

Constructive reflection principles for regular theories

Henrik Forssell Peter LeFanu Lumsdaine

April 19, 2016

Abstract

Classically, any structure for a signature Σ may be completed to a model of a desired regular theory \mathbb{T} by means of the *chase construction* or *small object argument*. Moreover, this exhibits $\text{Mod}(\mathbb{T})$ as weakly reflective in $\text{Str}(\Sigma)$.

We investigate this in the constructive setting. The basic construction is unproblematic; however, it is no longer a weak reflection. Indeed, we show that various reflection principles for models of regular theories are equivalent to choice principles in the ambient set theory. However, the embedding of a structure into its chase-completion still satisfies a *conservativity* property, which suffices for applications such as the completeness of regular logic with respect to Tarski (i.e. set) models.

Unlike most constructive developments of predicate logic, we do not assume that equality between symbols in the signature is decidable. While in this setting, we also give a version of one classical lemma which is trivial over discrete signatures but more interesting here: the abstraction of constants in a proof to variables.

1 Introduction

Most developments of first-order logic are given in a classical meta-theory. Even presentations in constructive settings usually assume that the signature is *discrete*: that is, that decidability (excluded middle) holds for equality of the basic function and predicate symbols. This precludes various classical constructions, such as the diagram of a structure¹.

In many cases, this restriction can be dropped with little consequence. Many standard results and constructions, especially as developed in the categorical tradition, go through constructively for arbitrary signatures with no significant modification.

Some, however, do not adapt so straightforwardly. Here we set down constructive versions of two such results, for arbitrary signatures: the “chase” construction

¹Or, similarly in categorical logic, the internal language of a category

for producing models of regular (or “positive-primitive”) theories; and the abstraction of constants in a proof to variables.

Most substantially, we investigate the *chase construction* for regular theories.

This is one of a family of similar constructions which have been invented independently in several fields. We draw the name *chase* from database theory [AHV95, §8.4]; categorically, it is a form of the *small object argument* [AHRT02].

In each case, the idea is to construct a model of a regular theory², starting from a given structure, by iteratively adjoining elements to witness all existential axioms, and updating the basic predicates as necessary. Classically, this provides a weak reflection from arbitrary structures into models of the theory. That is, it provides for each structure A a model M and homomorphism $\eta : A \rightarrow M$, such that any homomorphism $f : A \rightarrow N$ from A to a model extends along η to a homomorphism $\tilde{f} : M \rightarrow N$.

In Section 4 we give a functorial chase construction for regular theories, and show that this provides a \mathbb{T} -*conservative* map from any structure into a model of the theory. This suffices for applications including the regular and geometric completeness of regular logic (with respect to Tarski models).

However, unless choice holds, this cannot provide a weak reflection as it does classically. Indeed, in Section 5, we show (using variations of the chase) that this and related reflection principles for regular theories are equivalent to choice principles in the ambient set theory.

Chase constructions can be regarded and used syntactically as a kind of proof calculus (see e.g. [CLR01]): presentation formulas take the place of structures, and results of the process correspond to provable consequences. This exploits the interchangeability of variables and constants/elements, something classically too simple to usually warrant more than a throwaway remark: any proof mentions only finitely many constants of the language, so they (or any subset) may be replaced by free variables (“abstracted away”) throughout.

Over non-discrete signatures, this is no longer trivial, since the set of constants occurring in the proof may not be discrete, nor hence constructively finite. With a little more effort, however, a version of the result still holds, by looking at which *occurrences* of the constants are required to be equal by the proof, and using this to perform the abstraction. This is the main content of Section 3, accompanied by a sample application over diagrams of models.

Having recalled some background in Section 2, we treat the abstraction of constants first in Section 3, as a warm-up to the main results on the chase construction and reflection principles in Sections 4 and 5. The two topics are each self-contained, however, so the reader who wishes to skip Section 3 and cut to the chase may safely do so.

²Or, in categorical terms, to obtain an injectivity condition.

Acknowledgements

For helpful discussions and correspondence on the material of this paper, we would like to thank (in approximately chronological order) Christian Espíndola, Erik Palmgren, Håkon Robbestad Gylterud, Jiří Rosický, Michael Rathjen, and Marek Zawadowski.

2 Preliminaries

As mentioned in the introduction, constructive treatments of syntax (e.g. [TvD88, §1.1], [Fit69, §4.1], [Vel76]) typically assume that the sets of basic function and operation symbols are discrete, and often moreover enumerable, though these assumptions are rarely used. This is at odds with classical and categorical model theory, where it is standard to consider signatures consisting of arbitrary sets; and this extra generality is needed for instance in taking the diagram of a structure, where one adds the elements of the structure to the language as constants.

In fact, especially in categorical treatments ([Joh02, D1.1], [Dia08, §3.1]), many proofs and constructions over arbitrary signatures are already fully constructive, and this is generally well-known in the field. However, we are unaware of any source that lays out the fundamentals of first-order model theory over arbitrary signatures in an explicitly constructive setting. We therefore recall basic background results in slightly more detail than usual, to put on record the fact that they and their standard proofs are constructively unproblematic.

Precisely, we work throughout in the constructive set theory IZF [TvD88, §11.8]. Sections 4 and 3 may comfortably also be read as working over IHOL_∞ (the logic of an elementary topos with a natural numbers object [LS86]), or over extensional type theory with an impredicative universe of propositions [Hof95]. Section 5, however, deals essentially with issues of unbounded quantification, which cannot be so straightforwardly translated to those settings.

Syntax

Throughout, we say a set X is *discrete* if it has decidable equality: for all $x, y \in X$, either $x = y$ or $x \neq y$.³ By *finite*, we mean *cardinal-finite*, i.e. bijectable with some set of the form $\{1, \dots, n\}$. A set is *Kuratowski-finite* or *K-finite* if it admits a surjection from some cardinal-finite set. A set is *finite* just if it is K-finite and discrete.

A *family* means by default a set-indexed family.

A (finitary) signature Σ consists of sets $\Sigma_{\text{Fun}}, \Sigma_{\text{Rel}}$ of function and relation symbols, with a natural-number arities function $n_{(-)} : \Sigma_{\text{Fun}} + \Sigma_{\text{Rel}} \rightarrow \mathbb{N}$. (We consider only single-sorted signatures.)

Syntax over such a signature is now constructed as usual (specifically, we follow [Joh02, D1.1.3]), using a fixed set V of variables, assumed discrete and constructively

³This is distinct from the homotopical sense of “discrete” in Homotopy Type Theory.

infinite⁴. We view the syntax as stratified into sets $\text{Tm}_\Sigma(I)$, $\text{Form}_\Sigma(I)$, for $I \subseteq_{\text{fin}} V$, of “terms (resp. formulas) in context I ”, i.e. with $FV(t), FV(\varphi) \subseteq I$.

As in [Joh02, D1.1.3], we consider fragments of logic including all atomic formulas but varying selections of connectives. *Horn* logic has just \wedge and \top ⁵; *regular* logic has these, plus \exists ; *geometric* logic has additionally has arbitrarily-indexed disjunctions, within finite contexts; and *first-order* logic has all usual (finitary) connectives and quantifiers.

While our contexts are formally finite subsets of V , we will often write them as lists $\mathbf{x} = x_1, \dots, x_n$ of distinct variables, representing the set $\{x_1, \dots, x_n\}$. Similarly, our fundamental form of (simultaneous, capture-free) substitution is formally $\varphi[f]$, where $\varphi \in \text{Form}(I)$ and $f : I \rightarrow \text{Tm}(J)$; but for a finite context \mathbf{x} , we will often write this as $\varphi[\mathbf{s}/\mathbf{x}]$, where $s_i = f(x_i)$. In the special case of a subset inclusion function $f : I' \subseteq I$, for $\varphi \in \text{Form}(I')$, we will write the weakening $\varphi[f] \in \text{Form}(I)$ just as φ .

Finally, we will often display variables explicitly, introducing a formula as e.g. $\phi(\mathbf{x})$ to indicate $\varphi \in \text{Form}(\mathbf{x})$, and having done so, writing $\varphi(\mathbf{s})$ for the substitution $\varphi[\mathbf{s}/\mathbf{x}]$.

A *sequent* consists of a context and a pair of formulas in that context, written as the formal expression $\varphi \vdash_I \psi$. A sequent is *regular*, *geometric*, etc. if its formulas lie in that fragment.

A (*regular, geometric, etc.*) *theory* \mathbb{T} is a set of sequents of the specified fragment, the *axioms* of \mathbb{T} . We write $\varphi \vdash_I^{\mathbb{T}} \psi$ to indicate that the sequent $\varphi \vdash_I \psi$ is derivable from axioms of \mathbb{T} , in the fragment of logic under consideration.

A *normal regular sequent* is one of the form $\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$, with φ, ψ Horn formulas, and such that $\psi(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^0 \varphi(\mathbf{x})$ is derivable. A *normal regular theory* is a family of normal regular sequents. By [Joh02, D1.3.10], every regular sequent or theory can be canonically transformed to an equivalent normal one, which we call its *normalisation*.

Semantics

Structures for a signature, and the subsequent interpretation of first-order logic therein, are defined as usual. We write $\text{Str}(\Sigma)$ for the (complete and co-complete) category of Σ -structures, and $\text{Mod}(\mathbb{T})$ for the full subcategory of \mathbb{T} -models.

A map of signatures $f : \Sigma \rightarrow \Sigma'$ (i.e. functions $f_{\text{Fun}} : \Sigma_{\text{Fun}} \rightarrow \Sigma'_{\text{Fun}}$, $f_{\text{Rel}} : \Sigma_{\text{Rel}} \rightarrow \Sigma'_{\text{Rel}}$, preserving arities) induces on the one hand a translation of syntax $(-)^f : \text{Form}_\Sigma \rightarrow \text{Form}_{\Sigma'}$, and on the other hand a reduct functor $f^* : \text{Str}(\Sigma') \rightarrow \text{Str}(\Sigma)$. These satisfy $\llbracket \mathbf{x} \mid \varphi \rrbracket^{f^*A} = \llbracket \mathbf{x} \mid \varphi^f \rrbracket^A$; so for any theory \mathbb{T} , f^* restricts to a functor $f^* : \text{Mod}(\mathbb{T}^f) \rightarrow \text{Mod}(\mathbb{T})$.

For any Horn formula $\varphi(\mathbf{x})$ (i.e. conjunction of atomic formulas), there is a structure $\llbracket \mathbf{x} \mid \varphi \rrbracket$ which *represents* the interpretation of $\varphi(\mathbf{x})$: there is a natural isomorphism $\llbracket \mathbf{x} \mid \varphi \rrbracket^A \cong \text{Hom}_{\text{Str}(\Sigma)}(\llbracket \mathbf{x} \mid \varphi \rrbracket, A)$. In case the signature is *relational*

⁴I.e. with an operation giving, for any finite list from V , some element not in the list.

⁵Classical Horn clauses correspond to Horn *sequents* in this sense.

