A Constructive Approach to the Resource Semantics of Substructural Logics

Jason Reed and Frank Pfenning

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Abstract. We propose a constructive approach to the resource semantics of substructural logics via proof-preserving translations into a fragment of focused first-order intuitionistic logic with a preorder. Using these translations, we can obtain uniform proofs of cut admissibility, identity expansion, and the completeness of focusing for a variety of logics. We illustrate our approach on linear, ordered, and bunched logics.

Key words: Resource Semantics, Focusing, Substructural Logics, Constructive Logic

1 Introduction

Substructural logics derive their expressive power from subverting our usual intuitions about hypotheses. While assumptions in everyday reasoning are things can reasonably be used many times, used zero times, and used in any order, hypotheses in linear logic [Gir87] must be used exactly once, in affine logic at most once, in relevance logic at least once, in ordered logics [Lam58,Pol01] in a specified order, and in bunched logic [OP99] they must obey a discipline of occurrence and disjointness more complicated still.

The common device these logics employ to enforce restrictions on use of hypotheses is a *structured context*. The collection of active hypotheses is in linear logic a multiset, in ordered logic a list, and in bunched logic a tree. Similar to, for instance, display logic [Bel82] we aim to show how to unify diverse substructural logics by isolating and reasoning about the algebraic properties of their context's structure. Unlike these other approaches, we so do without introducing a new logic that itself has a sophisticated notion of structured context, but instead by making use of the existing concept of *focused* proofs [And92] in a more familiar nonsubstructural logic. We in fact give a translation of various substructural logics into in a rather simple fragment of focused first-order intuitionistic logic, equipped with a binary relation and algebraic operations on its first-order domain particular to the substructural logic being translated.

Such a translation constitutes a novel *constructive resource semantics* for substructural logics. A resource semantics generally gives the meaning of a proposition as a statement in a (typically classical) ambient logic of mathematical definitions concerning *resource labels* (sometimes called 'worlds' because of their role similar to that of Kripke worlds in the semantics of modal logics) that belong to a label algebra that reflects the structure of substructural context. Our approach is similar, except that our representation language is a focused, constructive proof system, so we are able to formulate and prove much stronger claims about the semantics. Not only does provability correspond to provability back and forth across translation, but proofs correspond bijectively to proofs, and focusing phases to focusing phases.

Our present approach generalizes prior work [Ree07] that can be seen as achieving this program for just the negative (in the sense of Girard) fragment of linear logic, where only negative connectives such as implication and additive conjunction are allowed. The key insight to the design of a system that handles all logical connectives, both negative and positive, is that one needs not only a notion of *worlds* to label resource-like hypotheses, but also a dual notion of what we call *frames* to label conclusions.

For clarity we hereafter refer to the substructural logic being encoded as the *object language* and the language it is encoded into as the *representation language*. In the following sections, we first describe the representation language, and give as a central example how to encode focused (intuitionistic, as is the case for all object languages considered here) linear logic into it. Subsequently we more briefly discuss how unfocused linear logic, ordered logic, and bunched logic can also be encoded.

2 The logic FF

Our representation language is a logic called FF, for **F**ocused **F**irst-order intuitionistic logic.

The notion of *focusing*, introduced by Andreoli [And92], is a way of narrowing eligible proofs down to those that decompose connectives in maximal contiguous runs of logical connectives of the same *polarity* (a trait of propositional connectives which divides those that can be eagerly decomposed as goals from those that can be eagerly decomposed as assumptions, among other properties) while remaining complete compared to ordinary proofs: there is a focused proof of a proposition iff there is an ordinary proof, but there are generally fewer *distinct* focused proofs. It is by using the tight control over proof search and proof identity that focusing permits that we are able to faithfully mimic not only of *which* propositions are provable the object language, but *how* they are proved.

The representation language is parametrized over the exact structure of its first-order domain. The data pertaining to the object language we must specify are:

1. A collection of function symbols that give the syntax of two syntactic sorts of worlds p and frames f. These syntactic sorts include at minimum world variables α and frame variables ϕ respectively, which arise from universal quantifiers.
2. A specification of a reflexive, transitive relation $\sqsubseteq$ on pairs $f \lhd p$ pairs of a world and a frame. These pairs are called *structures*.

The basic intuition is that world variables are abstract labels for substructural *hypotheses*, and frame variables are labels for *conclusions*. The algebraic manipulation of expressions built out of these labels abstractly recreates the behavior of substructural contexts. A world expression p describes a context of potentially many hypotheses, a frame expression describes a sequent with a propositional hole in its context, and a structure $f \lhd p$ describes the shape of an entire sequent of the object language. For example, in linear logic, the expression $\phi \lhd (\alpha_1 * \cdots * \alpha_n)$ will represent a sequent with a conclusion labelled ϕ and n assumptions each labelled α_i for $i \in \{1, \ldots, n\}$. The preorder $\sqsubseteq$ says of two sequent shapes that one is deductively weaker than or equal to the other, for example, one might have more hypotheses amenable to weakening than the other. In fact, for most of this paper we will consider relations $\sqsubseteq$ which are symmetric as well, and will be written $\equiv$. In this case, the relation simply tells which context structures are *equivalent*.

2.1 Syntax

The syntax of the representation language is as follows.

Negative Props A	$::= B \Rightarrow A \mid A \wedge A \mid \top \mid \forall\alpha.A \mid \forall\phi.A \mid s^-$
Positive Props B	$::= s^+ \mid \downarrow A$
Negative Atoms s^-	$::= f \lhd p$
Positive Atoms s^+	$::= a^- @ f \mid a^+ @ p$
Worlds p	$::= \alpha \mid \cdots$
Frames f	$::= \phi \mid \cdots$
Contexts Γ	$::= \cdot \mid \Gamma, B$

The bulk of the propositions are built of *negative* logical connectives: implication, conjunction, truth, universal quantification, and negative atomic propositions s^-. The first argument to implication is as usual a *positive* proposition. Ordinarily these might include existential quantification and disjunction, but for our purposes we only need positive atomic propositions s^+, and an inclusion of negative propositions back into positives, via the shift operator $\downarrow$ that interrupts focus phases.

Negative atomic propositions s^- consist of the notion of *structures* described above, a pair of a frame and a world, and positive atoms are one of two kinds of pairs, where one element is an *object-language* negative or positive atomic proposition, written respectively a^- or a^+, and the other element is a frame or world.

As stated above, the syntax of worlds and frames is specified per object language — the representation language works uniformly regardless of what it is. It is here left open except to note again that we must at least include in expressions the ability to use variables.

Contexts Γ are built out of positive propositions, and are themselves not at all substructural: they are subject to tacit weakening, contraction, and exchange.

The judgments of the system are:

Right Focus	$\Gamma \vdash [B]$
Left Focus	$\Gamma; [A] \vdash s^-$
Right Inversion	$\Gamma \vdash A$
Structure Relation	$s^- \sqsubseteq s_0^-$

The focus judgments are used when we have selected a proposition and have committed to continue decomposing it until we reach a polarity shift. Inversion takes place here when we are trying to prove a negative proposition, and we apply right rules eagerly, because all right rules for negative propositions are characteristically invertible.

