Chapter 4

The Formal System I: Basic Rules

4.1 Rules Inherited from First Order Logic

With the language of logical possibility $\mathcal{L}$ in place, I will now introduce some inference rules for reasoning about logical possibility. I will recursively define the set of strings which constitute proofs in my deductive system by listing closure conditions in this chapter and the next.[1] Let $\Gamma, \Gamma_1, \Gamma_2$ be finite sets of formulas, and $\Gamma \vdash \theta$ express the claim that one can prove θ given the assumptions in Γ.

My closure conditions begin, boringly, with the following principles

[1]As usual, I will say that all variables that occur in an atomic formula are free. If a variable occurs free (or bound) in θ or in ψ, then that same occurrence is free (or bound) in $\neg\theta$, $(\theta \wedge \psi)$, $(\theta \vee \psi)$, and $(\theta \rightarrow \psi)$ and $\Diamond\theta$ and $\Box\theta$. That is, the (unary and binary) connectives do not change the status of variables that occur in them. All occurrences of the variable v in θ are bound in $\forall v\theta$ and $\exists v\theta$. Any free occurrences of v in θ are bound by the initial quantifier. All other variables that occur in θ are free or bound in $\forall v\theta$ and $\exists v\theta$, as they are in θ.

corresponding to standard inference rules for first order logic,(which I take from the Stanford Encyclopedia article on classical logic[2]).

(As) If ϕ is a member of Γ, then $\Gamma \vdash \phi$.

($\wedge$I) If $\Gamma_1 \vdash \theta$ and $\Gamma_2 \vdash \psi$, then $\Gamma_1, \Gamma_2 \vdash (\theta \wedge \psi)$.

($\wedge$E) If $\Gamma \vdash (\theta \wedge \psi)$ then $\Gamma \vdash \theta$; and if $\Gamma \vdash (\theta \wedge \psi)$ then $\Gamma \vdash \psi$.

($\vee$I) If $\Gamma \vdash \theta$ then $\Gamma_1 \vdash \theta \vee \psi$; if $\Gamma \vdash \psi$ then $\Gamma \vdash \theta \vee \psi$.

($\vee$E) If $\Gamma_1 \vdash (\theta \vee \psi)$, $\Gamma_2, \theta \vdash \phi$ and $\Gamma_3, \psi \vdash \phi$, then $\Gamma_1, \Gamma_2, \Gamma_3 \vdash \phi$.

($\rightarrow$I) If $\Gamma, \theta \vdash \psi$, then $\Gamma \vdash (\theta \rightarrow \psi)$.

($\rightarrow$E) If $\Gamma_1 \vdash (\theta \rightarrow \psi)$ and $\Gamma_2 \vdash \theta$, then $\Gamma_1, \Gamma_2 \vdash \psi$.

($\neg$I) If $\Gamma_1, \theta \vdash \psi$ and $\Gamma_2, \theta \vdash \neg\psi$, then $\Gamma_1, \Gamma_2 \vdash \neg\theta$.

(DNE) If $\Gamma \vdash \neg\neg\theta$ then $\Gamma \vdash \theta$.

($\forall$E) If $\Gamma \vdash \forall v \theta$, then $\Gamma \vdash \theta(v|v')$, provided that v' is free for v in θ.[3]

($\forall$I) If $\Gamma \vdash \theta$ and the variable v does not occur free in any member of Γ, then $\Gamma \vdash \forall v \theta$.

(=I) $\Gamma \vdash v = v$, where v is any variable.

(=E) If $\Gamma_1 \vdash v_1 = v_2$ and $\Gamma_2 \vdash \theta$, then $\Gamma_1, \Gamma_2 \vdash \theta'$, where θ' is obtained from θ by replacing zero or more occurrences of v_1 with v_2, provided that no bound variables are replaced, and all substituted occurrences of v_2 are free.

($\perp$ I) If $\Gamma \vdash \psi \wedge \neg\psi$ then $\Gamma \vdash \perp$.