(i.e. $\Sigma_{\text{Fun}} = \emptyset$), the domain $|\langle \mathbf{x} \mid \varphi \rangle|$ is finite. Explicitly, the domain of $\langle \mathbf{x} \mid \varphi \rangle$ is the quotient of \mathbf{x} by the (decidable) equivalence relation generated by the equality statements of φ ; that is, such that $x_i \sim x_j$ iff $\varphi \vdash_{\mathbf{x}}^{\emptyset} x_i = x_j$. For $R \in \Sigma$,

$$\llbracket R \rrbracket := \{ \langle [x_{i_0}], \dots, [x_{i_{n_{R-1}}}] \rangle \in (\mathbf{x}/\sim)^{n_R} \mid \varphi \vdash_{\mathbf{x}}^{\emptyset} R(x_{i_0}, \dots, x_{i_{n_{R-1}}}) \}.$$

For any regular formula $\theta(\mathbf{x})$, $\langle \mathbf{x} \mid \varphi \rangle \models \theta([x_0], \dots, [x_{n-1}])$ just if $\varphi(\mathbf{x}) \vdash_{\mathbf{x}}^{\emptyset} \theta(\mathbf{x})$.

A structure validates a regular normal sequent $\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$ just if it is *injective* with respect to the canonical map $\langle \mathbf{x} \mid \varphi \rangle \longrightarrow \langle \mathbf{x}, \mathbf{y} \mid \psi \rangle$ (cf. [AR94, Ex. 5.e]).

3 Replacing constants by variables

As a first warm-up with non-discrete signatures, we give a version of the classical technique of abstracting away constants in a derivation to variables, together with a sample application.

Before giving the general construction, we take an example, both to convince the reader that the statement is not quite trivial, and to illustrate the procedure used in the proof.

Suppose we have a derivation, from some theory \mathbb{T} , of $\top \vdash R(c_1, c_2)$. We want to abstract this to a derivation not mentioning c_1, c_2 , but with some free variables y instead, from which the original derivation can be recovered by substitution.

One's knee-jerk reaction might be to replace c_1 by y_1 throughout, and c_2 by y_2 . Of course, this is wrong: it works only if $c_1 \neq c_2$. Classically, one can salvage this approach by working by cases. Either $c_1 = c_2$ or $c_1 \neq c_2$, and in each case one gets a bijection from $\{c_1, c_2\}$ to some finite set of variables, which one can abstract along.

Constructively, however, we cannot in general make this case distinction, and hence cannot find such a bijection. To find an abstraction, we must look not just at the conclusion but at the entire derivation.

For instance, if the derivation obtains $\top \vdash R(c_1, c_2)$ from an axiom $\top \vdash_{x_1, x_2} R(x_1, x_2)$, then we abstract it to a derivation of $\top \vdash_{y_1, y_2} R(y_1, y_2)$. In case $c_1 = c_2$, our abstraction is more general than the original derivation; but in any case it is general enough that substituting c_i for y_i recovers the original.

If instead $\top \vdash R(c_1, c_2)$ is obtained from an axiom $\top \vdash_{x_1} R(x_1, x_1)$, then this shows that in fact c_1 and c_2 must be equal. So in this case we abstract to a derivation of $\top \vdash R(y_1, y_1)$, and know that substituting c_1 for y_1 gives back the original.

The key point is that while “ $c_1 = c_2$?” may not be decidable, the weaker question “does this derivation *require* $c_1 = c_2$?” always is, essentially since our syntax and deduction system is finitary. So we replace all *occurrences* of constants by variables, distinct as far as possible, and equal just when the derivation requires them to be.

Lemma 3.1. *Let Σ be a signature, C a decidable subset of its constants, and \mathbb{T} a theory over $\Sigma \setminus C$. Suppose $\varphi(\mathbf{x}), \psi(\mathbf{x})$ are formulas over Σ such that $\varphi(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \psi(\mathbf{x})$.*

Then there exist formulas $\bar{\varphi}(\mathbf{x}, \mathbf{y})$ and $\bar{\psi}(\mathbf{x}, \mathbf{y})$ over $\Sigma \setminus C$, and a function $f : \mathbf{y} \longrightarrow C$, such that $\bar{\varphi}[f] = \varphi$, $\bar{\psi}[f] = \psi$, and $\bar{\varphi}(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^{\mathbb{T}} \bar{\psi}(\mathbf{x}, \mathbf{y})$.

Moreover, the derivation of $\bar{\varphi}(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^{\mathbb{T}} \bar{\psi}(\mathbf{x}, \mathbf{y})$ may be taken to use the same logical rules and axiom schemes as the original derivation of $\varphi(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \psi(\mathbf{x})$; in particular, to lie in the same fragment of logic.

Proof. Recall that we work formally with the rules set out in [Joh02, D1.3.1], i.e. the standard intuitionistic sequent rules, presented in terms of sequents with a single antecedent.

We work directly by induction over the form of the derivation of φ .

(Axioms of \mathbb{T} .) If the derivation consists of just an axiom $\varphi \vdash_{\mathbf{x}} \psi$ of \mathbb{T} , then since \mathbb{T} does not mention C , we are done by taking $\mathbf{y} = \emptyset$, $\bar{\varphi} = \varphi$, $\bar{\psi} = \psi$.

(Structural and logical axioms.) If the derivation is an \wedge -elimination axiom

$$\frac{}{\varphi_0 \wedge \varphi_1 \vdash_{\mathbf{x}} \varphi_0}$$

we take \mathbf{y} to be fresh variables (i.e. distinct from \mathbf{x}) corresponding to the (finite) set of occurrences of constants from C in either φ_0 or φ_1 , and $f : \mathbf{y} \longrightarrow C$ the function sending each variable to the constant appearing at the corresponding occurrence. Then taking $\bar{\varphi}_i$ to be φ_i with each such occurrence replaced by the corresponding variable, and $\bar{\varphi} = \bar{\varphi}_0 \wedge \bar{\varphi}_1$, $\bar{\psi} = \bar{\varphi}_0$, we are done.

Other axioms are entirely analogous.

(Substitution rule.) If the derivation concludes with a substitution

$$\frac{\varphi' \vdash_{\mathbf{x}'} \psi'}{\varphi'[\mathbf{s}/\mathbf{x}'] \vdash_{\mathbf{x}} \psi'[\mathbf{s}/\mathbf{x}']}$$

where $\mathbf{s} : \mathbf{x}' \longrightarrow \text{Tm}(\mathbf{x})$, then by induction we have \mathbf{y}' , $f' : \mathbf{y}' \longrightarrow C$, and $\bar{\varphi}' \vdash_{\mathbf{x}', \mathbf{y}'} \bar{\psi}'$ as in the original statement.

Choose fresh variables \mathbf{y}'' corresponding to occurrences of constants from C in terms in \mathbf{s} , $f'' : \mathbf{y}'' \longrightarrow C$ sending each variable to the corresponding constant, and $\bar{\mathbf{s}} : \mathbf{x}' \longrightarrow \text{Tm}(\mathbf{x}, \mathbf{y}'')$ to be \mathbf{s} with occurrences replaced by the corresponding variables.

Now taking $\mathbf{y} = (\mathbf{y}', \mathbf{y}'')$, $f = f' \cup f''$, $\bar{\varphi} = \bar{\varphi}'[\bar{\mathbf{s}}/\mathbf{x}']$, $\bar{\psi} = \bar{\psi}'[\bar{\mathbf{s}}/\mathbf{x}']$, we are done, concluding by the substitution

$$\frac{\bar{\varphi}' \vdash_{\mathbf{x}', \mathbf{y}'} \bar{\psi}'}{\bar{\varphi}'[\bar{\mathbf{s}}/\mathbf{x}'] \vdash_{\mathbf{x}, \mathbf{y}', \mathbf{y}''} \bar{\psi}'[\bar{\mathbf{s}}/\mathbf{x}']}$$

(Double rules.) The double rules for \exists , \forall , and \Rightarrow are straightforward, needing no modification of the \mathbf{y} , f provided by the inductive hypothesis.

(Multi-premise rules.) Suppose the derivation concludes with a cut

$$\frac{\varphi \vdash_{\mathbf{x}} \chi \quad \chi \vdash_{\mathbf{x}} \psi}{\varphi \vdash_{\mathbf{x}} \psi}$$

By induction we have abstractions \mathbf{y}' , f' , and $\bar{\varphi}' \vdash_{\mathbf{x}, \mathbf{y}'}^{\mathbb{T}} \bar{\chi}'$ of the first subderivation, and \mathbf{y}'' , f'' , and $\bar{\chi}'' \vdash_{\mathbf{x}, \mathbf{y}''}^{\mathbb{T}} \bar{\psi}''$ of the second. Without loss of generality, we may assume \mathbf{y}' , \mathbf{y}'' disjoint.

Since $\bar{\chi}'[f'] = \bar{\chi}''[f''] = \chi$, occurrences of variables from \mathbf{y}' in $\bar{\chi}'$ correspond to occurrences of variables from \mathbf{y}'' in $\bar{\chi}''$. Let \sim be the equivalence relation on $\mathbf{y}', \mathbf{y}''$ generated by setting $y'_i \sim y''_j$ whenever some occurrence of y'_i corresponds to some occurrence of y''_j .

As a finitely generated equivalence relation on a finite set, \sim is decidable, so its quotient is finite, and can be represented by some fresh variables \mathbf{y} , with quotient map $q : \mathbf{y}', \mathbf{y}'' \rightarrow \mathbf{y}$. Now $f' \cup f''$ factors uniquely through q as $f : \mathbf{y} \rightarrow C$, and $\chi'[q] = \chi''[q]$, so we are done, concluding

$$\frac{\frac{\bar{\varphi}' \vdash_{\mathbf{x}, \mathbf{y}'} \bar{\chi}'}{\bar{\varphi}'[q] \vdash_{\mathbf{x}, \mathbf{y}} \bar{\chi}'[q]} \quad \frac{\bar{\chi}'' \vdash_{\mathbf{x}, \mathbf{y}''} \bar{\psi}'}{\bar{\chi}''[q] \vdash_{\mathbf{x}, \mathbf{y}} \bar{\psi}'[q]}}{\bar{\varphi}'[q] \vdash_{\mathbf{x}, \mathbf{y}} \bar{\psi}''[q]}.$$

The \wedge -introduction and \vee -elimination rules are analogous. \square

We remark that this proof adapts directly to other standard forms of sequent calculus and natural deduction (including versions with proof-terms), and indeed to other finitary extensions of predicate logic (e.g. by modal operators and corresponding rules).

A typical application of Lemma 3.1 is in analysing provability over diagrams of models, where having added constants to the signature, one may wish to re-abstract them.

Definition 3.2. Let A be a Σ -structure. Then $\Sigma + |A|$ is the signature obtained by adding constants to Σ for all elements of A (i.e. with $(\Sigma + |A|)_{\text{Fun}} = \Sigma_{\text{Fun}} + |A|$), and \mathbb{D}_A , the *diagram* of A , is the theory over $\Sigma + |A|$ consisting of the sequents $\top \vdash \varphi(\mathbf{a})$, for each atomic predicate instance $\varphi(\mathbf{a})$ that holds in A .