For uniformity, we write $\Gamma \vdash J$ to stand for either $\Gamma; [A] \vdash s^-$ or $\Gamma \vdash A$. and on occasion when we need to contrast the judgment of FF with that of the object language, decorate the turnstile as $\vdash_{\text{FF}}$.

2.2 Proofs

The valid deductions of this judgment are defined inductively by the inference rules as follows:

$$\frac{}{\Gamma, s^+ \vdash [s^+]}\, s^+\mathsf{R} \qquad \frac{s^- \sqsubseteq s_0^-}{\Gamma; [s^-] \vdash s_0^-}\, s^-\mathsf{L} \qquad \frac{\Gamma \vdash A}{\Gamma \vdash [{\downarrow}A]}\, {\downarrow}\mathsf{R} \qquad \frac{\Gamma, {\downarrow}A; [A] \vdash s^-}{\Gamma, {\downarrow}A \vdash s^-}\, {\downarrow}\mathsf{L} \qquad \frac{}{\Gamma \vdash \top}\, \top R$$

$$\frac{\Gamma \vdash A}{\Gamma \vdash \forall\alpha.A}\, \forall R^{\alpha} \qquad \frac{\Gamma; [\{p/\alpha\}A] \vdash s^-}{\Gamma; [\forall\alpha.A] \vdash s^-}\, \forall L \qquad \frac{\Gamma, B \vdash A}{\Gamma \vdash B \Rightarrow A}\, {\Rightarrow}R \qquad \frac{\Gamma \vdash [B] \quad \Gamma; [A] \vdash s^-}{\Gamma; [B \Rightarrow A] \vdash s^-}\, {\Rightarrow}L$$

$$\frac{\Gamma \vdash A}{\Gamma \vdash \forall\phi.A}\, \forall R^{\phi} \qquad \frac{\Gamma; [\{f/\phi\}A] \vdash s^-}{\Gamma; [\forall\phi.A] \vdash s^-}\, \forall L \qquad \frac{\Gamma \vdash A_1 \quad \Gamma \vdash A_2}{\Gamma \vdash A_1 \wedge A_2}\, \wedge R \qquad \frac{\Gamma; [A_i] \vdash s^-}{\Gamma; [A_1 \wedge A_2] \vdash s^-}\, \wedge L_i$$

They are mostly standard, but we make some comments about the behavior of focusing: when we are focused on a negative atomic proposition s^-, the current conclusion s_0^- must be related to s^- according to the relation $\sqsubseteq$; when focused on a positive atom s^+, that atom must already be found in the current context. Encountering ${\downarrow}A$ on the right blurs focus, and begins inversion of A. Decomposing $\downarrow$ on the left begins a focus phase, which is only allowed when the conclusion has 'stabilized' to a negative atomic proposition s^-.

The right rules for the quantifiers are assumed to have the standard side-conditions about the freshness of variables they introduce. We write $\{p/\alpha\}$ and $\{f/\phi\}$ for substitution of a world or frame expression for the appropriate variable. We may also without any further difficulty support more universal quantifiers defined in the same uniform way over new first-order domains if necessary; an application of this generality can be found in Section 4.

2.3 Metatheory

This calculus satisfies the usual pair of properties that establish its internal soundness (cut admissibility) and internal completeness (identity expansion). Because of the relation allowed at negative atoms, we must first show a form of monotonicity with respect to the preorder.

Lemma 1 (Monotonicity). *Suppose* $s^- \sqsubseteq s_0^-$.

1. *If* $\Gamma \vdash s^-$, *then* $\Gamma \vdash s_0^-$
2. *If* $\Gamma; [A] \vdash s^-$, *then* $\Gamma; [A] \vdash s_0^-$

Proof. By induction on the derivation. Use transitivity for the case s^-L. ∎

The admissibility of cut now follows.

Theorem 1 (Cut Admissibility). *The following rules are admissible:*

$$\frac{\Gamma \vdash [B] \qquad \Gamma, B \vdash J}{\Gamma \vdash J} \qquad \frac{\Gamma \vdash A \qquad \Gamma, {\downarrow}A \vdash J}{\Gamma \vdash J} \qquad \frac{\Gamma \vdash A \qquad \Gamma; [A] \vdash s^-}{\Gamma \vdash s^-}$$

Proof. By a standard structural cut admissibility proof, using lexicographic induction on the cut formula A and the derivations involved. In the first rule, if B is an atom we are done, by the admissibility of contraction. Otherwise analyze and inductively decompose the second premise. In the third rule, both premises are decomposed in lockstep. ∎

We can also separately obtain the result that shows every proposition (not just an atom) entails itself:

Theorem 2 (Identity Expansion). $\Gamma, {\downarrow}A \vdash A$ *and* $\Gamma, B \vdash [B]$ *for all* Γ, A, B.

Proof. Deferred to the appendix.

3 Encoding Linear Logic

In this section we show how to encode focused linear logic into FF in a proof-preserving way. This result is equally applicable, then, to ordinary linear logic; we need only apply the usual *polarization* of an unpolarized linear logic proposition in a way that focused proof search on the result reproduces proof search on the original proposition.

3.1 Focused Linear Logic

In focused linear logic, the propositions are also polarized into negative propositions, and positive, with polarity shift operators $\uparrow$ and $\downarrow$ passing between them. Their syntax is as follows.

Negatives	$N ::= \uparrow P \mid N \& N \mid \top \mid P \multimap N \mid a^-$
Positives	$P ::= \downarrow N \mid P \otimes P \mid 1 \mid P \oplus P \mid 0 \mid a^+$
Left Stable	$\bar{N} ::= N \mid a^+$
Right Stable	$\bar{P} ::= P \mid a^-$
Linear Contexts	$\Delta ::= \cdot \mid \Delta, \bar{N}$
Asynch. Contexts	$\Omega ::= \cdot \mid \Omega, P$

Because positive atoms a^+ and negative propositions N are both 'stable' on the left since we can perform no more inversion on them there, we group them together and denote them as $\bar{N}$ — they are precisely the propositions that remain in linear contexts. Conversely negative atoms a^- and positive propositions P, are stable on the right, as conclusions. The exponential $!$ is considered in Section 3.4.

The judgments of focused linear logic are:

Right Focus	$\Gamma \vdash [P]$
Left Focus	$\Gamma; [N] \vdash \bar{P}$
Right Inversion	$\Gamma; \Omega \vdash N$
Left Inversion	$\Gamma; \Omega \vdash \bar{P}$

Where inversion utilizes an ordered context Ω decompose positive connectives in a fixed order. The inference rules defining focused proof search are fairly standard. If the reader is not familiar with them, they are provided in the appendix in Section A.4.

3.2 Encoding Focused Linear Logic

To encode focused linear logic into FF, we must first 'instantiate' FF by choosing how worlds and frames are built, and what the relation $\sqsubseteq$ on them is. The eventual goal, the *adequacy theorem* for the encoding, is to get out of these choices a compositional bijection between linear logic proofs of a given sequent, and FF proofs of its translation. This bijection will be obtained as the computational content of the theorem that proofs can be translated back and forth between the object and representation languages.