($\perp$ E) If $\Gamma, \theta \vdash \perp$ then $\Gamma \vdash \neg\theta$.

For convenience, I will also include the following inference rules for $\exists$

[2] http://plato.stanford.edu/entries/logic-classical/ , with straightforward simplifications arising from the fact that my language $\mathcal{L}$ does not contain any constants

[3] That is, if substituting v with v' does not lead to any variable which was antecedently free becoming bound. Here $\theta(v|v')$ stands for the result of substituting *all* free instances of v in θ with instances of v'.

whose validity is straightforward to demonstrate using the definition of $\exists$ (as an abbreviation for $\neg\forall\neg$).

($\exists$I) If $\Gamma \vdash \theta$, then $\Gamma \vdash \exists v\theta'$, where θ' is obtained from θ by substituting the variable v' for zero or more occurrences of a variable v, provided that (1) all of the replaced occurrences of v are free in θ, and (2) all of the substituted occurrences of v' are free in θ.

($\exists$E) If $\Gamma_1 \vdash \exists v\theta$ and $\Gamma_2, \theta \vdash \phi$, then $\Gamma_1, \Gamma_2 \vdash \phi$, provided that v does not occur free in ϕ, nor in any member of Γ_2.

In order to state analogous inference rules for the $\Box$ and $\Diamond$, I will define a sense in which a sentence can be *content restricted* to a finite list of relations $\mathcal{L}$. Note that, just like the relations subscripted by a $\Diamond$ or $\Box$, the order of the relations in $\mathcal{L}$ does not matter, so we may freely take intersections or talk of one list containing another without concern for order.

4.2 Basic $\Diamond$ Rules

4.2.1 Content-restriction

In reasoning about logical possibility, it will be useful to distinguish a class of sentences whose truth depends only on the facts about a list $\mathcal{L}$ of relations, i.e., those sentences ϕ such that $\Diamond_{\mathcal{L}}\phi$ intuitively entails ϕ. We will call such a sentence *content restricted to* $\mathcal{L}$. For example, if $\mathcal{L}$ is the list $\text{Person}(\cdot), \text{Loves}(\cdot,\cdot)$ then the claim 'every person loves something', i.e., $(\forall x)\,[\text{Person}(x) \implies (\exists y)\text{Loves}(x,y)]$, is content restricted to $\mathcal{L}$. In contrast the sentence 'every thing loves some thing', i.e., $(\forall x)(\exists y)(\text{Loves}(x,y))$, is not content restricted as it's truth depends on the existence of objects that neither Person nor Loves[4] applies to. As these examples suggest, sentences are content restricted if only the relations from $\mathcal{L}$ are mentioned and every quantifier is restricted to range over elements that belong to some tuple in the extension of a relation in $\mathcal{L}$. The following definitions capture this intuition.

Definition Let $y \in \mathrm{Ext}(R_1, \ldots R_n)$ abbreviate the formula

$$\bigvee_{\substack{1 \le i \le n \\ 1 \le j \le l_i}} (\exists x_1) \ldots, (\exists x_{j-1}), (\exists x_{j+1}), \ldots, (\exists x_{l_i}) R_i(x_1, \ldots, x_{j-1}, y, x_{j+1}, \ldots, x_{l_i})$$

where l_i is the arity of R_i and $\bigvee_{\substack{1 \le i \le n \\ 1 \le j \le l_i}} \phi_{i,j}$ indicates the disjunction $\phi_{i,j}$ over all indicated values for i and j.

Thus, $y \in \mathrm{Ext}(R_1, \ldots R_n)$ is the formula asserting that some tuple $\vec{v}$ including y satisfies some $R_i(\vec{v})$.