Proposition 3.3. For any regular formula $\varphi(\mathbf{x})$ and $\mathbf{a} \in A^{\mathbf{x}}$, $A \models \varphi(\mathbf{a})$ if and only if $\top \vdash^{\mathbb{D}_A} \varphi(\mathbf{a})$.

Proof. Immediate once $\varphi(\mathbf{x})$ is replaced by some equivalent $\exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$, with ψ Horn. \square

Lemma 3.4. Let A be a Σ -structure, \mathbb{T} a theory over Σ . Suppose $\psi(\mathbf{x}), \varphi(\mathbf{x}, \mathbf{y})$ are (arbitrary first-order) formulas over Σ , and $\mathbf{a} \in A^{\mathbf{y}}$, such that $\varphi(\mathbf{x}, \mathbf{a}) \vdash_{\mathbf{x}}^{\mathbb{T} \cup \mathbb{D}_A} \psi(\mathbf{x})$.

Then there is some regular formula $\chi(\mathbf{y})$ over Σ , such that $\top \vdash^{\mathbb{D}_A} \chi(\mathbf{a})$ (equivalently, $A \models \chi(\mathbf{a})$), and $\chi(\mathbf{y}) \wedge \varphi(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^{\mathbb{T}} \psi(\mathbf{x})$.

Proof. Fix some derivation of $\varphi(\mathbf{x}, \mathbf{a}) \vdash_{\mathbf{x}}^{\mathbb{T} \cup \mathbb{D}_A} \psi(\mathbf{x})$. Let ξ be the conjunction of all sentences σ (over $\Sigma + |A|$) occurring in axioms $\top \vdash \sigma$ from \mathbb{D}_A in the derivation. Then $\xi \wedge \varphi(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^{\mathbb{T}} \psi(\mathbf{x})$ is derivable.

By Lemma 3.1, we can abstract this to a derivation of some sequent $\bar{\xi}(\mathbf{w}) \wedge \bar{\varphi}(\mathbf{x}, \mathbf{w}) \vdash_{\mathbf{x}, \mathbf{w}}^{\mathbb{T}} \psi(\mathbf{x})$ just over Σ , along with $f : \mathbf{w} \rightarrow A$ such that $\bar{\xi}(\mathbf{w})[f] = \xi$ and $\bar{\varphi}(\mathbf{x}, \mathbf{w})[f] = \varphi(\mathbf{x}, \mathbf{a})$. Without loss of generality, \mathbf{w} is disjoint from \mathbf{y} . Let $\rho(\mathbf{y}, \mathbf{w})$ be the conjunction of equalities $y_i = w_j$, for y_i and w_j occurring in corresponding places of φ and $\bar{\varphi}$. Then $\rho(\mathbf{y}, \mathbf{w}) \wedge \varphi(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}, \mathbf{w}}^{\emptyset} \bar{\varphi}(\mathbf{x}, \mathbf{w})$.

But now $\exists \mathbf{w}. (\rho(\mathbf{y}, \mathbf{w}) \wedge \bar{\xi}(\mathbf{w})) \wedge \varphi(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}}^{\mathbb{T}} \psi(\mathbf{x})$, and f witnesses that $\top \vdash^{\mathbb{D}A} \exists \mathbf{w}. (\rho(\mathbf{a}, \mathbf{w}) \wedge \bar{\xi}(\mathbf{w}))$; so we are done. \square

4 Conservativity and completeness results for regular theories

Let \mathbb{T} be a regular theory over a signature Σ . Classically, $\text{Mod}(\mathbb{T})$ is a weakly reflective subcategory of $\text{Str}(\Sigma)$ (see e.g. [AR94, Thm. 4.8, Ex. 5.e]). Indeed, it is functorially so: every Σ -structure has a natural weakly reflective embedding $A \rightarrow W_{\mathbb{T}}(A)$ into a model of \mathbb{T} . The standard proofs involve choice, and we will show in Section 5 that this is unavoidable.

In this section, we show (constructively) a slightly weaker property, which nonetheless suffices for many applications: every Σ -structure has a natural \mathbb{T} -conservative embedding $\eta_A : A \rightarrow \text{Ch}_{\mathbb{T}}(A)$ into a model of \mathbb{T} .

Definition 4.1. For any Σ, \mathbb{T} , a homomorphism of structures $f : A \rightarrow B$ is \mathbb{T} -conservative (w.r.t. regular formulas) if, for every regular formula $\varphi(\mathbf{x})$ and any elements $\mathbf{a} \in A^{l(\mathbf{x})}$ such that $B \models \varphi(f(\mathbf{a}))$, there is some regular formula $\psi(\mathbf{x})$ such that $A \models \psi(\mathbf{a})$ and $\psi \vdash_{\mathbf{x}}^{\mathbb{T}} \varphi$.

Proposition 4.2. *Identity morphisms are \mathbb{T} -conservative. Composites of \mathbb{T} -conservative morphisms are \mathbb{T} -conservative.*

Definition 4.3. Let \mathbb{T} be a regular theory over a signature Σ . A *chase functor* for \mathbb{T} is a functor $\text{Ch}_{\mathbb{T}} : \text{Str}(\Sigma) \rightarrow \text{Mod}(\mathbb{T})$ together with a natural transformation $\eta : 1_{\text{Str}(\Sigma)} \rightarrow I \cdot \text{Ch}_{\mathbb{T}}$ (where I is the inclusion $\text{Mod}(\mathbb{T}) \hookrightarrow \text{Str}(\Sigma)$), such that η_A is \mathbb{T} -conservative for every A .

We borrow the term “chase” from the forward chaining algorithm of database theory of that name, as in e.g. [AHV95, §8.4]. Our construction can be seen as an adaptation of that method, or of the categorical *small-object argument* as given in e.g. [AHRT02].

The basic step of the traditional chase construction is the “axiom-induced” extension of a structure. Let $\sigma = (\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y}))$ be a sequent, let M be a structure, and let $\mathbf{m} \in M$ such that $M \models \varphi[\mathbf{m}/\mathbf{x}]$. Then “the extension of M induced by σ and \mathbf{m} ” is the structure obtained by adding new elements \mathbf{b} corresponding to the variables \mathbf{y} , and updating the predicates just as required to make $\psi(\mathbf{a}, \mathbf{b})$ hold. The map from M to this extension is evidently weakly orthogonal to models of σ .

If M is enumerable, or more generally well-orderable, then (classically) one can iterate this construction over all possible arguments \mathbf{m} , eventually obtaining a

weak reflection $h : M \rightarrow \text{Ch}_\sigma(M)$ into a model of the sequent σ ; and similarly for any well-orderable set of sequents.

In the current setting, where theories and structures need not be finite, enumerable, or even discrete, this one-at-a-time approach is insufficient. Instead, at each step, we simultaneously adjoin new elements for all possible applications of axioms of \mathbb{T} to arguments in the structure (as in the small-object argument). We then iterate this step ω -many times, to obtain a model of \mathbb{T} .

The tricky part is proving conservativity. For this, we perform the construction directly just in the case where Σ is purely relational (i.e. has no constants or function symbols), and \mathbb{T} does not mention equality. This renders analysis of the basic extension step more tractable. We then obtain the general version from this restricted case, by means of elimination of equality and function symbols in the syntax.

Definition 4.4. Let Σ be a purely relational signature, and \mathbb{T} a regular normal equality-free theory over Σ .

For a Σ -structure A , the *one-step \mathbb{T} -extension* of A , denoted $\text{ex}_{\mathbb{T}}(A)$ or just $\text{ex}(A)$, is the Σ -structure defined as follows:

- $|\text{ex}(A)|$ is the disjoint union of $|A|$ with the set of all triples (τ, \mathbf{a}, j) , where $\tau = (\varphi \vdash_x \exists y. \psi)$ is an axiom of \mathbb{T} , $\mathbf{a} \in \llbracket \mathbf{x} \mid \varphi \rrbracket^A$, and $0 \leq j < l(y)$. Write $\iota_A : |A| \rightarrow |\text{ex}(A)|$ for the inclusion.
- For each predicate symbol R , $\llbracket R \rrbracket^{\text{ex}(A)}$ consists of $\iota_A(\llbracket R \rrbracket^A)$, together with for each axiom $\tau = (\varphi \vdash_x \exists y. \psi)$, each $\mathbf{a} \in \llbracket \mathbf{x} \mid \varphi \rrbracket^A$, and each occurrence of R as a conjunct $R(\mathbf{t}(\mathbf{x}, \mathbf{y}))$ in ψ , the tuple $\llbracket \mathbf{x}, \mathbf{y}, \mathbf{t} \rrbracket^{\text{ex}(A)}(\iota(\mathbf{a}), \mathbf{b})$, where \mathbf{b} is the *canonical witnessing tuple* $((\tau, \mathbf{a}, i) \mid 0 \leq i < l(\mathbf{y})) \in \text{ex}(A)^{\mathbf{x}, \mathbf{y}}$.

Moreover, this is evidently functorial in A ; and ι_A gives a natural homomorphism $A \rightarrow \text{ex}(A)$.

Categorically, $\text{ex}(A)$ is the pushout of A with a coproduct of copies of the structure inclusions $\langle \mathbf{x} \mid \varphi \rangle \rightarrow \langle \mathbf{x}, \mathbf{y} \mid \psi \rangle$ representing the axioms of \mathbb{T} (cf. $C(K)$ in [AHRTo2, II.4]). Equality-freeness of \mathbb{T} ensures that this pushout can be presented simply as a disjoint union, with no quotienting required.

Some terminology will be useful for working explicitly in $\text{ex}(A)$. Given a newly adjoined element (τ, \mathbf{a}, j) , call τ the *justification* for this element, \mathbf{a} its *arguments*, and j its *index*.

Note that whenever an instance of a basic predicate holds in $\text{ex}(A)$, either it already holds in A , or else all new elements occurring in it have the same justification τ and arguments \mathbf{a} , and the instance comes from the conclusion of τ applied to \mathbf{a} .

Definition 4.5. (Σ relational, \mathbb{T} regular normal equality-free.)

For a Σ -structure A , the *\mathbb{T} -chase of A* , written $\text{Ch}_{\mathbb{T}}(A)$, is the colimit of the sequence of structures

$$A \xrightarrow{\iota} \text{ex}(A) \xrightarrow{\iota} \text{ex}(\text{ex}(A)) \xrightarrow{\iota} \dots$$

Concretely, since each ι is a complemented inclusion, $\text{Ch}(A)$ may be taken to consist of pairs (i, x) , where $i \in \mathbb{N}$, and $x \in \text{ex}^i(A) \setminus \text{ex}^{i-1}(A)$ for $i > 0$, or $x \in A$ in case $i = 0$. Then $\text{Ch}(A) \models R((i_1, x_1), \dots, (i_r, x_r))$ just when there is some $j \geq i_1, \dots, i_r$ such that $\text{ex}^j(A) \models R(\iota^{j-i_1}(x_1), \dots, \iota^{j-i_r}(x_r))$.