We define the syntax of worlds and frames as follows:

Worlds	$p ::= \alpha \mid \epsilon \mid p * p$
Frames	$f ::= \phi \mid f \circledast p$

Worlds are built from variables, the empty world (which specifies no resources), and a binary operator $*$ of *resource combination*. Frames consist of frame variables, and an operation $f \circledast p$ that adjoins a world p to f. This operation can be seen as the permission, when using a linear logic left rule, to choose a proposition to decompose from inside the linear context, and to set all the other propositions (represented by p) aside.

For the relation $\sqsubseteq$ we take a notion of equivalence $\equiv$ on these expressions and structures $f \lhd p$, determined by associativity, commutativity, and unit laws on $*$, and an additional associativity property that captures the meaning of $\circledast$. Specifically, the relation $\equiv$ is axiomatized by

$$\epsilon * p \equiv p \qquad p * q \equiv q * p \qquad p * (q * r) \equiv (p * q) * r \qquad (f \circledast p) \triangleleft q \equiv f \triangleleft (q * p)$$

plus symmetry, reflexivity, and transitivity of $\equiv$, and all congruence laws as expected; for example, $p * q \equiv p' * q'$ when $p \equiv p'$ and $q \equiv q'$.

Since in this case the relation $\equiv$ was chosen to be symmetric, we can simplify proofs by positing that all worlds, frames, and negative atoms s^- are identified up to $\equiv$, and lifting this identification to all propositions. It it easy to see that sequents that are the same up to $\equiv$ on the atoms within them have isomorphic proofs.

For the encoding we will speak of *continuations* k, which are defined to be abstractions $\alpha.A$, that is, an FF proposition abstracted over a bound world variable α. We write $k(p)$ for application of k to an argument p. By definition $(\alpha.A)(p)$ is the same as $\{p/\alpha\}A$.

Encoding Propositions We define a pair of mutually recursive functions as follows from focused linear logic propositions to negative propositions of FF. We translate N with respect to a world p, and write the translation $(N@p)$, and we translate P with respect to a continuation k, written $(P@k)$.

P	$(P@k)$
0	$\top$
$P_1 \oplus P_2$	$(P_1@k) \wedge (P_2@k)$
1	$k(\epsilon)$
$P_1 \otimes P_2$	$(P_2@\alpha.(P_1@\beta.k(\alpha * \beta)))$
a^+	$\forall\alpha.a^+@\alpha \Rightarrow k(\alpha)$
$\downarrow N$	$\forall\alpha.\downarrow(N@\alpha) \Rightarrow k(\alpha)$

N	$(N@p)$
$\top$	$\top$
$N_1 \,\&\, N_2$	$(N_1@p) \wedge (N_2@p)$
$P \multimap N$	$(P@\alpha.(N@(p * \alpha)))$
a^-	$\forall\phi.a^-@\phi \Rightarrow \phi \triangleleft p$
$\uparrow P$	$\forall\phi.\downarrow(P@\alpha.(\phi \triangleleft \alpha)) \Rightarrow \phi \triangleleft p$

This gives a resource semantics for linear logic. The expression $(N@p)$ describes the truth of N 'at world p', or 'given resources p'. We see sensibly that an additive conjunction $N_1 \,\&\, N_2$ is true at a world iff both its conjuncts are true at that world. The interpretation of positive connectives is given dually in terms of their role as assumptions. A positive proposition passes to its world-continuation a world describing what resources it stands for. Thus we see that the meaning of $P \multimap N$ at world p is the meaning of N at the larger world $p * \alpha$, under the supposition that P is true given resources α. Resource combination also takes place in the interpretation of $\otimes$: the resources passed to the final continuation for $P_1 \otimes P_2$ are those passed by P_1 combined, via $*$, with those given by P_2.

We can also see the dual nature of positives in that the disjunctive connectives in linear logic are translated as conjunctions in FF, and in that the polarity shift operators effect a species of negation in both directions.

Note that for a positive object-language atom a^+, the expression $(a^+@k)$ is an invocation of the translation function, while $(a^+@p)$ is a representation language positive atom. Conversely, $(a^-@p)$ is a call to translation, but $(a^-@f)$ is an atom. We take advantage of this overloading the definitions below.

Before discussing how to translate sequents with contexts in full generality, we can already give a suggestive example of the way in which linear logic derivations resemble their counterparts in FF: for as a special case a proof $\vdash N$ of a single negative proposition in the empty context will correspond to a proof $\vdash_{\mathrm{FF}} (N@\epsilon)$ of the translation of N at the empty world. So we can compare for example a proof of $\downarrow\uparrow 1 \otimes \downarrow\uparrow 1 \multimap \uparrow 1$ and of its translation: (eliding some intermediate asynchronous steps)

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\cdot \vdash [1]}}{\cdot \vdash 1}}{[\uparrow 1] \vdash 1}}{\uparrow 1 \vdash 1}}{\uparrow 1; [\uparrow 1] \vdash 1}}{\dfrac{\uparrow 1, \uparrow 1 \vdash 1}{\vdash \downarrow\uparrow 1 \otimes \downarrow\uparrow 1 \multimap \uparrow 1}}
$$

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\phi \triangleleft \epsilon \equiv \phi \triangleleft \epsilon}{[\phi \triangleleft \epsilon] \vdash \phi \triangleleft \epsilon}}{\Gamma \vdash \phi \triangleleft \epsilon}}{\Gamma \vdash [\downarrow \phi \triangleleft \epsilon]} \quad \dfrac{\phi \triangleleft \beta \equiv (\phi \circledast \beta) \triangleleft \epsilon}{\Gamma; [\phi \triangleleft \beta] \vdash (\phi \circledast \beta) \triangleleft \epsilon}}{\Gamma; [\downarrow \phi \triangleleft \epsilon \Rightarrow \phi \triangleleft \beta] \vdash (\phi \circledast \beta) \triangleleft \epsilon}}{\Gamma; [(\uparrow 1 @ \beta)] \vdash (\phi \circledast \beta) \triangleleft \epsilon}}{\Gamma \vdash (\phi \circledast \beta) \triangleleft \epsilon}}{\Gamma \vdash [\downarrow (\phi \circledast \beta) \triangleleft \epsilon]} \quad \dfrac{(\phi \circledast \beta) \triangleleft \alpha \equiv \phi \triangleleft (\alpha * \beta)}{\Gamma; [(\phi \circledast \beta) \triangleleft \alpha] \vdash \phi \triangleleft (\alpha * \beta)}}{\dfrac{\dfrac{\dfrac{\Gamma; [\downarrow (\phi \circledast \beta) \triangleleft \epsilon \Rightarrow (\phi \circledast \beta) \triangleleft \alpha] \vdash \phi \triangleleft (\alpha * \beta)}{\Gamma; [(\uparrow 1 @ \alpha)] \vdash \phi \triangleleft (\alpha * \beta)}}{\downarrow(\uparrow 1 @ \alpha), \downarrow(\uparrow 1 @ \beta), \downarrow \phi \triangleleft \epsilon \vdash \phi \triangleleft (\alpha * \beta)}}{\vdash \forall \alpha. \downarrow(\uparrow 1 @ \alpha) \Rightarrow \forall \beta. \downarrow(\uparrow 1 @ \beta) \Rightarrow (\uparrow 1 @ (\alpha * \beta))}}
$$

where Γ abbreviates $\downarrow(\uparrow 1@\alpha), \downarrow(\uparrow 1@\beta), \downarrow\phi \triangleleft \epsilon$, and frequently only lazily expanding the translation of $\uparrow 1$ to save space.