Definition I will say that a sentence ϕ is **explicitly content-restricted** to a list $\mathcal{L}$ if it is a member of the smallest set S satisfying:

- $\perp$ is in S
- If v_i, v_j are variables the formula $v_i = v_j$ is in S
- If v_i is a variable and $R_i \in \mathcal{L}$ then $R_i(v_j)$ is in S
- If $\psi \in S$ and $\rho \in S$ then $\neg\psi$, $\psi \vee \rho$, $\psi \wedge \rho$ and $\psi \to \rho$ are all in S
- If $\psi \in S$ and $\mathcal{L}$ is non-empty then $\exists y(y \in \mathrm{Ext}(\mathcal{L}) \wedge \psi)$ is in S

[4] By this, I mean objects which are not part of any pair in the extension of Loves.

- If $\psi \in S$ and $\mathcal{L}$ is non-empty then $\forall y(y \in \text{Ext}(\mathcal{L}) \rightarrow \psi)$ is in S

- If $\phi = \Diamond_{\mathcal{L}'}\psi$, where ψ is a sentence and $\mathcal{L}' \subseteq \mathcal{L}$ then $\phi \in S$. Note that *ψ need not be in S*

The last clause is motivated by the fact that the truthvalue of $\Diamond_{\mathcal{L}'}\psi$ is completely determined by facts about the relations in $\mathcal{L}'$. Furthermore, as no free variables are allowed in $\Diamond_{\mathcal{L}'}\psi$ its truth value is unaffected by any external quantification.

Thus, for example, if $\mathcal{L}$ is a list that contains (exactly) a two-place relation R and a one place relation Q, then $(\forall x)(\forall y)(x = y)$ is not content-restricted to $\mathcal{L}$. Neither is $(\exists x)(Q(x) \wedge K(x))$. But $(\forall x)[x \in \text{Ext}(R) \rightarrow (\forall y)(y \in \text{Ext}(R) \rightarrow [R(x,y) \rightarrow Q(y)]])$[5] (which is first order logically equivalent to $(\forall x)(\forall y)[R(x,y) \rightarrow Q(y)]$) is content-restricted to $\mathcal{L}$. And so is $\Diamond_R[(\forall x)(R(x,x) \wedge (\exists y)S(x,y))]$.

Also note the following consequences of the definition above:

- If $\mathcal{L}$ is a sublist of $\mathcal{L}'$, then all formulae ϕ which are content restricted to $\mathcal{L}$ are also content restricted to $\mathcal{L}'$.

- A sentence is content restricted to the empty list $\mathcal{E}$ iff it is a truth functional combination of unsubscripted $\Box$ or $\Diamond$ sentences or $\bot$.

As you may have noticed, explicitly content-restricted sentences are generally long and unwieldy. This can be annoying when writing up proofs whose inference steps can only (strictly speaking) be applied to sentences

[5]i.e. $(\forall x)[(\exists k)(R(x,k) \vee R(k,x)) \rightarrow (\forall y)[(\exists k')(R(y,k') \vee R(k',y)) \rightarrow (R(x,y) \rightarrow Q(x))]]$

which are content-restricted to some list $\mathcal{L}$. To avoid this annoyance, I make the following definition.

Definition A sentence ϕ is **implicitly content-restricted** to $\mathcal{L}$ if there is a sentence ψ explicitly content restricted to $\mathcal{L}$ and $\phi \leftrightarrow \psi$ can be derived (using no assumptions) using the above inference rules.

I will then frequently use the short hand of applying rules which (strictly speaking) can only be applied to content-restricted sentences to implicitly content restricted sentences – taking the work of using first order logic to deduce the explicitly content-restricted form of a sentence before applying the relevant rule (and then transforming it back after applying the rule) for granted.

4.2.2 Rules

I can now introduce the core inference rules and axiom schemas which govern reasoning with $\Box$ and $\Diamond$ in my formal system.

($\Diamond$ I) Diamond Introduction. If $\Gamma \vdash \theta$ and θ is a sentence, then $\Gamma \vdash \Diamond_{\mathcal{L}}\theta$

This rule captures the idea that what is actual must also be logically possible, even while holding fixed the facts any list of relations $\mathcal{L}$ one might care to specify.

Examples:

- "There are two cats" $\Rightarrow$ "It is logically possible, given what cats there are, that there are two cats".