This too is functorial in A , and the colimit inclusions $v_i : \text{ex}^i(A) \longrightarrow \text{Ch}(A)$ are natural. We write $\eta_A : A \longrightarrow \text{Ch}(A)$ for the 0th such inclusion v_0 .

Proposition 4.6. (Σ relational, \mathbb{T} regular normal equality-free.) $\text{Ch}(A)$ is a model of \mathbb{T} .

Proof. For each axiom $\tau = (\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y}))$ of \mathbb{T} , suppose $\mathbf{a} \in \text{Ch}(A)^{l(\mathbf{x})}$, and $\text{Ch}(A) \models \varphi(\mathbf{a})$. Then there is some n and some $\mathbf{b} \in \text{ex}^n(A)^{l(\mathbf{x})}$ such that $v_n(\mathbf{b}) = \mathbf{a}$, and $\text{ex}^n(A) \models \varphi(\mathbf{b})$. Now by construction, $\text{ex}^{n+1}(A) \models \psi(\iota(\mathbf{b}), (\tau, \mathbf{b}, i)_{1 \leq i \leq l(\mathbf{y})})$; so $\text{Ch}(A) \models \psi(\mathbf{a}, v_{n+1}(\tau, \mathbf{b}, i)_{1 \leq i \leq l(\mathbf{y})})$, and so $\text{Ch}(A)$ validates the conclusion of τ . \square

Recall Definition 4.1 of \mathbb{T} -conservativity for homomorphisms. We will show in 4.8 that ι_A is \mathbb{T} -conservative; first we establish a restricted special case.

Lemma 4.7. (Σ relational, \mathbb{T} regular normal equality-free.) Suppose:

- $\varphi(\mathbf{x}, \mathbf{y})$ is some non-empty conjunction of atomic formulas, each containing some variable from \mathbf{y} ;
- $\mathbf{a} \in A^{\mathbf{x}}$, $\mathbf{b} \in \text{ex}(A)^{\mathbf{y}}$;
- all elements of \mathbf{b} are in $\text{ex}(A) \setminus A$, and have the same justification and argument;
- $\text{ex}(A) \models \varphi(\iota(\mathbf{a}), \mathbf{b})$.

Then there is some regular formula $\psi(\mathbf{x})$ such that $A \models \psi(\mathbf{a})$ and $\psi(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \exists \mathbf{y}. \varphi(\mathbf{x}, \mathbf{y})$.

Proof. Take $\tau = (\chi \vdash_{\mathbf{z}} \exists \mathbf{w}. \zeta)$ to be the shared justification of the elements of \mathbf{b} , and $\mathbf{c} \in A^{\mathbf{z}}$ their shared argument. Write $\text{idx}(b)$ for the index of b .

Write \mathbf{d} for their canonical witness tuple $(\tau, \mathbf{c}, j)_{j \in \mathbf{w}} \in \text{ex}(A)^{\mathbf{w}}$. Define $f : \mathbf{y} \longrightarrow \mathbf{w}$ by $i \mapsto \text{idx}(b_i)$.

Consider any non-equality atomic conjunct $R(\mathbf{t}(\mathbf{x}, \mathbf{y}))$ of φ , where $\mathbf{t} \in \text{Tm}(\mathbf{x}, \mathbf{y})^{n_R}$. We know $R(\mathbf{t}(\mathbf{x}, \mathbf{y}))$ contains some variable from \mathbf{y} ; so its interpretation in $\text{ex}(A)$ must arise from τ applied to \mathbf{c} . That is, there is some conjunct of $\zeta(\mathbf{z}, \mathbf{w})$ of the form $R(\mathbf{s}(\mathbf{z}, \mathbf{w}))$, where $\mathbf{s} \in \text{Tm}(\mathbf{z}, \mathbf{w})^{n_R}$, such that $\llbracket \mathbf{s}(\mathbf{z}, \mathbf{w}) \rrbracket^{\text{ex}(A)}(\iota(\mathbf{c}), \mathbf{d}) = \llbracket \mathbf{t}(\mathbf{x}, \mathbf{y}) \rrbracket^{\text{ex}(A)}(\iota(\mathbf{a}), \mathbf{b})$.

Now certainly $\zeta(\mathbf{z}, \mathbf{w}) \wedge \bigwedge_{i \in n_R} t_i(\mathbf{z}, \mathbf{w}) = s_i(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}} R(\mathbf{t}(\mathbf{x}, \mathbf{y}))$.

For each $i \in n_R$, we have $\llbracket t_i(\mathbf{x}, \mathbf{y}) \rrbracket^{\text{ex}(A)}(\iota(\mathbf{a}), \mathbf{b}) = \llbracket s_i(\mathbf{z}, \mathbf{w}) \rrbracket^{\text{ex}(A)}(\iota(\mathbf{c}), \mathbf{d})$. Since each term is a variable, and the images of \mathbf{b}, \mathbf{d} are distinct from $\iota(\mathbf{a}), \iota(\mathbf{c})$, either $t_i(\mathbf{x}, \mathbf{y})$ is of the form x_j while $s_i(\mathbf{z}, \mathbf{w})$ is of the form z_k , or else $t_i(\mathbf{x}, \mathbf{y})$ is of the form y_j while $s_i(\mathbf{z}, \mathbf{w})$ is exactly $w_{f(j)}$.

So let $\epsilon(\mathbf{x}, \mathbf{z})$ be the conjunction, over all i such that $t_i(\mathbf{x}, \mathbf{y})$ is some x_j , of the equalities $t_i(\mathbf{x}) = s_i(\mathbf{z})$. Then $\zeta(\mathbf{z}, \mathbf{w}) \wedge \epsilon(\mathbf{x}, \mathbf{z}) \wedge \bigwedge_{1 \leq i \leq l(\mathbf{y})} y_i = w_{f(i)} \vdash_{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}} R(\mathbf{t}(\mathbf{x}, \mathbf{y}))$; and since ι is injective, $A \models \epsilon(\mathbf{a}, \mathbf{c})$.

Take $\rho(\mathbf{x}, \mathbf{z})$ to be the conjunction of these formulas $\epsilon(\mathbf{x}, \mathbf{z})$ over all non-equality atomic conjuncts of φ . Then

$$\rho(\mathbf{x}, \mathbf{z}) \wedge \zeta(\mathbf{z}, \mathbf{w}) \wedge \bigwedge_{1 \leq i \leq l(\mathbf{y})} y_i = w_{f(i)} \vdash_{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}} \varphi(\mathbf{x}, \mathbf{y}) :$$

the left-hand side implies all non-equality conjuncts of φ by the construction of ρ , and all equality conjuncts since if $b_i = b_j$ then $f(i) = f(j)$. Meanwhile, $\text{ex}(A) \models \rho(\iota(\mathbf{a}), \iota(\mathbf{c}))$, so since ι is an injection and ρ a conjunction of equalities, $A \models \rho(\mathbf{a}, \mathbf{c})$.

Finally, take the desired formula $\psi(\mathbf{x})$ to be $\exists \mathbf{z}. \rho(\mathbf{x}, \mathbf{z}) \wedge \chi(\mathbf{z})$. Then $A \models \psi(\mathbf{a})$, witnessed by \mathbf{c} , and moreover

$$\begin{aligned} \psi(\mathbf{x}) &\vdash_{\mathbf{x}}^{\mathbb{T}} \exists \mathbf{z}, \mathbf{w} \rho(\mathbf{x}, \mathbf{z}) \wedge \zeta(\mathbf{z}, \mathbf{w}) \\ &\vdash_{\mathbf{x}}^{\mathbb{T}} \exists \mathbf{z}, \mathbf{w}, \mathbf{y} \rho(\mathbf{x}, \mathbf{z}) \wedge \zeta(\mathbf{z}, \mathbf{w}) \wedge \bigwedge_i y_i = w_{f(i)} \\ &\vdash_{\mathbf{x}}^{\mathbb{T}} \exists \mathbf{z}, \mathbf{w}, \mathbf{y} \varphi(\mathbf{x}, \mathbf{y}) \\ &\vdash_{\mathbf{x}}^{\mathbb{T}} \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y}). \quad \square \end{aligned}$$

Proposition 4.8. (Σ relational, \mathbb{T} regular normal equality-free.) *The map $\iota : A \longrightarrow \text{ex}(A)$ is \mathbb{T} -conservative, for every Σ -structure A .*

Proof. Let $\varphi(\mathbf{x})$ be a regular formula, and $\mathbf{a} \in A^{l(\mathbf{x})}$ be such that $\text{ex}(A) \models \varphi(\iota(\mathbf{a}))$. We seek some $\psi(\mathbf{x})$ such that $A \models \psi(\mathbf{a})$ and $\psi(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \varphi(\mathbf{x})$.

As usual, φ is provably equivalent to some formula $\exists \mathbf{y}. \varphi'$, with $\varphi'(\mathbf{x}, \mathbf{y})$ a finite conjunction of atomic formulas. By adjoining instances of reflexivity if necessary, we may assume that each variable of \mathbf{y} occurs somewhere in φ' . Now take some $\mathbf{b} \in \text{ex}(A)^{l(\mathbf{y})}$ such that $\text{ex}(A) \models \varphi'(\mathbf{a}, \mathbf{b})$.

Since $\iota : A \longrightarrow \text{ex}(A)$ is a decidable injection, we can decide which values of \mathbf{b} are in its image, and correspondingly reorder \mathbf{y} into the form \mathbf{x}', \mathbf{y}' , such that the corresponding reordering of \mathbf{b} is of the form $\iota(\mathbf{a}'), \mathbf{b}'$, with $\mathbf{a}' \in A^{l(\mathbf{x}')}$, and \mathbf{b}' lying entirely in the complement of ι .

Now let \sim be the equivalence relation on \mathbf{y}' generated by setting $y'_i \sim y'_j$ whenever there is some atomic conjunct of φ' in which both y'_i and y'_j occur. This is decidable, so allows us to reorder \mathbf{y}' as $\mathbf{y}^1, \dots, \mathbf{y}^n$, where each \mathbf{y}^i is an equivalence class of \sim . Write $\mathbf{b}^1, \dots, \mathbf{b}^n$ for the corresponding reordering of \mathbf{b}' . By the definition of \sim and the note following Definition 4.4, for each i , all the elements in \mathbf{b}^i must have the same justification and argument.

By reordering quantifiers and conjuncts in $\exists \mathbf{y}'. \varphi'(\mathbf{x}, \mathbf{y}')$, we now have $\varphi(\mathbf{x})$ equivalent to some formula

$$\exists \mathbf{x}'. \exists \mathbf{y}^1. \dots \exists \mathbf{y}^n. \varphi_0(\mathbf{x}, \mathbf{x}') \wedge \varphi_1(\mathbf{x}, \mathbf{x}', \mathbf{y}^1) \wedge \dots \wedge \varphi_n(\mathbf{x}, \mathbf{x}', \mathbf{y}^n)$$

where each φ_i is a non-empty conjunction of atomic formulas, and for $i > 0$, each conjunct of φ_i containing some variable from \mathbf{y}^i .