The way the FF proof works is that the negative atom in its conclusion tracks the shape of the linear logic sequent — in particular the shape of its linear context. All linear logic hypotheses (resp. the conclusion) appear as FF hypotheses, each translated at a distinct world (resp. a frame) variable. This is made more explicit in the translation of contexts below.

Observe that there are an equal number of focus phases in both proofs, in this case three. In linear logic we focus twice on a resource of $\uparrow 1$, and in FF we focus twice on its translation, first at world α, then at β, each time choosing (via the $\forall L$ rule) a frame, first $\phi \circledast \beta$, then ϕ, in which that world occurs — that is, which satisfies the appropriate relation $\equiv$ with the current context structure, namely the negative atom that is the FF conclusion — and then replace that world with ϵ, modeling the deletion of 1 from the context. Finally, the third focus phase checks that the context is empty in linear logic, and in FF checks that the current context structure is $\equiv$ to the structure $\phi \triangleleft \epsilon$.

Encoding Sequents The translation of a linear logic sequent will involve an FF context

$$(\bar{N}_1@\alpha_1),\ldots,(\bar{N}_n@\alpha_n),(\bar{P}_1@\phi_1),\ldots,(\bar{P}_m@\phi_m)$$

recalling that $\bar{N}$ and $\bar{P}$ allow for opposite-polarity atoms, where $(\bar{P}@\phi)$ in case $\bar{P}$ is a positive proposition P is defined to be the lifting of ϕ to a continuation $(P@\alpha.(\phi \triangleleft \alpha))$. We say a context Γ is *regular* if it is of this form. Such a Γ represents a collection of object language hypotheses and conclusions, each uniquely labelled by a distinct world or frame variable, respectively. Since Γ is itself not substructural, object language hypotheses and (even multiple) conclusions persist in it during bottom-up proof search: it is only the world and frame discipline that prevents resources from being inappropriately reused.

Specifically, a world, frame, or structure can be used to select from Γ a collection of substructural hypotheses and/or conclusion, according to the following definition. Suppose Γ is regular, and that Δ is the linear logic context $\bar{N}_{i_1},\ldots,\bar{N}_{i_\ell}$ for distinct $i_1,\ldots,i_\ell \in \{1,\ldots,n\}$. Then we write:

1. $\Gamma \sim_p \Delta$ iff $p \equiv \alpha_{i_1} * \cdots * \alpha_{i_\ell}$
2. $\Gamma \sim_f (\Delta \vdash P_i)$ iff $f \equiv \phi_i \circledast (\alpha_{i_1} * \cdots * \alpha_{i_\ell})$
3. $\Gamma \sim_{f \triangleleft p} (\Delta \vdash P_i)$ iff $f \triangleleft p \equiv \phi_i \triangleleft (\alpha_{i_1} * \cdots * \alpha_{i_\ell})$

We must also account for how asynchronous contexts Ω are translated. Define the function $((\Omega; N)@p)$, which yields a proposition from Ω, N, p, and describes inversion of N on the right with Ω on the left, at world p. It is defined by

$$((\Omega, P; N)@p) = (P@\alpha.((\Omega; N)@(p * \alpha))) \qquad ((\cdot; N)@p) = (N@p)$$

The function $(\Omega@f)$ yields a proposition from Ω, f, and describes the inversion of Ω on the left, at frame f. It is defined by

$$((\Omega, P)@f) = (P@\alpha.(\Omega@(f \circledast \alpha))) \qquad (\cdot@f) = f \triangleleft \epsilon$$

3.3 Adequacy

Recall that to show the encoding is correct, we seek a compositional bijection between proofs before and after translation. This bijection is obtained as the computational content of the theorem that proofs can be so translated. There are five parts to the translation, corresponding to the five states of focused proof search: neutral sequents, negative focus, positive focus, negative inversion, and positive inversion.

Theorem 3 (Adequacy). *Suppose Γ is regular.*

1. $\Gamma \vdash_{FF} s^-$ *iff there are* $\Delta, \bar{P}$ *such that* $\Gamma \sim_{s^-} (\Delta \vdash \bar{P})$ *and* $\Delta \vdash \bar{P}$.
2. $\Gamma[(N@p)] \vdash_{FF} s^-$ *iff there are* $f, \Delta, \bar{P}$ *such that* $\Gamma \sim_f (\Delta \vdash \bar{P})$ *and* $\Delta[N] \vdash \bar{P}$ *and* $f \triangleleft p \equiv s^-$.
3. $\Gamma[(P@k)] \vdash_{FF} s^-$ *iff there are* p, Δ *such that* $\Gamma \sim_p \Delta$ *and* $\Delta \vdash [P]$ *and* $\Gamma[k(p)] \vdash_{FF} s^-$.

4. $\Gamma \vdash_{FF} ((\Omega; N)@p)$ *iff there is* Δ *such that* $\Gamma \sim_p \Delta$ *and* $\Delta \vdash \Omega > N$.
5. $\Gamma \vdash_{FF} (\Omega@f)$ *iff there are* $\Delta, \bar{P}$ *such that* $\Gamma \sim_f (\Delta \vdash \bar{P})$ *and* $\Delta \vdash \Omega > \bar{P}$.

Proof. Deferred to the appendix.

Theorem 4 (Adequacy of Provability). *The function from derivations of* $\Gamma \vdash_{FF} s^-$ *and derivations of* $\Delta \vdash \bar{P}$ *such that* $\Gamma \sim_{s^-} (\Delta \vdash \bar{P})$ *given by the constructive proof of the previous theorem is a bijection.*

Proof. By structural induction on the proof that the previous theorem yields a given result. ■

Corollary 1. *Focused linear logic satisfies cut admissibility and identity expansion.*

Proof. By appeal to the same properties of FF, for we have just shown that focused linear logic is faithfully embedded there. ■

3.4 Adding the Exponential

An advantage of our approach is that adding the exponential $!$ to the existing translation is quite easy. Syntactically, we add to linear logic $\cdots \mid !N$, and translate it via saying $(!N@k) = (N@\epsilon) \Rightarrow k(\epsilon)$.

The intuition is that propositions under bang should be true *absolutely*, with respect to *no* resources, and ϵ represents exactly this lack of resources. We can see immediately that the translations of $!\top$ and 1 coincide (up to eliminating a vacuous $\top \Rightarrow$) and so too $!N_1 \otimes !N_2$ and $!(N_1 \mathbin{\&} N_2)$, because of the equivalence $\epsilon \equiv \epsilon * \epsilon$. The adequacy theorem can be extended straightforwardly to relate this encoding of the exponential with intuitionistic linear logic with exponential in judgmental style [Bar97,CCP03].

3.5 Unpolarized Linear Logic

We can also represent ordinary unpolarized linear logic, and the usual unfocused sequent calculus for it, by composing our translation above with the well-known embedding of unpolarized propositions into polarized ones that inserts shift operators between every connective. The details of this translation are routine, and deferred to the appendix.