- "There are two cats" ⇒ "It is logically possible, given what dogs there are, that there are two cats".

(◇ E) Diamond Elimination. If $\Gamma \vdash \Diamond_L \theta$ and θ is content-restricted to $\mathcal{L}$ then $\Gamma \vdash \theta$

This rule expresses the idea that when θ is content-restricted to $\mathcal{L}$, the truth value of θ is totally determined by the facts about $\mathcal{L}$.

For instance:

- "It is logically possible, given what cats there are, that there are two cats" ⇒ "There are two cats"
- BUT NOT: "It is logically possible, given what dogs there are, that there are two cats" ⇒ "There are two cats"

Note that the second inference is not permitted by my rule because θ ("there are two cats") is not content-restricted to the list $\{dog(\cdot)\}$

(In◇) Inner Diamond.

Suppose $\Gamma_1 \vdash \Diamond_L \theta$. If $\Gamma_2, \theta \vdash \phi$, where $\Gamma_2 = \gamma_1 ... \gamma_m$ is a list of sentences which are content-restricted to $\mathcal{L}$. Then $\Gamma_1, \Gamma_2 \vdash \Diamond_{\mathcal{L}'} \phi$ for any $\mathcal{L}' \subseteq \mathcal{L}$.

This inference rule captures reasoning of the following form. Given the facts about $\mathcal{L}$, it's possible that θ. Any scenario where θ is true while the facts about $\mathcal{L}$ are held fixed must also be one in which the premises $\gamma_1 \wedge ... \wedge \gamma_m$ are true (because these sentences are content-restricted to $\mathcal{L}$). Thus it must be possible, given the facts about $\mathcal{L}$, that $\theta \wedge \gamma_1 \wedge ... \wedge \gamma_m$. As a matter of logic, any scenario in which $\theta \wedge \gamma_1 \wedge ... \wedge \gamma_m$ is one in which ϕ. So, it must be possible given the facts about $\mathcal{L}$ that ϕ. And since $\mathcal{L}' \subseteq \mathcal{L}$, ϕ must also but possible holding fixed only the facts about $\mathcal{L}'$.

I will use some visually suggestive notation to keep track of inferences of this form, as follows:

1	γ	Assump [1]
2	$\Diamond_{\mathcal{L}}\theta$	Assump [2]
3	$\Diamond$ | θ $[\mathcal{L}]$	2, In$\Diamond$I [2]
4	| γ	1, import [1]
5	| ...	
6	| ϕ	[1,2]
7	$\Diamond_{L}\phi$	2,3-6 In$\Diamond$E [1,2]

Intuitively speaking, the forked line going from 3-6 above separates off a location for reasoning about a logically possible scenario in which θ is true while all the facts about $\mathcal{L}$ in our current context are preserved.

A line ρ can be written down inside the "$\Diamond_{\mathcal{L}}$ context" governed by the claim that $\Diamond_{\mathcal{L}}\theta$ if

- $\rho = \theta$
- $\rho = \gamma$ for some γ which is content-restricted to $\mathcal{L}$ and occurs on an earlier line in the proof which is in the same context as the $\Diamond_{\mathcal{L}}\theta$ statement used to introduce this inner diamond context.

- ρ follows from previous lines within this ◇ context by one of the inference rules for reasoning about logical possibility presented in this paper.

One can leave $\Diamond_{\mathcal{L}}$ context above by going from knowledge that ϕ holds within this context to the conclusion that $\Diamond_{\mathcal{L}'}\phi$ holds outside it, provided that $\mathcal{L}'$ is a sublist of $\mathcal{L}$.

Example of Deploying In◇:

Consider the following very short argument.

> Given what cats and hunters there are, its logically possible that something is both a cat and a hunter.⇒ Given what cats there are, its logically possible that something is both a cat and a hunter.

We can capture this argument in my system as follows.