This is in turn equivalent to

$$\exists \mathbf{x}' . \left(\varphi_0(\mathbf{x}, \mathbf{x}') \wedge (\exists \mathbf{y}^1 . \varphi_1(\mathbf{x}, \mathbf{x}', \mathbf{y}^1)) \wedge \cdots \wedge (\exists \mathbf{y}^n . \varphi_n(\mathbf{x}, \mathbf{x}', \mathbf{y}^n)) \right).$$

To unify the form of these conjuncts, take \mathbf{y}^0 to be the empty sequence of variables.

It is now sufficient to show that for each $0 \leq i \leq n$, there is some formula $\psi_i(\mathbf{x}, \mathbf{x}')$ such that $A \models \psi_i(\mathbf{a}, \mathbf{a}')$, and $\psi_i \vdash_{\mathbf{x}, \mathbf{x}'}^{\mathbb{T}} \exists \mathbf{y}^i . \varphi_i$. Given these, choosing some such ψ_i for each i and taking $\psi(\mathbf{x}) := \exists \mathbf{x}' . \bigwedge_i \psi_i(\mathbf{x}, \mathbf{x}')$ completes the proof.

For $i \neq 0$, this is exactly Lemma 4.7.

In the case $i = 0$, \mathbf{y}^0 is empty and so $\exists \mathbf{y}^0 . \varphi_0(\mathbf{x}, \mathbf{x}', \mathbf{y}^0)$ is just $\varphi_0(\mathbf{x}, \mathbf{x}')$. Here it is enough to show that for each atomic conjunct $\alpha(\mathbf{x}, \mathbf{x}')$ of φ_0 , there is some $\sigma(\mathbf{x}, \mathbf{x}')$ with $A \models \sigma(\mathbf{a}, \mathbf{a}')$ and $\sigma(\mathbf{x}, \mathbf{x}') \vdash_{\mathbf{x}, \mathbf{x}'}^{\mathbb{T}} \alpha(\mathbf{x}, \mathbf{x}')$. Given this, we are done by choosing some such σ for each α and taking ψ_0 to be their conjunction.

So, take some such $\alpha(\mathbf{x}, \mathbf{x}')$. By the explicit description of the structure on $\text{ex}(A)$, the fact that $\text{ex}(A) \models \alpha(\iota(\mathbf{a}), \iota(\mathbf{a}'))$ must arise either because $A \models \alpha(\mathbf{a}, \mathbf{a}')$, or else from some axiom τ , some arguments \mathbf{c} , and some atomic conjunct of the conclusion of τ applied to \mathbf{c} . In the former case, we are done just by taking σ to be α itself.

In the latter case, say τ is of the form $\chi \vdash_z \exists \mathbf{w} . \zeta$, and $R(\mathbf{s}(\mathbf{z}, \mathbf{w}))$ the atomic conjunct of ζ giving rise to the fact that $\text{ex}(A) \models \alpha(\iota(\mathbf{a}), \iota(\mathbf{a}'))$.

Now $\alpha(\mathbf{x}, \mathbf{x}')$ must itself be of the form $R(\mathbf{t}(\mathbf{x}, \mathbf{x}'))$, and with $\mathbf{t}(\iota(\mathbf{a}), \iota(\mathbf{a}')) = \mathbf{s}(\iota(\mathbf{c}), (\tau, \mathbf{c}, j)_{1 \leq j \leq l(\mathbf{w})})$, in $\text{ex}(A)$; so \mathbf{s} must be just of the form $\mathbf{s}(\mathbf{z})$, not mentioning \mathbf{w} , and since ι is injective, $\mathbf{t}(\mathbf{a}, \mathbf{a}') = \mathbf{s}(\mathbf{c})$ in A . So taking $\rho(\mathbf{x}, \mathbf{x}', \mathbf{z})$ to be the conjunction of equalities $\bigwedge_{1 \leq j \leq n(R)} t_j(\mathbf{x}, \mathbf{x}') = s_j(\mathbf{z})$, and taking $\sigma(\mathbf{x}, \mathbf{x}')$ to be $\exists \mathbf{z} . (\rho(\mathbf{x}, \mathbf{x}', \mathbf{z}) \wedge R(\mathbf{s}(\mathbf{z})))$, we are done. \square

Corollary 4.9. (Σ relational, \mathbb{T} regular normal equality-free.) *The embedding $\eta_A : A \longrightarrow \text{Ch}(A)$ is \mathbb{T} -conservative, for every Σ -structure A ; so $\text{Ch}_{\mathbb{T}}$ and η form a chase functor for \mathbb{T} .*

Proof. Since \mathbb{T} -conservative maps are closed under composition and identities, $\iota^n : A \longrightarrow \text{ex}^n(A)$ is \mathbb{T} -conservative, for all $n \geq 0$. But given $\varphi(\mathbf{x})$ and \mathbf{a} such that $\text{Ch}(A) \models \varphi(\eta(\mathbf{a}))$, by finitariness there is some n such that $\text{ex}^n(A) \models \iota^n(\mathbf{a})$, at which point we are done by conservativity of ι^n . \square

We can now ease our restrictions on \mathbb{T} and Σ , using first *elimination of equalities* to remove “ \mathbb{T} equality-free”, and then *elimination of function symbols* to remove “ Σ purely relational”. These techniques are standard, but formulations of them vary widely, so we recall carefully the specific versions that we require.

Definition 4.10 (Elimination of equalities; cf. [BM77, §2.11]). Let Σ be any signature; take Σ_E to be the extension of Σ by a new binary predicate symbol E , and let \mathbb{E}_{Σ} be the theory stating that E is an equivalence relation and all other predicate symbols respect E .

For a formula φ over Σ , let φ^E be the formula over Σ_E given by replacing equality in φ with E ; conversely, for a formula ψ over Σ_E , let ${}^E\psi$ be the formula

over Σ given by replacing E in ψ with equality. These translations extend to sequents and theories in the obvious way.

Proposition 4.11. *Both translations preserve provability, modulo \mathbb{E}_Σ . That is, if $\varphi_1 \vdash_{\mathbf{x}}^{\mathbb{T}} \varphi_2$ (over Σ), then $\varphi_1^E \vdash_{\mathbf{x}}^{\mathbb{T}^E \cup \mathbb{E}_\Sigma} \varphi_2^E$; and if $\psi_1 \vdash_{\mathbf{x}}^{\mathbb{T} \cup \mathbb{E}_\Sigma} \psi_2$ (over Σ_E), then ${}^E\psi_1 \vdash_{\mathbf{x}}^{{}^E\mathbb{T}} {}^E\psi_2$. Also, ${}^E(\varphi^E) = \varphi$, for any formula φ over Σ . \square*

Proposition 4.12. *The composite forgetful functor $u : \text{Mod}(\mathbb{E}_\Sigma) \longrightarrow \text{Str}(\Sigma_E) \longrightarrow \text{Str}(\Sigma)$ has a left adjoint $e : \text{Str}(\Sigma) \longrightarrow \text{Mod}(\mathbb{E}_\Sigma)$, which extends any Σ -structure to a Σ_E -structure by interpreting E as equality; and e has a further left adjoint $q : \text{Mod}(\mathbb{E}_\Sigma) \longrightarrow \text{Str}(\Sigma)$, which quotients a structure by E . Moreover, the adjunction $q \dashv e$ is a reflection: its counit $q(e(A)) \longrightarrow A$ is a natural isomorphism.*

Let $\varphi(\mathbf{x})$ be any formula over Σ , $A \in \text{Str}(\Sigma)$, and $\mathbf{a} \in A^{\mathbf{x}}$. Then $A \models \varphi(\mathbf{a})$ if and only if $e(A) \models \varphi^E(\mathbf{a})$.

Let $\psi(\mathbf{x})$ be any formula over Σ_E , $B \in \text{Str}(\Sigma_E)$, and $\mathbf{b} \in B^{\mathbf{x}}$. If $B \models \psi(\mathbf{b})$, then $q(B) \models {}^E\psi([\mathbf{b}])$.

Finally, for any theory \mathbb{T} over σ , the adjunction $q \dashv e$ restricts to an adjunction between the subcategories $\text{Mod}(\mathbb{T})$ and $\text{Mod}(\mathbb{T}^E \cup \mathbb{E}_\Sigma)$; and $q : \text{Mod}(\mathbb{E}_\Sigma) \longrightarrow \text{Str}(\Sigma)$ sends \mathbb{T}^E -conservative maps to \mathbb{T} -conservative maps. \square

Proposition 4.13. *Any regular theory \mathbb{T} (possibly involving equality) over a purely relational signature Σ has a chase functor.*

Proof. By Corollary 4.9, $\mathbb{T}^E \cup \mathbb{E}_\Sigma$ has a chase functor, which we write as Ch_E, η^E .

Consider these as restricted to the subcategory $\text{Mod}(\mathbb{E}_\Sigma)$. Then take $\text{Ch}_{\mathbb{T}}$ to be the composite $q \cdot \text{Ch}_E \cdot e : \text{Str}(\Sigma) \longrightarrow \text{Mod}(\mathbb{E}_\Sigma) \longrightarrow \text{Mod}(\mathbb{T}^E \cup \mathbb{E}_\Sigma) \longrightarrow \text{Mod}(\mathbb{T})$; and take $\eta_A^{\mathbb{T}}$ to be the composite of $q(\eta_{e(A)}^E) : q(e(A)) \longrightarrow q(\text{Ch}_{\mathbb{T}^E}(e(A)))$ with the natural isomorphism $A \cong q(e(A))$. \square

(Note that when \mathbb{T} is already equality-free, $\text{Ch}_{\mathbb{T}}$ provided by this lemma is naturally isomorphic to the previous version.)

Proposition 4.14 (Elimination of function symbols; cf. [BM77, §2.10]). *Let Σ be any signature. Let $\bar{\Sigma}$ be the signature obtained from Σ by replacing each n -ary function symbol by an $(n+1)$ -ary predicate symbol, and let $\mathbb{F}_{\bar{\Sigma}}$ the regular theory over $\bar{\Sigma}$ asserting that the new predicate symbols are functional. Then there is a translation $\bar{\cdot} : \text{Form}_\Sigma(-) \longrightarrow \text{Form}_{\bar{\Sigma}}(-)$, preserving the regular fragment of logic, and conservative modulo $\mathbb{F}_{\bar{\Sigma}}$, i.e. such that $\varphi_1, \dots, \varphi_n \vdash_{\mathbf{x}} \psi$ if and only if $\bar{\varphi}_1, \dots, \bar{\varphi}_n \vdash_{\mathbf{x}}^{\mathbb{F}_{\bar{\Sigma}}} \bar{\psi}$, for all suitable $\mathbf{x}, \varphi_i, \psi$. Moreover, this translation induces an equivalence of categories $\text{Str}(\Sigma) \simeq \text{Mod}(\mathbb{F}_{\bar{\Sigma}})$. \square*

Theorem 4.15. *Let Σ be an arbitrary signature, and \mathbb{T} a regular theory over Σ . Then there exists a chase functor $\text{Ch}_{\mathbb{T}} : \text{Str}(\Sigma) \longrightarrow \text{Mod}(\mathbb{T})$ for \mathbb{T} .*

Proof. Let $\bar{\Sigma}$ and the translation $\bar{\cdot}$ be as in Proposition 4.14. Take $\bar{\mathbb{T}}$ to be the normalisation of the regular theory

$$\mathbb{F}_{\bar{\Sigma}} \cup \{\bar{\varphi} \vdash_{\mathbf{x}} \bar{\psi} \mid (\varphi \vdash_{\mathbf{x}} \psi) \in \mathbb{T}\}.$$

The equivalence $\text{Str}(\Sigma) \simeq \text{Mod}(\mathbb{F}_\Sigma)$ restricts to an equivalence $\text{Mod}(\mathbb{T}) \simeq \text{Mod}(\bar{\mathbb{T}})$, so the chase functor $\text{Ch}_{\bar{\mathbb{T}}}$ provided by Proposition 4.13 restricts to $\text{Mod}(\mathbb{F}_\Sigma)$ and transfers along these equivalences, yielding a chase functor for \mathbb{T} . \square

(Note that in case Σ was already purely relational and \mathbb{T} normal, we once again have not changed $\text{Ch}_{\mathbb{T}}$, since in this case $\bar{\Sigma} = \Sigma$, $\bar{\mathbb{T}} = \mathbb{T}$, and the translations and equivalences involved are identities.)