To show that the unfocused and focused proof systems are equivalent — that is, to show the completeness of focusing for linear logic — we need only show that double-shifts, the unnecessary pauses in focus, are eliminable.

Lemma 2. $(\uparrow\downarrow N@p) \dashv\vdash (N@p)$ *and* $(\downarrow\uparrow P@k) \dashv\vdash (P@k)$.

This follows by a straightforward induction on N and P, constructing derivations in FF.

This method is even more useful for logics, like the two described below, for which a focusing system is not as well known, or not known at all. For then we can show adequacy of the object language with respect to a polarizing translation that inserts shifts everywhere, and simply *read off* a complete focusing system for the object language from what is obtained by removing redundant double-shifts.

4 Encoding Ordered Logic

The study of noncommutative logics goes back to Lambek [Lam58], but we take the more modern treatment in Polakow's thesis [Pol01] as a reference point, which includes modalities ! and ¡ that allow unrestricted and mobile (i.e. satisfying the structural law of exchange) hypotheses to be used alongside ordered hypotheses. Although focused proof search for ordered logic is not as well-known as for linear logic, it is implicit in, taken together, the ordered logic programming systems of Pfenning and Simmons [PS09] and Polakow [Pol00].

Without its exponentials, the encoding of ordered logic would be almost entirely the same that of linear logic, except that we would drop the axiom that makes $*$ commutative. We wish to illustrate, however, how to accommodate the variety of different substructural hypotheses present in full ordered logic.

The syntax of polarized propositions in ordered logic is

Negatives	$N ::= \uparrow P \mid N \mathbin{\&} N \mid \top \mid P \rightarrowtail N \mid P \twoheadrightarrow N \mid a^-$
Positives	$P ::= \downarrow N \mid P \bullet P \mid 1 \mid P \oplus P \mid 0 \mid a^+ \mid !N \mid ¡N$

The judgment of ordered logic has three zones of hypotheses, an unrestricted Γ, a linear Δ, and an ordered Ω. To encode this structure, we take the syntax of worlds and frames to be defined by

Linear Worlds	$\tilde{p} ::= \tilde{\alpha} \mid \tilde{\epsilon} \mid \tilde{p} * \tilde{p}$
Worlds	$p ::= \alpha \mid \epsilon \mid p \cdot p \mid \iota\tilde{p}$
Frames	$f ::= \phi \mid f \odot p \mid p \odot f$

with an injection ι from linear worlds into ordered worlds, with $\equiv$ axiomatized by

$$\epsilon \cdot p \equiv p \qquad p \cdot (q \cdot r) \equiv (p \cdot q) \cdot r \qquad (f \odot p) \lhd q \equiv f \lhd (q \cdot p) \qquad (p \odot f) \lhd q \equiv f \lhd (p \cdot q)$$

$$\tilde{\epsilon} * \tilde{p} \equiv \tilde{p} \qquad \tilde{p} * \tilde{q} \equiv \tilde{q} * \tilde{p} \qquad \tilde{p} * (\tilde{q} * \tilde{r}) \equiv (\tilde{p} * \tilde{q}) * \tilde{r} \qquad \iota\tilde{\epsilon} \equiv \epsilon \qquad \iota(\tilde{p} * \tilde{q}) \equiv \iota\tilde{p} \cdot \iota\tilde{q}$$

which makes $*$ commutative, $\cdot$ generally noncommutative, (but commutative on the range of ι) and which makes ι an algebra homomorphism. We have essentially axiomatized a free noncommutative monoid with a commutative submonoid. The language of frames has expanded not because of the modalities, but just because order introduces the possibility that left rules take place to the left or right of other, uninvolved resources.

The translation of propositions for ordered logic is in Figure 1. and we may rehearse suitable versions of all of the adequacy results proved above for linear logic.

A context $\Gamma; \Delta; \Omega$ is then modeled by a world expression that translates $N \in \Gamma$ to $(N @ \epsilon)$, and $N \in \Delta$ to $(N @ \iota\tilde{\alpha})$ (for a fresh $\tilde{\alpha}$) and $N \in \Omega$ to $(N @ \alpha)$ (for a fresh α).

The points to take away from this example are that a variety of different sorts of substructural hypotheses can be modeled by different syntactic sorts of

P	$(P @ k)$
0	$\top$
$P_1 \oplus P_2$	$(P_1 @ k) \wedge (P_2 @ k)$
1	$k(\epsilon)$
$P_1 \bullet P_2$	$(P_1 @ \alpha.(P_2 @ \beta.k(\alpha \cdot \beta)))$
a^+	$\forall \alpha. a^+ @ \alpha \Rightarrow k(\alpha)$
$\downarrow N$	$\forall \alpha. \downarrow(N @ \alpha) \Rightarrow k(\alpha)$
$¡N$	$\forall \tilde{\alpha}.(N @ \iota\tilde{\alpha}) \Rightarrow k(\iota\tilde{\alpha})$
$!N$	$(N @ \epsilon) \Rightarrow k(\epsilon)$

N	$(N @ p)$
$\top$	$\top$
$N_1 \,\&\, N_2$	$(N_1 @ p) \wedge (N_2 @ p)$
$P \rightarrowtail N$	$(P @ \alpha.(N @ (\alpha \cdot p)))$
$P \twoheadrightarrow N$	$(P @ \alpha.(N @ (p \cdot \alpha)))$
a^-	$\forall \phi. a^- @ \phi \Rightarrow \phi \triangleleft p$
$\uparrow P$	$\forall \phi. \downarrow(P @ \alpha.\phi \triangleleft \alpha) \Rightarrow \phi \triangleleft p$

Fig. 1. Translation of Ordered Logic

algebraic expressions, and that modalities that mediate between them such as ¡ operate by being essentially a shift operator that quantifies over a different type of variable — even ! can be construed as quantifying over exactly the worlds (of which there is only one) that *are equal to* ϵ.

5 Bunched Logic

The logic of bunched implications [OP99] features two ways of combining contexts, one multiplicative, denoted by a comma, and one additive, denoted by a semicolon, which lead to contexts being tree-shaped, since the two operations, while both separately associative, do not distribute over one another. There are structural rules to make these operations behave appropriately, the bunched versions of weakening, contraction, and a generalized version of exchange

$$\frac{\Gamma(\Delta;\Delta) \vdash A}{\Gamma(\Delta) \vdash A} \qquad \frac{\Gamma(\Delta_1) \vdash A}{\Gamma(\Delta_1;\Delta_2) \vdash A} \qquad \frac{\Delta_1 \vdash A \qquad \Delta_1 \equiv \Delta_2}{\Delta_2 \vdash A}$$

for Γ being any bunched logic context-with-hole (exactly what our notion of *frames* are meant to represent) and where $\equiv$ is equivalence of bunched contexts, which includes for example associativity and commutativity of ',' and ';'.

