L}_1$ have a high computational complexity. That might be a serious problem if we want to determine whether or not some large formula is valid. But generally we will still be able to reason about simple formulas. If $\mathcal{L}_1$ is not expressive enough we cannot even do that. If we want to reason about, say, necessity, knowledge, time or action, a basic propositional logic will be completely useless; even the simplest sentences regarding these concepts cannot be formalized in a logic that is not expressive enough.

Still, despite the somewhat privileged position of expressivity we should not forget that there are other criteria for preferring one logic over another. These other criteria are worthy of study, just like expressivity. Studying them is outside the scope of this thesis though, as one might expect given its title.

Now that we are on the subject of the title of this thesis it may be worthwhile to explain why we are considering the expressivity of logics of knowledge and action instead of other logics. The main reason for looking at logics of knowledge and action is that there are many of them that use the same class of models. We can only apply Definition 1.1 to logics that share a class of models, so a large set of logics with the same class of models is a very good place to look for interesting expressivity results.

Of course logics of knowledge and action are not the only area where a lot of logics use the same models. Another such area is temporal logic, and it is no coincidence that the expressivity of temporal logics has also been studied extensively (most famously in [Kamp, 1968]).

Remark. Above I mentioned that we do not have a good definition for the absolute expressivity $e(\mathcal{L})$ of a logic $\mathcal{L}$. But even if we would manage to find such a definition that would not help us much, since we would have no hope of exactly determining the value of $e(\mathcal{L})$ for a specific logic $\mathcal{L}$.

The expressivity $e(\mathcal{L})$ of a logic $\mathcal{L}$ should be the set of those things that can be expressed in a logic. In order to properly define e we would therefore first have to determine the category C of things that could possibly be expressed, and what it means for a logic to express some $c \in C$. Finding this C would be hard, perhaps even impossible. Worse, even if we were to find this C we would still be unable to determine the exact expressivity of a logic.

In the example we used logics to express the two sentences "Alice is either hungry or tired" and "If Alice is tired then, necessarily, she is not hungry." There is nothing special about these sentences, we could have taken any other declarative sentences and asked whether the logics can express those sentences. This means that C must include, possibly in addition to other things, all declarative sentences in natural language. But natural language is too imprecise for us to have any hope of determining exactly which sentences can be expressed in a logic and which cannot.

1.2 Outline of this Thesis

This thesis consists of two parts. The first part contains expressivity results, the second part results that are not themselves expressivity results[6] but that are related to expressivity.

Part I consists of Chapters 2 and 3. In Chapter 2 we compare the logic $\mathcal{L}_{\text{CP}}$ (an epistemic logic with *common knowledge* and *public announcements*) to a logic $\mathcal{L}_{\text{CPS}}$ (an epistemic logic with *common knowledge*, *public announcements* and *public substitutions*). Of these two logics it was already known from [Kooi, 2007] that $\mathcal{L}_{\text{CPS}}$ is more expressive than $\mathcal{L}_{\text{CP}}$ over the class **K** of models. We show that $\mathcal{L}_{\text{CPS}}$ is also more expressive than $\mathcal{L}_{\text{CP}}$ over the classes **KD45**, **S4** and **S5** of models, if the set of agents is large enough.

In Chapter 3 we consider a large number of logics, most notably $\mathcal{L}_{\text{R}}$ (epistemic logic with *relativized common knowledge*), $\mathcal{L}_{\text{CU}}$ (epistemic logic with *common knowledge* and *arrow updates*) and $\mathcal{L}_{\text{U}^*}$ (epistemic logic with *arrow common knowledge*). I show that $\mathcal{L}_{\text{U}^*}$ is more expressive than $\mathcal{L}_{\text{R}}$. Additionally, I show that $\mathcal{L}_{\text{CU}}$ is equally expressive as $\mathcal{L}_{\text{U}^*}$. The proof of this second claim is extremely long, technical and complicated. Many of the details are therefore given not in Chapter 3 but in the appendix.