We end this section by drawing, as corollaries, the regular and geometric completeness of regular theories.

Corollary 4.16. *Let \mathbb{T} be a regular theory over a signature Σ , and σ a regular sequent valid in all models of \mathbb{T} . Then $\mathbb{T} \vdash \sigma$.*

Proof. We assume without loss of generality that Σ is relational (by Proposition 4.14), and that σ is of the form $\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \psi(\mathbf{x})$, with φ Horn.

Write S for the representing structure $\langle\langle \mathbf{x} \mid \varphi \rangle\rangle$, and $[\mathbf{x}]$ for its canonical tuple $[x_0], \dots, [x_{n-1}] \in S^{\mathbf{x}}$. Now consider $\eta_S : S \rightarrow \text{Ch}_{\mathbb{T}}(S)$. Certainly $S \models \varphi([\mathbf{x}])$, so since φ is Horn, $\text{Ch}_{\mathbb{T}}(S) \models \varphi(\eta_S([\mathbf{x}]))$. By hypothesis, $\text{Ch}_{\mathbb{T}}(S) \models \sigma$, so $\text{Ch}_{\mathbb{T}}(S) \models \psi(\eta_S([\mathbf{x}]))$. So by \mathbb{T} -conservativity of η_S , there is some regular $\theta(\mathbf{x})$ such that $S \models \theta([\mathbf{x}])$ and $\theta(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \psi(\mathbf{x})$.

But by the characterisation of validity in S , this implies that $\varphi(\mathbf{x}) \vdash_{\mathbf{x}}^{\emptyset} \theta([\mathbf{x}])$, and hence $\varphi(\mathbf{x}) \vdash_{\mathbf{x}}^{\mathbb{T}} \psi(\mathbf{x})$. \square

In fact, in this proof, the consequent of σ could have been an arbitrary disjunction $\bigvee_{i \in I} \psi_i$ of regular formulas: $\text{Ch}_{\mathbb{T}}(S) \models \bigvee_{i \in I} \psi_i(\eta_S([\mathbf{x}]))$ means that there exists some $i \in I$ such that $\text{Ch}_{\mathbb{T}}(S) \models \psi_i(\eta_S([\mathbf{x}]))$, whence we proceed as before. We thus obtain a both constructive and (Tarski-)semantical proof of the disjunction property (cf. [Joh02, D3.3.11]) and geometric completeness for regular theories:

Scholium 4.17. *Let \mathbb{T} be a regular theory and σ a sequent of the form $\varphi \vdash_{\mathbf{x}} \bigvee_I \psi_i$, where φ and ψ_i are regular formulas. If $\mathbb{T} \vdash \sigma$, then there is some $i \in I$ such that $\varphi \vdash_{\mathbf{x}}^{\mathbb{T}} \psi_i$. \square*

Scholium 4.18. *Let \mathbb{T} be a regular theory over a signature Σ . Then any geometric sequent valid in all models of \mathbb{T} is provable from \mathbb{T} .*

Proof. For sequents of the form $\varphi \vdash_{\mathbf{x}} \bigvee_I \exists \mathbf{y}_i. \psi_i$, with φ and all ψ_i Horn, this is a minor adjustment of the proof of Corollary 4.16, as sketched above.

By [Joh02, D1.3.8], any geometric sequent is provably equivalent to one of the form $\bigvee_J \exists \mathbf{z}_j. \varphi_j \vdash_{\mathbf{x}} \bigvee_I \exists \mathbf{y}_i. \psi_i$, with all φ_j, ψ_i Horn. But validity/provability of σ is equivalent to the validity/provability of all the sequents $\varphi_j \vdash_{\mathbf{x}, \mathbf{z}_j} \bigvee_I \exists \mathbf{y}_i. \psi_i$; so this reduces to the preceding case, and we are done. \square

5 Reflection and choice principles

In a classical meta-theory the construction of $\text{Ch}_{\mathbb{T}}(A)$ shows that the category of models for a regular theory is weakly reflective in the category of structures [AHRT02, II.10]. That is to say, that any homomorphism from a structure A to a \mathbb{T} -model M factors (not necessarily uniquely) through $\eta_A : A \rightarrow \text{Ch}_{\mathbb{T}}(A)$.

$$\begin{array}{ccc} A & \longrightarrow & M \\ \eta_A \downarrow & \nearrow h & \\ \text{Ch}_{\mathbb{T}}(A) & & \end{array}$$

The classical proof of this uses the axiom of choice; it turns out that this is unavoidable. We show in this section that various reflection properties for regular theories are in fact equivalent to choice principles in the ambient set theory.

Background

We start by recalling some further background in category theory and constructive choice principles.

Definition 5.1. Let \mathcal{C} be a category, with an initial object 0 and terminal object 1 .

- For maps m, p , say m is *weakly (left) orthogonal* to p if for every commutative square from m to p , there exists some diagonal filler.

$$\begin{array}{ccc} A & \longrightarrow & C \\ m \downarrow & \nearrow & \downarrow p \\ B & \longrightarrow & D \end{array}$$

- A family of maps with common domain, $m_i : A \rightarrow B_i$ ($i \in I$), is *jointly weakly (left) orthogonal* to $p : C \rightarrow D$ if for every $h : A \rightarrow C$ and family $k_i : B_i \rightarrow D$ such that $k_i m_i = p h$, there is some $i \in I$ for which there is a diagonal filler $e : B_i \rightarrow C$.
- A map (resp. family) is *weakly (resp. jointly weakly) orthogonal* to an object C if it is (jointly) weakly orthogonal to the unique map $!_C : C \rightarrow 1$.⁶
- A family of objects A_i is *jointly weakly initial* if for every object X , for some i there is a map $A_i \rightarrow X$. (Equivalently, if the maps $0 \rightarrow A_i$ are jointly weakly orthogonal to all objects.)
- An object X is *projective* if the map $0 \rightarrow X$ is weakly orthogonal to all epimorphisms.

⁶Viewed from the right-hand side, these are exactly injectivity of an object with respect to a map or cone, in the sense of [AR94, 4.1, 4.14].

We recall various choice principles, and the relationships between them.

Definition 5.2.

- The *axiom of choice* (AC) states just that all sets are projective.
- The *presentation axiom* (PAx) [Acz78, Rato6], also known as CoSHEP (“category of sets has enough projectives”), states that for every set X , there is some cover (i.e. surjection) $s : Y \rightarrow X$, with Y projective.
- The axiom WIC (“*weakly initial cover*”) states that for every set X there is some cover $s : Y \rightarrow X$, weakly initial in the category of covers of X .
- The axiom WISC (“*weakly initial set of covers*”) (given in [vdBM14], there called the weak axiom of multiple choice) states that for every set X , there is some family of covers $(s_i : Y_i \rightarrow X)_{i \in I}$, jointly weakly initial in the category of covers of X .

(To our knowledge, WIC has not been previously considered in the literature.)

Definition 5.3.

- *Dependent choice* (DC) is the statement that for any set A and entire set-relation $R \subseteq A \times A$, if R , there is some function $f : \mathbb{N} \rightarrow A$ such that $f(0) = x_0$ and for all $n \in \mathbb{N}$, $(x_n, x_{n+1}) \in R$.
- *Relativised dependent choice* (RDC), which might also be called *class dependent choice*, is the scheme asserting for all formulas $\alpha(x)$, $\rho(x, y)$ (possibly with further free variables) that if ρ is an entire relation on α , then there is some function f on \mathbb{N} such that $f(0) = x_0$ and for all $n \in \mathbb{N}$, $\alpha(x_n)$ and $\rho(x_n, x_{n+1})$.
- The *relation reflection scheme* (RRS) [Aczo8] is the scheme asserting for each pair of formulas $\alpha(x)$, $\rho(x, y)$ (possibly with further free variables) that if ρ is an entire relation on α , then for any set A that is a subclass of α , there is some set B , also a subclass of α , such that $A \subseteq B$ and the restriction of ρ to B is still entire.

We recall known relationships between these, and add WIC to the mix.

Proposition 5.4.

1. $AC \Rightarrow PAx \Rightarrow WIC \Rightarrow WISC$.
2. $RDC \Leftrightarrow DC + RRS$.
3. $WIC \Rightarrow DC$.

Proof. $AC \Rightarrow PAx \Rightarrow WISC$ is given in [Rato6]. The interpolation of WIC is straightforward: any cover by a projective set is weakly initial, and the singleton of any weakly initial cover is a jointly weakly initial set.

$RDC \Leftrightarrow DC + RRS$ is [Aczo8, Thm. 2.4].

For $WIC \Rightarrow DC$, note that the argument for $PAx \Rightarrow DC$ in [Bla79, Thm. 6.2] requires only weak initiality of the cover. \square

Proposition 5.5. (*Meta-theorem.*) *Even under $ZF + REA$, $WISC$ does not imply WIC .*

Proof. [Rato6, Cor. 6.11] shows $ZF + AMC + REA \not\Rightarrow AC_\omega$. But by Proposition 5.4, $AMC \Rightarrow WISC$, and $WIC \Rightarrow DC \Rightarrow AC_\omega$, under CZF. So $ZF + REA + WISC \not\Rightarrow WIC$. \square

In fact, WIC suffices for many applications of PAx in the literature. It seems likely to us that WIC is strictly weaker than PAx, but this does not seem to be obvious from existing results.

We will use DC reformulated in terms of graphs (in the category-theorist’s sense).

Definition 5.6. A *graph* G consists a set G_0 and a function $G_1 : (G_0 \times G_0) \rightarrow \mathbf{Set}$. A *class graph* (meta-definition) is a formula $\gamma_0(x)$ together with a formula $\gamma_1(x, y, f)$, possibly with further free variables (“parameters”). Any graph $G = (G_0, G_1)$ may be considered as a class graph, given by the formulas $x \in G_0, f \in G_1(x, y)$.