We can make use of the full generality of allowing FF atoms to be compared by a non-symmetric relation $\sqsubseteq$ to account for the asymmetry of the first two of these three rules. We say that polarized bunched logic propositions are given by

Negatives	$N ::= \uparrow P \mid N \wedge N \mid \top \mid P \to N \mid P \mathrel{-\!\!*} N \mid a^-$
Positives	$P ::= \downarrow N \mid P * P \mid I \mid P \vee P \mid \bot \mid a^+$

and then we choose the syntax of worlds and frames to be

Worlds	$p ::= \alpha \mid \epsilon_* \mid \epsilon_\wedge \mid p * p \mid p \wedge p$
Frames	$f ::= \phi \mid f \circledast p \mid f \owedge p$

(where multiplicative operations are marked with $*$ and additives marked with $\wedge$) with $\sqsubseteq$ axiomatized by (where $p \equiv q$ means $p \sqsubseteq q$ and $q \sqsubseteq p$)

$$\epsilon_* * p \equiv p \qquad p * (q * r) \equiv (p * q) * r \qquad (f \circledast p) \triangleleft q \equiv f \triangleleft (q * p)$$

$$\epsilon_\wedge \wedge p \equiv p \qquad p \wedge (q \wedge r) \equiv (p \wedge q) \wedge r \qquad (f \owedge p) \triangleleft q \equiv f \triangleleft (q \wedge p)$$

$$f \triangleleft (p \wedge p) \sqsubseteq f \triangleleft p \qquad f \triangleleft p \sqsubseteq f \triangleleft (p \wedge q)$$

These last two axioms are effectively direct rewritings in our syntax of the contraction and weakening rules. The remaining axioms correspond directly to the axiomatization of $\equiv$ in the BI literature. The translation on propositions is then simply

P	$(P@k)$	N	$(N@p)$
$\bot$	$\top$	$\top$	$\top$
$P_1 \vee P_2$	$(P_1@k) \wedge (P_2@k)$	$N_1 \wedge N_2$	$(N_1@p) \wedge (N_2@p)$
1	$k(\epsilon_*)$	$P \to N$	$(P@\alpha.(N@(p \wedge \alpha)))$
$P_1 * P_2$	$(P_1@\alpha.(P_2@\beta.k(\alpha * \beta)))$	$P \mathrel{-\!\!*} N$	$(P@\alpha.(N@(p * \alpha)))$
a^+	$\forall\alpha.a^+@\alpha \Rightarrow k(\alpha)$	a^-	$\forall\phi.a^-@\phi \Rightarrow \phi \triangleleft p$
$\downarrow N$	$\forall\alpha.\downarrow(N@\alpha) \Rightarrow k(\alpha)$	$\uparrow P$	$\forall\phi.\downarrow(P@\alpha.\phi \triangleleft \alpha) \Rightarrow \phi \triangleleft p$

and we go through a similar series of adequacy theorems as before. The main difference in this case is that between every focusing phase (and if we target unfocused BI by polarizing propositions by inserting shifts everywhere, this effectively means between every logical connective decomposed) we must match up via $\sqsubseteq$ the current context and the one instantiating a quantifier coming from a shift operation, and this relationship is now asymmetric and rather nontrivial. This corresponds exactly to the ability in BI to apply structural rules at any point in a derivation.

6 Related Work

There has been a significant amount of work on the Kripke semantics for substructural logics, including the modalities—see Kamide [Kam02] for a systematic study and further references. Similarly, the logic of bunched implication was conceived from the beginning with a resource semantics [OP99]. The classical nature of the metalanguage in which these interpretations are formulated becomes particularly apparent when the objective of the translation is theorem proving: proof search then proceeds in a classical logic where formulas have been augmented in a systematic way to encompass worlds, be it for linear [MO99], noncommutative [GN03], or bunched [GMP05] logic. At the root of these concrete interpretations we can find labeled deduction [BDG+00]. Our work shares with these the idea of resource combination via algebraic operators and partial orders for resource entailment.

The above models capture a notion of truth with respect to resources or worlds. In many applications, however, we are interested in the precise structure of proofs. Besides well-known computational interpretations of proofs, their fine structure also determines the behavior of logic programs where computation proceeds by proof search. Capturing proofs is usually the domain of logical frameworks such as LF [HHP93], or its substructural extensions such as RLF [IP98], LLF [CP02], or OLF [Pol01], where proofs are reified as objects. In case of these frameworks, however, natural representations (namely those mapping the substructural consequence of the object logic to consequence in the metalogic) are limited by the substructural properties of the framework. Moreover, even if (small-step) proofs can be represented in this manner, focused proofs in an object language seem to require even more substructural expressiveness in the metalanguage. Instead of escalating the number of substructural judgments and modalities in the metalanguage, we propose here to slice through the knot using just an intuitionistic framework and capturing substructural properties algebraically. The framework here is first-order, but we conjecture, based on our experience with in HLF [Ree07], that reifying proofs in a dependent version of the present proposal should not present much difficulty.

7 Conclusion

We have presented a method of interpreting substructural logics into a first-order focused constructive logic which preserves proofs and focusing structure. It isolates the algebra of contexts from the machinery of propositional inference, and so allows further comparison between otherwise superficially different logics — although intuitively it is quite clear that linear logic's $\otimes$, ordered logic's $\bullet$, and bunched logic's $*$ and $\wedge$ all internalize an operation on contexts, we can give this intuition formal content by saying precisely that they share the same translation up to a different choice of algebra. In future work, we hope to extend this semantics to classical variants of substructural logics, as well as modal logics such as judgmental S4 [PD01], by analogy with linear $!$. It also seems like it might be possible to unify linear with bunched logic along the same lines as ordered and linear logic, by using a modality to mediate between the two respective algebraic structures.

References

[And92] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347, 1992.

[Bar97] Andrew Graham Barber. *Linear Type Theories, Semantics and Action Calculi*. PhD thesis, University of Edinburgh, 1997. Available as Technical Report ECS-LFCS-97-371.

[BDG+00] D. Basin, M. D'Agostino, D.M. Gabbay, S. Matthews, and L. Viganò, editors. *Labelled Deduction*, volume 17 of *Applied Logic Series*. Kluwer Academic Publishers, 2000.

[Bel82] Nuel Belnap. Display logic. *Journal of philosophical logic*, 11:375–417, 1982.

[CCP03] Boy-Yuh Evan Chang, Kaustuv Chaudhuri, and Frank Pfenning. A judgmental analysis of linear logic. Technical Report CMU-CS-03-131R, Carnegie Mellon University, Department of Computer Science, December 2003.

[CP02] Iliano Cervesato and Frank Pfenning. A linear logical framework. *Information & Computation*, 179(1):19–75, 2002.

[Gir87] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, 1987.

[GMP05] Didier Galmiche, Daniel Méry, and David J. Pym. The semantics of BI and resource tableaux. *Mathematical Structures in Computer Science*, 15(6):1033–1088, 2005.

[GN03] Didier Galmiche and Jean-Marc Notin. Connection-based proof construction in non-commutative logic. In M.Vardi and A.Voronkov, editors, *Proceedings of the 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR'03)*, pages 422–436, Almaty, Kazakhstan, September 2003. Springer LNCS 2850.

[HHP93] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the Association for Computing Machinery*, 40(1):143–184, January 1993.

[IP98] Samin Ishtiaq and David Pym. A relevant analysis of natural deduction. *Journal of Logic and Computation*, 8(6):809–838, 1998.

[Kam02] Norihiro Kamide. Kripke semantics for modal substructural logics. *Journal of Logic, Language and Information*, 11(4):453–470, 2002.