1	$\Diamond_{cat,hunter}(\exists x)(cat(x) \wedge hunter(y))$		[1]
2	◇ | $(\exists x)(cat(x) \wedge hunter(y))$	$[cat, hunter]$	1 In◇I, [1,2]
3	| $(\exists x)(cat(x) \wedge hunter(y))$		2 repetition[6] [1,2]
4	$\Diamond_{cat}(\exists x)(cat(x) \wedge hunter(y))$		2-3, In◇E [1]

Thus $\Diamond_{cat,hunter}(\exists x)(cat(x) \wedge hunter(y)) \vdash \Diamond_{cat}(\exists x)(cat(x) \wedge hunter(y))$.

[6] strictly speaking this repetition is not necessary

Note that $\{cat\}$ is a sublist of $\{cat, hunter\}$ and no extra premises are used in the deduction of ϕ from ϕ, so the requirements for Inn$\Diamond$E are satisfied.

Note: the requirement that each γ_i be content-restricted to $\mathcal{L}$ prevents us from importing facts about objects which don't satisfy any of the relations in $\mathcal{L}$ into our reasoning about what the relevant scenario $\mathcal{L}$ must be like. For example, consider the following *invalid* inference.

> "There is something that is not a cat." "It is logically possible, given what cats there are, that everything is a cat." ⇒ "It is logically possible given what cats there are, that everything is a cat and something is not a cat."

1		$(\exists x)\neg cat(x)$	[1]
2		$\Diamond_{cat}(\forall y)cat(y)$	[2]
3	$\Diamond$	$(\forall y)cat(y)$ $[cat]$	2, In$\Diamond$I
4		$(\exists x)\neg cat(x)$	2 import [1] (**INVALID**)
5		$(\exists x)\neg cat(x) \wedge (\forall y)cat(y)$	3,4 $\wedge$ I [1,2,3]
6		$\Diamond[(\exists x)\neg cat(x) \wedge (\forall y)cat(y)]$	1, 3-5 In$\Diamond$ E [1,2]

We cannot import [4] from line 1 because only sentences content restricted to [cat()] can be imported.

(◇ Ign) ◇ Ignoring. Suppose θ is content-restricted to $\mathcal{L}, R_1, \ldots R_n$ and $S_1 \ldots S_m$ are relations not among $\mathcal{L}, R_1, \ldots R_n$. If $\mathit{\Gamma}_1 \vdash \Diamond_{\mathcal{L}}\theta$ then $\mathit{\Gamma}_1 \vdash \Diamond_{\mathcal{L},S_1\ldots S_m}\theta$. Conversely, if $\mathit{\Gamma}_1 \vdash \Diamond_{\mathcal{L},S_1\ldots S_m}\theta$ then $\mathit{\Gamma}_1 \vdash \Diamond_{\mathcal{L}}\theta$.

Remember that when a formula is content-restricted to $\mathcal{L}$, its truth depends only on facts about $\mathcal{L}$. This principle reflects this intuition by allowing one to ignore other facts.

Examples:

- It is possible, given what cats there are, that there every cat admires a distinct dog $\leftrightarrow$ It is possible, given what cats and dolphins there are, that every cat admires a different dog.

- But NOT: It is possible, given what cats there are, that there are exactly 3 objects $\leftrightarrow$ It is possible, given what cats and dolphins there are, that there are exactly 3 objects.

  This inference is not permitted because the claim that there are exactly 3 objects is not content restricted to any list of relations including $cats()$ but not $dolphin()$.

- And NOT: It is possible, given what cats there are, that every cat admires a distinct dog $\leftrightarrow$ It is possible, given what cats and dogs there are, that every cat admires a distinct dog.

  Here θ is content restricted to {cat, dog, admires}, but for this inference to be permitted θ would have to be content restricted to a list that didn't include the relation dog.

For each of the $\Diamond$ principles above, an analogous inference involving $\Box$ can be justified, exploiting the fact that $\Box_{\mathcal{L}}\phi$ abbreviates $\neg\Box_{\mathcal{L}}\neg\phi$. See Appendix 4.4 for details. Also like $\exists$ the choice to define $\Box_{\ldots}$ in terms of $\Diamond_{\ldots}$ rather than vice-versa was arbitrary and either choice yields the same results.