A class graph $(\gamma_0(x), \gamma_1(x, y, f))$ is *entire* if for each x such that $\gamma_0(x)$ holds, there exist some y, f such that $\gamma_0(y)$ and $\gamma_1(x, y, f)$.

A *branch*⁷ in a class graph $(\gamma_0(x), \gamma_1(x, y, f))$ is a function b from \mathbb{N} to pairs (x_n, f_n) such that for each n , both $\gamma_0(x_n)$ and $\gamma_1(x_n, x_{n+1}, f_n)$ holds. For a (set) graph G , write $Br_G(x)$ for the set of branches starting from $x \in G_0$.

Proposition 5.7. (*Over CZF.*) *DC is equivalent to the statement: for every entire graph (G_0, G_1) and $x_0 \in G_0$, there is some branch in (G_0, G_1) starting from x_0 . RDC is equivalent to the scheme asserting for each class graph $(\gamma_0(x), \gamma_1(x, y, f))$ that for all values of the parameters, if $(\gamma_0(x), \gamma_1(x, y, f))$ is entire, then for all x_0 such that $\gamma_0(x_0)$ holds, there is some branch starting from x_0 .* \square

Reflection principles for regular theories


With this background recalled, we are ready to start on the meat of this section: showing that reflection principles for regular theories are equivalent to choice principles. The main work is in building suitable weak reflections. For this, we will proceed along similar lines to Section 4 above: we will give generalisations of the chase construction, first in the purely relational equality-free case, and then deduce from this the general case.

⁷*Path* and *trail* are more established for this; but they are also often used just for *finite* paths, so to avoid ambiguity we borrow *branch* from the special case of trees.

For the next few statements, fix a theory \mathbb{T} over a signature Σ , and work in the category of Σ -structures.

Definition 5.8.

- A map $f : A \longrightarrow B$ is *weakly \mathbb{T} -reflective* if it is weakly orthogonal to all models of \mathbb{T} .



- A family of maps $A \longrightarrow B_i$ is *jointly weakly \mathbb{T} -reflective* if it is jointly weakly orthogonal to all models of \mathbb{T} .

Definition 5.9. (Reflection principles for theories.)

- \mathbb{T} satisfies *weak reflection* if for every Σ -structure A there is some weakly \mathbb{T} -reflective map from A into a model of \mathbb{T} .
- \mathbb{T} satisfies *joint weak reflection* if for every Σ -structure A there is some jointly weakly \mathbb{T} -reflective set of maps from A into models of \mathbb{T} .
- \mathbb{T} satisfies *functorial weak reflection* if there is a functor $F : \text{Str}(\Sigma) \longrightarrow \text{Mod}(\mathbb{T})$ and natural transformation $\eta : 1_{\text{Str}(\Sigma)} \longrightarrow I \cdot F$ such that each $\eta_A : A \longrightarrow F(A)$ is weakly \mathbb{T} -reflective. (Here I is the inclusion functor $\text{Mod}(\mathbb{T}) \hookrightarrow \text{Str}(\Sigma)$.)

In order to generalise the chase, we abstract the property of the one-step extension that ensures the chase yields a model of \mathbb{T} :

Definition 5.10. (\mathbb{T} regular normal.) A map $f : A \longrightarrow B$ is \mathbb{T} -*productive* if for each axiom $\varphi(\mathbf{x}) \vdash_{\mathbf{x}} \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$ of \mathbb{T} and $\mathbf{a} \in \llbracket \mathbf{x} \mid \varphi \rrbracket^A$, there is some $\mathbf{b} \in B^y$ such that $(f(\mathbf{a}), \mathbf{b}) \in \llbracket \mathbf{x}, \mathbf{y} \mid \psi \rrbracket^B$.

Proposition 5.11. (\mathbb{T} regular normal.) *The colimit of an ω -chain of \mathbb{T} -productive maps is a model of \mathbb{T} .*

Proof. Immediate by finitariness of the axioms of \mathbb{T} . □

Proposition 5.12. (Σ relational, \mathbb{T} regular normal equality-free.)

1. Assuming WISC, every Σ -structure has some jointly weakly \mathbb{T} -reflective set of \mathbb{T} -productive extensions.
2. Assuming WIC, every Σ -structure has some \mathbb{T} -productive and weakly \mathbb{T} -reflective extension.
3. Assuming AC, the functorial extension $\iota_A : A \longrightarrow \text{ex}(A)$ is always weakly \mathbb{T} -reflective (besides being \mathbb{T} -productive).

Proof. Let A be a Σ -structure. Recall from Definition 4.4 the one-step \mathbb{T} -extension $\iota_A : A \rightarrow \text{ex}_{\mathbb{T}}(A)$; we now generalise this construction.

Take M_A to be the set of pairs (α, \mathbf{a}) such that α is an axiom of \mathbb{T} , with the form $\varphi(\mathbf{x}) \vdash \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$, and $\mathbf{a} \in \llbracket \mathbf{x} \mid \varphi(\mathbf{x}) \rrbracket^A$. For any map $p : X \rightarrow M_A$, define a structure $A[p]$ by taking $|A[p]|$ to be the disjoint union of A with the set of pairs (x, i) such that $x \in X$ and $(p(x), i) \in \text{ex}(A)$, i.e. $i \in \mathbf{y}$, where the axiom component of $p(x)$ has the form $\varphi(\mathbf{x}) \vdash \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$. Take the structure on $A[p]$ to be induced by the evident map $|A[p]| \rightarrow |\text{ex}(A)|$. There is an evident homomorphism $\iota_A^p : A \rightarrow A[p]$, and if p is surjective, then ι^p is \mathbb{T} -productive.

Now if $\langle p_i : M_i \rightarrow M_A \rangle_{i \in I}$ is a weakly initial set of covers, we claim that the extensions $\iota^{p_i} : A \rightarrow A[p_i]$ are jointly weakly \mathbb{T} -reflective. Suppose $f : A \rightarrow B$, where $B \models \mathbb{T}$. Then for each $(\alpha, \mathbf{a}) \in M_A$, where α has form $\varphi(\mathbf{x}) \vdash \exists \mathbf{y}. \psi(\mathbf{x}, \mathbf{y})$, there is some $\mathbf{b} \in B^{\mathbf{y}}$ such that $(f(\mathbf{a}), \mathbf{b}) \in \llbracket \mathbf{y} \mid \psi(\mathbf{x}, \mathbf{y}) \rrbracket^B$. So by weak initiality, for some i there is a function giving for each $m \in M_i$ some \mathbf{b} witnessing this existential for $p(m)$. This induces a homomorphism $A[p_i] \rightarrow B$ extending f .

In particular, if $p : M' \rightarrow M$ is a single weakly initial cover, then ι^p is weakly \mathbb{T} -reflective.

Finally, assuming AC, the identity 1_{M_A} is itself weakly initial, so $\iota^{1_{M_A}}$ is weakly \mathbb{T} -reflective. But $\iota^{1_{M_A}} \cong \iota_A$, under A . \square

The preceding proposition gives us suitably reflective one-step extensions, or families thereof. With the aid of (R)DC, we can now take colimits of chains of these to give suitably reflective (families of) maps into models of \mathbb{T} .

Proposition 5.13. *Assume DC. Let \mathbf{C} be a locally small category with colimits of ω -chains, and X an object of \mathbf{C} .*

1. *For any ω -chain $F : \omega \rightarrow \mathbf{C}$ of maps weakly orthogonal to X , the transfinite composite $F(0) \rightarrow \varinjlim_i F(i)$ is again weakly orthogonal to X .*
2. *Let G be a graph, and $F : G \rightarrow \mathbf{C}$ a diagram such that for each vertex $i \in G_0$, the set $\langle F(a) : F(i) \rightarrow F(j) \rangle_{j \in G_0, a \in G_1(i, j)}$ is jointly weakly orthogonal to X . Then for any $i_0 \in G_0$ the set of all maps $\{F(i_0) \rightarrow F_b \mid b \in \text{Br}_G(i_0)\}$ is jointly weakly orthogonal to X . (Here F_b denotes the colimit of F along a branch b .)*

Proof. The first statement is straightforward.

For the second, consider the comma graph F/X , where $(F/X)_0$ consists of pairs (i, f) such that $i \in G_0$, $f : F(i) \rightarrow X$, and $(F/X)_1((i, f), (j, g))$ is the set of arrows $a \in G_1(i, j)$ such that $g \cdot F(a) = f$. The joint reflectivity assumption on F implies directly that F/X is entire.

So given any $f_0 : F(i_0) \rightarrow X$, DC gives us a branch in F/X starting from (i_0, f_0) . The first components of this form a branch b in G ; and the second components induce a map $F_b \rightarrow X$ extending f_0 . \square

Proposition 5.14. *(Σ relational, \mathbb{T} regular normal equality-free.) Assuming WIC and either RRS or RDC, every structure has some weakly \mathbb{T} -reflective map into a \mathbb{T} -model.*

Proof. (By Prop. 5.4, RRS is equivalent to RDC in the presence of WIC.)

Let A be any Σ -structure. By Proposition 5.12 together with RDC, we can take some branch $A = A_0 \longrightarrow A_1 \longrightarrow A_2 \longrightarrow \dots$, in which each step is \mathbb{T} -productive and weakly \mathbb{T} -reflective.

Now set $\bar{A} = \varinjlim_i A_i$. By Proposition 5.11, $\bar{A} \models \mathbb{T}$; and by Proposition 5.13, $A \longrightarrow \bar{A}$ is weakly \mathbb{T} -reflective. \square

Proposition 5.15. (*Σ relational, \mathbb{T} regular normal equality-free.*) *Assuming RDC and WISC, every structure has some jointly weakly \mathbb{T} -reflective set of maps into \mathbb{T} -models.*

Proof. First note that given any set \mathcal{X} of Σ -structures, there exists some set \mathcal{Y} of Σ -structures such that for each $A \in \mathcal{X}$, the set of all \mathbb{T} -productive maps from A into structures in \mathcal{Y} is jointly weakly \mathbb{T} -reflective. This follows from Prop. 5.12 together with Collection, and the fact that any superset of a jointly weakly reflective set is jointly weakly reflective.

Now let A_0 be any Σ -structure. By RRS applied to $\{A_0\}$ and the entire relation of the previous paragraph, there is some set \mathcal{X} of Σ -structures containing A_0 such that for each $A \in \mathcal{X}$, the set of all \mathbb{T} -productive maps from A into structures in \mathcal{X} is jointly weakly \mathbb{T} -reflective.

Take G to be the graph with objects \mathcal{X} , and with arrows all \mathbb{T} -productive maps between these. Then Prop. 5.11, applied to the evident diagram $G \longrightarrow \text{Str}(\Sigma)$, tells us that the maps $\langle A_0 \longrightarrow A_b \rangle_{b \in \text{Br}_G(A_0)}$ form a jointly \mathbb{T} -reflective family. But by Prop. 5.11, each such A_b is a model of \mathbb{T} ; so we are done. \square

Finally, we can put all the pieces together:

Theorem 5.16 (“Choice principles \Leftrightarrow reflection principles”).