[Lam58] Joachim Lambek. The mathematics of sentence structure. *American Mathematical Monthly*, 65(3):154–170, 1958.

[MO99] Heiko Mantel and Jens Otten. linTAP: A tableau prover for linear logic. In N.Murray, editor, *Proceedings of the International Conference on Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX'99)*, pages 217–231, Saratoga Springs, New York, 1999. Springer LNCS 1617.

[OP99] P.W. O'Hearn and D.J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.

[PD01] Frank Pfenning and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11(4):511–540, 2001.

[Pol00] Jeff Polakow. Linear logic programming with an ordered context. In *Proceedings of the 2nd International Conference on Principles and Practice of Declarative Programming (PPDP'00)*, pages 68–79. ACM Press, 2000.

[Pol01] Jeff Polakow. *Ordered Linear Logic and Applications*. PhD thesis, Carnegie Mellon University School of Computer Science, 2001.

[PS09] Frank Pfenning and Robert J. Simmons. Substructural operational semantics as ordered logic programming. In A.Pitts, editor, *Proceedings of 24th Annual Symposium on Logics in Computer Science (LICS'09)*, Los Angeles, California, August 2009. IEEE. To appear.

[Ree07] Jason Reed. Hybridizing a logical framework. In P. Blackburn, T. Bolander, T. Braner, V. de Paiva, and J. Villadsen, editors, *Proceedings of the International Workshop on Hybrid Logic (HyLo 2006)*, 2007.

A Appendix

A.1 Proof of Identity Expansion

We strengthen the induction hypothesis to:

For all Γ, A,

1. If there is a derivation
$$\begin{array}{c}\Gamma; [A] \vdash s^- \\ \vdots \\ \Gamma \vdash s^-\end{array}$$
parametric in s^-, then $\Gamma \vdash A$.
2. $\Gamma, B \vdash [B]$

The proof of it is by induction on A, with case 1 considered less than case 2 for equal A.

1. Split cases on A. We show some representative cases.

Case: s^-. Use reflexivity of $\sqsubseteq$.
$$\frac{s^- \sqsubseteq s^-}{\Gamma; [s^-] \vdash s^-}$$
Case: $A_1 \,\&\, A_2$. Apply the i.h. to
$$\begin{array}{c}\dfrac{\Gamma; [A_i] \vdash s^-}{\Gamma; [A_1 \,\&\, A_2] \vdash s^-} \\ \vdots \\ \Gamma \vdash s^-\end{array}$$
to get $\Gamma \vdash A_i$, and then observe
$$\frac{\Gamma \vdash A_1 \qquad \Gamma \vdash A_2}{\Gamma \vdash A_1 \,\&\, A_2}$$
Case: $B \Rightarrow A$. Apply the i.h. at A to
$$\begin{array}{c}\dfrac{\dfrac{}{\Gamma, B \vdash [B]}\,i.h. \qquad \Gamma, B; [A] \vdash s^-}{\Gamma, B; [B \Rightarrow A] \vdash s^-} \\ \vdots \\ \Gamma, B \vdash s^-\end{array}$$
(having used the i.h. part 2 at B to see that B entails itself) to obtain $\Gamma, B \vdash A$, and then prove
$$\frac{\Gamma, B \vdash A}{\Gamma \vdash B \Rightarrow A}$$

2. If B we are done. Otherwise, $B = {\downarrow}A$. In that case, blur on the right, and apply i.h. part 1 to the inference rule $\downarrow$ L.

■

A.2 Proof of Adequacy

By lexicographic induction on the object-language proposition (or in parts 4,5, context), and the derivation. Some representative cases:

1. In the forward direction, the only move available is focusing on some proposition in Γ. Since by assumption Γ is regular, it is either of the form $(\bar{N}@\alpha)$ or $(\bar{P}@\phi)$. But we cannot begin focus on a positive atom on the left, so the only possibilities left are $(N@\alpha)$ or $(P@\phi)$. For these apply the induction hypothesis part 2 or 3, respectively.
In the reverse direction, we likewise appeal to induction hypothesis 2 or 3 depending on whether a negative or positive atom is focused on.
2.

Case: $P \multimap N$ in the forward direction. By assumption $\Gamma[(P \multimap N@p)] \vdash_{\text{FF}} s^-$ which means $\Gamma[(P@\alpha.(N@(p * \alpha)))] \vdash_{\text{FF}} s^-$. By induction hypothesis part 3, there are q, Δ_1 such that $\Gamma \sim_q \Delta_1$ and $\Delta_1 \vdash [P]$ and $\Gamma[(N@(p * q))] \vdash_{\text{FF}} s^-$. By the induction hypothesis part 2, there are $f, \Delta_2, \bar{P}$ such that $\Gamma \sim_f (\Delta_2 \vdash \bar{P})$ and $\Delta_2 \vdash [N] > \bar{P}$ and $f \triangleleft (p * q) \equiv s^-$. To satisfy our obligations we produce the frame $f \circledast q$, the combined context (Δ_1, Δ_2) and conclusion $\bar{P}$, the fact that $\Gamma \sim_{f \circledast q} (\Delta_1, \Delta_2 \vdash \bar{P})$ by definition of $\sim$, the fact that $\Delta_1, \Delta_2 \vdash [P \multimap N] > \bar{P}$ by rule application, and that $(f \circledast q) \triangleleft p \equiv f \triangleleft (p * q) \equiv s^-$.

Case: $\uparrow P$ in the forward direction. We have

$$\Gamma[\forall \phi. \downarrow (P@\alpha.(\phi \triangleleft \alpha)) \Rightarrow \phi \triangleleft p] \vdash_{\text{FF}} s^-$$

By inversion, a proof of this chooses f to instantiate ϕ, and contains subderivations of

$$\Gamma \vdash (P@\alpha.(f \triangleleft \alpha))$$

$$\Gamma; [f \triangleleft p] \vdash s^-$$

The second of these will only succeed of $f \triangleleft p \equiv s^-$. We have satisfied part of our obligations by producing the frame f, and this equivalence. For the rest, we appeal to the induction hypothesis part 5, observing that $(P@\alpha.(f \triangleleft \alpha))$ and $(P@f)$ (the latter being the translation of asynchronous context that just happens to have one element) are the same proposition up to $\equiv$.

4.

Case: $\Omega, (P_1 \otimes P_2)$. We have $\Gamma \vdash_{\text{FF}} ((\Omega, (P_1 \otimes P_2); N)@p)$, whose conclusion by definition is

$$\begin{aligned} &= ((P_1 \otimes P_2)@\alpha.((\Omega; N)@(p * \alpha))) \\ &= (P_2@\alpha_2.(P_1@\alpha_1.((\Omega; N)@(p * (\alpha_1 * \alpha_2))))) \\ &\equiv (P_2@\alpha_2.(P_1@\alpha_1.((\Omega; N)@((p * \alpha_2) * \alpha_1))))) \\ &= (P_2@\alpha_2.((\Omega, P_1; N)@(p * \alpha_2))) \\ &= ((\Omega, P_1, P_2; N)@p) \end{aligned}$$

so we may appeal to the induction hypothesis to find Δ such that $\Gamma \sim_p \Delta$ and $\Delta \vdash \Omega, P_1, P_2 > N$, and use rule application to get $\Delta \vdash \Omega, (P_1 \otimes P_2) > N$.