4.3 Example: Pasting Lemma

Let us now get a little experience with how these basic inference rules work together, by using them to prove the following helpful lemma.

Lemma 4.3.1. ***(P) Pasting*** *Let $\mathcal{I}$, $\mathcal{J}$ and $\mathcal{L}$ be pairwise disjoint sets of relations. If $\Diamond_{\mathcal{L}}\phi$, where ϕ is content restricted to $\mathcal{L},\mathcal{I}$ and $\Diamond_{\mathcal{L}}\psi$, where ψ is content-restricted to $\mathcal{L},\mathcal{J}$, then $\Diamond_{\mathcal{L}}(\phi\wedge\psi)$.*

One cannot generally infer from $\Diamond_L\phi$ and $\Diamond_L\psi$ to $\Diamond_L(\phi\wedge\psi)$; consider the case where ϕ says there are exactly 8 million things and ψ says there are exactly 9 million things. However, this principle says that one *can* make this inference when ϕ and ψ describe suitably disjoint aspects of the universe (outside of the objects satisfying relations in $\mathcal{L}$).

We can prove this lemma using the basic inference rules and axiom schemas above as follows:

Proof. Let ϕ be content restricted to $\mathcal{L},\mathcal{I}$ and ψ to $\mathcal{L},\mathcal{J}$, as per the antecedent.

1	$\Diamond_{\mathcal{L}}\phi$	[1]
2	$\Diamond_{\mathcal{L}}\psi$	[2]
3	$\Diamond$ | $\phi \quad [\mathcal{L}]$	1, In$\Diamond$I [1]
4	| $\Diamond_{\mathcal{L}}\psi$	2, import [2]
5	| $\Diamond_{\mathcal{L},\mathcal{I}}\psi$	4, Ign [2]
6	| $\Diamond$ | $\psi \quad [\mathcal{L},\mathcal{I}]$	5, In$\Diamond$I [2]
7	| | ϕ	3, import [1]
8	| | $\phi \wedge \psi$	5,6 &I [1,2]
9	| $\Diamond_{\mathcal{L}}(\phi \wedge \psi)$	5,6-8 In$\Diamond$E [1,2]
10	$\Diamond_{\mathcal{L}}(\Diamond_{\mathcal{L}}(\phi \wedge \psi))$	1,3-9 In$\Diamond$E [1,2]
11	$\Diamond_{\mathcal{L}}(\phi \wedge \psi)$	10, $\Diamond$E [1,2]

$\square$

Informally, this deduction corresponds to the following reasoning:

Assume that $\Diamond_{\mathcal{L}}\phi$ and $\Diamond_{\mathcal{L}}\phi$. We can prove our claim by making two nested In$\Diamond$ arguments.

First enter the $\Diamond_{\mathcal{L}}$ context associated with $\Diamond_{\mathcal{L}}\phi$. In this context we clearly have ϕ. But we also know that $\Diamond_{\mathcal{L}}\psi$ must remain true, because it was true in our previous context and it is content restricted to $\mathcal{L}$. We can deduce from this that $\Diamond_{\mathcal{L},\mathcal{I}}\psi$ by Ignoring.

Now enter this second, interior, $\Diamond_{\mathcal{L},\mathcal{I}}$ context. Here we clearly have ψ. But we can import the fact that ϕ from the previous context, because it is content restricted to $\mathcal{L},\mathcal{I}$. So we can deduce $\phi \wedge \psi$.

Now, leaving this inner $\Diamond_{\mathcal{L},\mathcal{I}}$ context, we can conclude that $\Diamond_{\mathcal{L}}(\phi \wedge \psi)$ (because $\mathcal{L}$ is clearly a sublist of $\mathcal{L},\mathcal{I}$).