1. *AC is equivalent to “all regular theories have functorial weak reflection”.*
2. *Assuming RRS (or a fortiori RDC), WIC is equivalent to “all regular theories satisfy weak reflection”.*
3. *Assuming RDC, WISC is equivalent to “all regular theories satisfy joint weak reflection”.*

Proof. We tackle first the \Rightarrow directions of the three equivalences.

For Σ purely relational and \mathbb{T} equality-free, these are Prop. 5.12 (3), Prop. 5.14, and Prop. 5.15 respectively.

For Σ purely relational but \mathbb{T} not necessarily equality-free, these follow from the equality-free cases by elimination of equalities, as set up in Def. 4.10–Prop. 4.13. For weak reflection, given $A \in \text{Str}(\Sigma)$, consider it as a Σ_E structure $e(A)$, take a weakly $(\mathbb{T}^E \cup \mathbb{E}_\Sigma)$ -reflective map into a $(\mathbb{T}^E \cup \mathbb{E}_\Sigma)$ -model M , and apply q to get a map from A ($\cong q(e(A))$) into the quotient \mathbb{T} -model $q(M)$. The resulting map $A \longrightarrow q(M)$ is weakly \mathbb{T} -reflective, by the reflection $q \dashv e$. The cases for joint and functorial reflection are analogous.

Finally, the case for general Σ and \mathbb{T} reduce to the relational case by elimination of function symbols, Proposition 4.14. Replace Σ by an associated signature $\bar{\Sigma}$ and theory \mathbb{F}_{Σ} thereover, such that $\text{Str}(\Sigma) \simeq \text{Mod}(\mathbb{F}_{\Sigma})$, and moreover this restricts to $\text{Mod}(\mathbb{T}) \simeq \text{Mod}(\bar{\mathbb{T}})$ for some theory $\bar{\mathbb{T}}$ associated to \mathbb{T} . Then the (functorial, weak, joint) reflection from $\text{Str}(\bar{\Sigma})$ into $\text{Mod}(\bar{\mathbb{T}})$ restricts to a similar reflection from $\text{Str}(\Sigma)$ into $\text{Mod}(\mathbb{T})$.

This completes the \Rightarrow directions of each equivalence.

For the \Leftarrow directions, consider the signature Σ with just one unary predicate P and one binary predicate R , and the theory \mathbb{T} consisting just of the axiom $P(x) \vdash_x \exists y. R(x, y)$. There are evident functors $F : \mathbf{Set}^{\rightarrow} \rightarrow \text{Str}(\Sigma)$ and $G : \text{Str}(\Sigma) \rightarrow \mathbf{Set}^{\rightarrow}$, acting on objects as follows:

- F sends a function $f : X \rightarrow Y$ to the Σ -structure on $X + Y$ with $\llbracket P \rrbracket = v_1[X]$, $\llbracket R \rrbracket = \{(v_1(fy), v_2(y)) \mid y \in Y\}$;
- G sends a Σ -structure A to the projection map $\llbracket x, y \mid P(x) \wedge R(x, y) \rrbracket^A \rightarrow \llbracket x \mid P(x) \rrbracket^A$, which we denote by $g_A : G_1(A) \rightarrow G_0(A)$.

It is easy to check that F is left adjoint to G ; F is full and faithful; and F and G preserve and reflect the subcategories of \mathbb{T} -models and covers.

We identify any set X with the unique map $0 \rightarrow X$, as an object of $\mathbf{Set}^{\rightarrow}$. In particular, by 1 we mean for now the map $0 \rightarrow 1$, not the terminal object of $\mathbf{Set}^{\rightarrow}$.

We now take the three \Leftarrow directions in turn.

Firstly, suppose $W : \text{Str}(\Sigma) \rightarrow \text{Str}(\Sigma)$ is a functorial weak reflection, with unit $\eta : 1 \rightarrow W$; we will show that AC holds.

Take any singleton $1 = \{\star\}$. Then $W(F(1))$ contains some element \diamond , such that $W(F(1)) \models R(\eta(v_1(\star)), \diamond)$.

Now, let $p : Y \rightarrow X$ be any cover; we will construct a splitting of p .

There is a unique map $\iota : X \rightarrow p$ in $\mathbf{Set}^{\rightarrow}$ over 1_X . This induces a map $F\iota : F(X) \rightarrow F(p)$, which by the universal property of the weak reflector factors through $W(F(X))$ as $F\iota = f \cdot \eta_{F(X)}$, for some f .

We can now obtain a splitting of p , by first finding a witnessing function for the existentials in $W(F(X))$, and then composing this with f . Precisely, each $x \in X$ determines a map $\hat{x} : 1 \rightarrow X$, and hence a homomorphism $\ulcorner x \urcorner := F\hat{x} : F(1) \rightarrow F(X)$. For each $x \in X$, naturality of η ensures that $W(F(X)) \models R(\eta_{F(X)}(v_1(x)), W\ulcorner x \urcorner(\diamond))$. So the map $X \rightarrow X + Y (= |F(p)|)$ sending x to $f(W\ulcorner x \urcorner(\diamond))$ factors through Y ; and this factorisation gives a section of p .

$$\begin{array}{ccccc}
 W(F(1)) & \xrightarrow{W\ulcorner x \urcorner} & W(F(X)) & \xrightarrow{f} & F(p) \\
 \uparrow \eta_{F(1)} & & \uparrow \eta_{F(X)} & \nearrow F\iota & \\
 F(1) & \xrightarrow{\ulcorner x \urcorner} & F(X) & &
 \end{array}$$

Secondly, suppose weak reflection holds for \mathbb{T} ; we want to show that WIC holds. So let X be any set; we will find a weakly initial cover for X .

Consider the Σ -structure $F(X)$ as before, and let $\eta : F(X) \twoheadrightarrow A$ be some weak reflection into a \mathbb{T} -model. Applying G to A gives a cover $g_A : G_1(A) \rightarrow G_0(A)$, and transposing η under the adjunction yields a map $\hat{\eta} : X \twoheadrightarrow G_0(A)$. Pulling back g_A along $\hat{\eta}$ gives a cover $p : \bar{X} \twoheadrightarrow X$, and a factorisation of $\hat{\eta}$ through the canonical inclusion $\iota : X \twoheadrightarrow p$.

Transposing again, we get a factorisation of $\eta : F(X) \twoheadrightarrow A$ through $F(\iota) : F(X) \twoheadrightarrow F(p)$. Now $F(\iota) : X \twoheadrightarrow F(p)$ is weakly \mathbb{T} -reflective, since η was. Since F is full and faithful and sends covers to models of \mathbb{T} , it follows that $\iota : X \twoheadrightarrow p$ is left-orthogonal to all covers (viewed as objects of $\mathbf{Set}^{\rightarrow}$). That is, $p : \bar{X} \twoheadrightarrow X$ is a weakly initial cover.

Finally, to show that joint reflection for \mathbb{T} implies WISC, we perform the same construction as above on each map in a jointly reflective set $\eta_i : F(X) \twoheadrightarrow A_i$, and obtain (by essentially the same argument) a weakly initial set of covers p_i . \square

References

- [Acz78] Peter Aczel, *The type theoretic interpretation of constructive set theory*, Logic Colloquium '77 (Proc. Conf., Wrocław, 1977) (A. MacIntyre, L. Pacholski, and J. Paris, eds.), Stud. Logic Foundations Math., vol. 96, North-Holland, Amsterdam-New York, 1978, pp. 55–66, doi: [10.1016/S0049-237X\(08\)71989-X](https://doi.org/10.1016/S0049-237X(08)71989-X).
- [Aczo8] ———, *The relation reflection scheme*, Mathematical Logic Quarterly **54** (2008), no. 1, 5–11, doi: [10.1002/malq.200710035](https://doi.org/10.1002/malq.200710035).
- [AHRT02] Jiří Adámek, Horst Herrlich, Jiří Rosický, and Walter Tholen, *On a generalized small-object argument for the injective subcategory problem*, Cah. Topol. Géom. Différ. Catég. **43** (2002), no. 2, 83–106.
- [AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu, *Foundations of databases*, Addison-Wesley, 1995.
- [AR94] Jiří Adámek and Jiří Rosický, *Locally presentable and accessible categories*, London Mathematical Society Lecture Note Series, vol. 189, Cambridge University Press, Cambridge, 1994, doi: [10.1017/CBO9780511600579](https://doi.org/10.1017/CBO9780511600579).
- [Bla79] Andreas Blass, *Injectivity, projectivity, and the axiom of choice*, Trans. Amer. Math. Soc. **255** (1979), 31–59, doi: [10.2307/1998165](https://doi.org/10.2307/1998165).
- [BM77] J. L. Bell and M. Machover, *A course in mathematical logic*, North-Holland Publishing Co., Amsterdam-New York-Oxford, 1977.

- [CLR01] Michel Coste, Henri Lombardi, and Marie-Françoise Roy, *Dynamical method in algebra: effective Nullstellensätze*, Ann. Pure Appl. Logic **111** (2001), no. 3, 203–256, doi:10.1016/S0168-0072(01)00026-4.
- [Dia08] Răzvan Diaconescu, *Institution-independent model theory*, Studies in Universal Logic, Birkhäuser Verlag, Basel, 2008.
- [Fit69] Melvin Chris Fitting, *Intuitionistic logic, model theory and forcing*, Studies in Logic and the Foundations of Mathematics, North-Holland Publishing Co., Amsterdam-London, 1969, doi:10.1016/s0049-237x(08)70794-8.
- [Hof95] Martin Hofmann, *Extensional concepts in intensional type theory*, Ph.D. thesis, University of Edinburgh, 1995, <http://www.lfcs.inf.ed.ac.uk/reports/95/ECS-LFCS-95-327/>.
- [Joh02] Peter T. Johnstone, *Sketches of an elephant: a topos theory compendium. Vols. 1 and 2*, Oxford Logic Guides, vol. 43 and 44, The Clarendon Press Oxford University Press, New York, 2002.
- [LS86] J. Lambek and P. J. Scott, *Introduction to higher order categorical logic*, Cambridge Studies in Advanced Mathematics, vol. 7, Cambridge University Press, Cambridge, 1986.
- [Rat06] Michael Rathjen, *Choice principles in constructive and classical set theories*, Logic Colloquium '02, Lect. Notes Log., vol. 27, Assoc. Symbol. Logic, La Jolla, CA, 2006, pp. 299–326.
- [TvD88] Anne S. Troelstra and Dirk van Dalen, *Constructivism in mathematics*, Studies in Logic and the Foundations of Mathematics, vol. 121 and 123, North-Holland Publishing Co., Amsterdam, 1988.
- [vdBM14] Benno van den Berg and Ieke Moerdijk, *The axiom of multiple choice and models for constructive set theory*, J. Math. Log. **14** (2014), no. 1, 1450005, 28, doi:10.1142/S0219061314500056.
- [Vel76] Wim Veldman, *An intuitionistic completeness theorem for intuitionistic predicate logic*, J. Symbolic Logic **41** (1976), no. 1, 159–166.