Part II consists of Chapters 4, 5 and 6. In it we consider a number of results that are not themselves expressivity results but that are closely related to questions of expressivity. In Chapter 4 we look at Arbitrary Public Announcement Logic (APAL) and ask whether the *arbitrary public announcement* operator $\Diamond$ in that logic is truly arbitrary. In the end we conclude that it is not, in fact, truly arbitrary. The formulas and models used in the proof that $\Diamond$ is not fully arbitrary are very similar to those used in expressivity proofs. This is no coincidence; the fact that $\Diamond$ is not fully arbitrary depends strongly on the fact that APAL is more expressive than Public Announcement Logic (PAL), which was shown to be the case in [Balbiani et al., 2007].

In Chapter 5 we look at a number of deontic logics, logics intended to reason about obligations and permissions. But we do not compare the expressivity of these logics. Instead we investigate whether certain sentences can be expressed in particular logics. As mentioned in the remark in the previous section we cannot determine the exact set of sentences that can be expressed in a logic. We can however choose a few specific sentences and check whether those can be expressed in a particular logic. Specifically, in Chapter 5 we investigate whether any deontic logic that is based on the idea of a single "sanction" representing wrongdoing is capable of expressing a certain type of obligation called a *contrary-to-duty* obligation. We conclude that none of the logics under consideration have the tools required to represent such obligations.

Finally, in Chapter 6 we try to generalize expressivity. Recall that a logic $\mathcal{L}_2$ is at least as expressive as a logic $\mathcal{L}_1$ if for every formula φ_1 of $\mathcal{L}_1$ there is an equivalent formula φ_2 of $\mathcal{L}_2$, and that the word "equivalent" should be read as "having the same

[6] So these results do not show that $\mathcal{L}_1 \preceq \mathcal{L}_2$ or that $\mathcal{L}_1 \not\preceq \mathcal{L}_2$ for any logics $\mathcal{L}_1$ and $\mathcal{L}_2$.

truth value on every pointed model". So the formula φ_2 must be such that for every pointed model $\mathcal{M}, w$ we have $\mathcal{M}, w \models \varphi_1$ if and only if $\mathcal{M}, w \models \varphi_2$. But that means we cannot compare the expressivity of logics that have different classes of models. There are very good reasons for restricting expressivity in this way; if we carelessly generalize expressivity to allow different models we risk ending up with a trivial concept where almost all logics are as expressive as almost all other logics.

Still, a number of interesting results have been published that have strong similarities to expressivity results, even though they compare logics that have different models. This suggests that there is some generalization of expressivity that is nontrivial and interesting. In Chapter 6 we try to find a generalization of expressivity that covers these interesting results. Because we are generalizing a formal concept, instead of proving things about an existing formal concept, this chapter is more speculative than the other chapters.

1.3 Previous Work

Several of the chapters in this thesis are based on previously published papers. Specifically, Chapter 2 is based on the paper "The Expressivity of Factual Change in Dynamic Epistemic Logic" [Kuijer, 2014b], which appeared in the Review of Symbolic Logic. Chapter 3 is based on the paper "The Expressivity of Update Logic" [Kuijer, to appear], which is to appear in the Journal of Logic and Computation.[7] Chapter 4 is based on the paper "How Arbitrary are Arbitrary Public Announcements?"[Kuijer, 2013] which was presented at the student session of ESSLLI 2013. An extended version [Kuijer, 2014a] of that paper also appeared in a collection of selected papers from the ESSLLI 2012 and ESSLLI 2103 student sessions. Finally, Chapter 5 is based on the paper "Sanction Semantics and Contrary-to-Duty Obligations" [Kuijer, 2012], which was presented at the 11th International Conference on Deontic Logic in Computer Science.

The structure of the papers that these chapters are based on has, mostly, been preserved. This should allow most of the chapters of this thesis to be read separately. In particular, every chapter is self-contained in the sense that it contains all definitions that are needed in that chapter.

[7]The electronic version of [Kuijer, to appear] is already available at http://logcom.oxfordjournals.org/content/early/2014/09/04/logcom.exu047.

Part I

Expressivity Results