So, leaving the larger $\Diamond_{\mathcal{L}}$ context we can conclude that $\Diamond_{\mathcal{L}}(\Diamond_{\mathcal{L}}(\phi \wedge \psi))$ holds in the situation we were originally considering.

Finally, because $\Diamond_{\mathcal{L}}(\phi \wedge \psi)$ is content restricted to $\mathcal{L}$, we can use $\Diamond$E to draw the desired conclusion $\Diamond_{\mathcal{L}}(\phi \wedge \psi)$.

4.4 $\Box$ Inf. Rules

Although the $\Box$ is not an official item in our symbolism, but merely an abbreviation for $\neg \Diamond \neg$, it is often helpful to reason in terms of it. Thus we should note that the above inference rules can be used to vindicate analogous inference rules involving the $\Box$:

($\Box$ I) Box Introduction. If $\Gamma \vdash \theta$, where $\Gamma = \gamma_1 ... \gamma_m$ and for all i γ_i is content-restricted to $\mathcal{L}$ then $\Gamma \vdash \Box_{\mathcal{L}}\theta$.

As with In$\Diamond$, I will use some visually suggestive notation to keep track of inferences of this form, as follows:

1	γ	Assump [1]
2	$\Box$ $[\mathcal{L}]$	
3	γ	1, import [1]
4	...	
5	ϕ	[1]
6	$\Box_{\mathcal{L}}\phi$	2-5 $\Box$I [1]

Intuitively speaking, the forked line going from 3-6 above demarcates reasoning about what an arbitrary logically possible scenario in which all the facts about $\mathcal{L}$ (in our current context) are held fixed would have to be like.

A line ρ can be written down inside this "$\Box_{\mathcal{L}}$ introduction context" if

- $\rho = \gamma$ for some γ which is content-restricted to $\mathcal{L}$ and occurs on an earlier line in the proof in the same context as the intended conclusion of this $\Box_{\mathcal{L}}$I argument.

- ρ follows from previous lines in this $\Box_{\mathcal{L}}$ introduction context by one of the inference rules for reasoning about logical possibility presented in this paper.

One can leave $\Box_{\mathcal{L}}$ context above by going from knowledge that ϕ holds within this context to the conclusion that $\Box_{\mathcal{L}'}\phi$ holds outside it, provided that $\mathcal{L}'$ is a sublist of $\mathcal{L}$.

Proof. Suppose we have $\gamma_1 ... \gamma_m \vdash \theta$ as above. Then we can derive $\Box_{\mathcal{L}}\theta$ from Γ as follows.

1		$\gamma_1 ... \gamma_m$	[1]
2		$\Diamond_{\mathcal{L}} \neg\theta$	[2]
3	$\Diamond$	$\mid \neg\theta$	In$\Diamond$ I [2]
4		$\mid \gamma_1 ... \gamma_m$	import [Γ]
5		$\mid$ ...	
6		$\mid \theta$	[2, Γ]
7		$\mid \bot$	3, 6 $\bot$ I
8		$\Diamond_{\mathcal{L}} \bot$	2,3-7 In$\Diamond$E [2,Γ]
9		$\bot$	8 $\Diamond$E [2,Γ]
10		$\neg \Diamond_{\mathcal{L}} \neg\theta$	2-9 $\neg$I [Γ]
11		$\Box_{\mathcal{L}}\theta$	[Γ]

□

($\Box$ E) Box Elimination. If $\Gamma \vdash \Box_{\mathcal{L}}\theta$ then $\Gamma \vdash \theta$

1	$\Box_{\mathcal{L}}\theta$	$[\Gamma]$
2	$\neg\Diamond_{\mathcal{L}}\neg\theta$	$[\Gamma]$
3	$\neg\theta$	Assump. [3]
4	$\Diamond_{\mathcal{L}}\neg\theta$	4 $\Diamond$ I [3]
5	$\bot$	2, 4 $\bot$ I [3,Γ]
6	$\neg\neg\theta$	3-5 $\neg$I [Γ]
7	θ	6 $\neg$E [Γ]