Case: $\Omega = \cdot$ and $N = \uparrow P$. We know $\Gamma \vdash \forall\phi.\downarrow(P@\alpha.(\phi \lhd \alpha)) \Rightarrow \phi \lhd p$. By inversion

$$\Gamma, \downarrow(P@\alpha.(\phi \lhd \alpha)) \vdash \phi \lhd p$$

So we are able to appeal to the induction hypothesis part 1.

■

A.3 Polarizing Linear Logic

The polarization function consists of four mutually recursive pieces, defined in Figure 2. They are RU, LU, $\hat{R}U$, $\hat{L}U$, where U is an unpolarized linear logic proposition from the grammar

$$\text{Propositions } U ::= U \mathbin{\&} U \mid \top \mid U \multimap U \mid a^- \mid U \otimes U \mid 1 \mid U \oplus U \mid 0 \mid a^+$$

The functions R and $\hat{L}$ always yield a positive proposition result, and L and $\hat{R}$ always yield a negative.

U	RU	LU
$U_1 \otimes U_2$	$\downarrow\hat{R}U_1 \otimes \downarrow\hat{R}U_2$	$\uparrow(\hat{L}U_1 \otimes \hat{L}U_2)$
$U_1 \oplus U_2$	$\downarrow\hat{R}U_1 \oplus \downarrow\hat{R}U_2$	$\uparrow(\hat{L}U_1 \oplus \hat{L}U_2)$
1	1	$\uparrow 1$
0	0	$\uparrow 0$
$U_1 \mathbin{\&} U_2$	$\downarrow(\hat{R}U_1 \mathbin{\&} \hat{R}U_2)$	$\uparrow\hat{L}U_1 \mathbin{\&} \uparrow\hat{L}U_2$
$\top$	$\downarrow\top$	$\top$
$U_1 \multimap U_2$	$\downarrow(\hat{L}U_1 \multimap \hat{R}U_2)$	$\downarrow\hat{R}U_1 \multimap \uparrow\hat{L}U_2$
a^-	$\downarrow a^-$	a^-
a^+	a^+	$\uparrow a^+$

$$\hat{R}U = \begin{cases} a^- & \text{if } U = a^- \\ \uparrow RU & \text{otherwise.} \end{cases}$$

$$\hat{L}U = \begin{cases} a^+ & \text{if } U = a^+ \\ \downarrow LU & \text{otherwise.} \end{cases}$$

Fig. 2. Linear Logic Polarization

These have the property that

Lemma 3. *Focused linear logic proofs of* $LU_1, LU_2, \cdots LU_n \vdash RU$ *are in bijective correspondence with unfocused proofs of* $U_1, U_2, \cdots U_n \vdash U$

Proof. By induction on the respective derivations. ■

Since we have used a translation that translates propositions differently depending on whether they essentially appear on the left or right, cut elimination and identity no longer come entirely for free. Nonetheless they are not difficult to show; they follow from the following results concerning shift operators.

Lemma 4. $(\uparrow\downarrow\uparrow P@p) \dashv\vdash (\uparrow P@p)$ *and* $(\downarrow\uparrow\downarrow N@k) \dashv\vdash (\downarrow N@k)$.

Lemma 5. $(\uparrow RU@p) \dashv\vdash (LU@p)$ *and* $(RU@k) \dashv\vdash (\downarrow LU@k)$.

The first of these two results is not unlike Brouwer's theorem $\neg\neg\neg A \dashv\vdash \neg A$. The second follows from applying it inductively at every connective, and shows that the translations L and R only differ up to the insertion of shift operators. As a consequence we recover identity expansion and cut admissibility hold for the object language. For example,

Lemma 6. $U \vdash U$

Proof. This sequent polarizes to $LU \vdash RU$, which under translation corresponds to $(LU@\alpha), (RU@\phi) \vdash \phi \triangleleft \alpha$. But this is provable iff $(LU@\alpha) \vdash \forall\phi.(RU@\phi) \Rightarrow \phi \triangleleft \alpha = (\uparrow RU@\alpha)$ is by inversion, which we know from Lemma 5. ∎

The cut admissibility proof works similarly: the import of Lemma 5 is that, while shifts were inserted in slightly different places to exactly match unfocused proof search, they have no effect on provability.

A.4 Focused Linear Logic Proofs

$$\frac{}{\Delta, a^+ \vdash [a^+]}\ a^+\mathsf{R} \qquad \frac{\Delta, a^+; \Omega \vdash N}{\Delta; \Omega, a^+ \vdash N}\ a^+\mathsf{L}$$

$$\frac{}{\Delta; [a^-] \vdash a^-}\ a^-\mathsf{L}$$

$$\frac{\Delta \vdash N}{\Delta \vdash [\downarrow N]}\ \downarrow\mathsf{R} \qquad \frac{\Delta, N; \Omega \vdash N'}{\Delta; \Omega, \downarrow N \vdash ; N'}\ \downarrow\mathsf{L}$$

$$\frac{\Delta \vdash P}{\Delta \vdash \uparrow P}\ \uparrow\mathsf{R} \qquad \frac{\Delta, P; \Omega \vdash P'}{\Delta; \Omega, \uparrow P \vdash ; P'}\ \uparrow\mathsf{L}$$

$$\frac{}{\Delta \vdash ; \top}\ \top R$$

$$\frac{\Delta; P \vdash ; N}{\Delta \vdash ; P \multimap N}\ \multimap R \qquad \frac{\Delta \vdash [P] \quad \Delta; [N] \vdash \bar{P}}{\Delta; [P \multimap N] \vdash \bar{P}}\ \multimap L$$

$$\frac{\Delta \vdash ; N_1 \quad \Delta \vdash ; N_2}{\Delta \vdash ; N_1 \,\&\, N_2}\ \& R \qquad \frac{\Delta; [N_i] \vdash \bar{P}}{\Delta; [N_1 \,\&\, N_2] \vdash \bar{P}}\ \& L_i$$

$$\frac{\Delta \vdash [P_i]}{\Delta \vdash [P_1 \oplus P_2]}\ \oplus\mathsf{R}_i \qquad \frac{\Delta; \Omega, P_1 \vdash N \qquad \Delta; \Omega, P_2 \vdash N}{\Delta; \Omega, P_1 \oplus P_2 \vdash N}\ \oplus\mathsf{L}$$

$$\frac{}{\Delta; \Omega, 0 \vdash N}\ 0\mathsf{L}$$

$$\frac{\Delta_1 \vdash [P_1] \qquad \Delta_2 \vdash [P_2]}{\Delta_1, \Delta_2 \vdash [P_1 \otimes P_2]}\ \otimes\mathsf{R} \qquad \frac{\Delta; \Omega, P_1, P_2 \vdash N}{\Delta; \Omega, P_1 \otimes P_2 \vdash N}\ \otimes\mathsf{L}$$

$$\frac{\Delta; N \vdash \bar{P}}{\Delta, N \vdash \bar{P}}\ \mathsf{foc} \qquad \frac{\Delta \vdash N}{\Delta; \cdot \vdash N}\ \mathsf{blur}$$