

Principal Type Inference under a Prefix

A Fresh Look at Static Overloading

Microsoft Research Technical Report, MSR-TR-2024-34, 2025-05-19 (v7)

DAAN LEIJEN, Microsoft Research, USA

WENJIA YE, National University of Singapore, Singapore, The University of Hong Kong, China

At the heart of the Damas-Hindley-Milner (HM) type system lies the abstraction rule which derives a function type for a lambda expression. In this rule, the type of the parameter can be “guessed”, and can be any type that fits the derivation. The beauty of the HM system is that there always exists a most general type that encompasses all possible derivations – Algorithm W is used to infer these most general types in practice.

Unfortunately, this property is also the bane of the HM type rules. Many languages extend HM typing with additional features which often require complex side conditions to the type rules to maintain principal types. For example, various type systems for impredicative type inference, like HMF, FreezeML, or Boxy types, require let-bindings to always assign most general types. Such a restriction is difficult to specify as a logical deduction rule though, as it ranges over all possible derivations. Despite these complications, the actual implementations of various type inference algorithms are usually straightforward extensions of algorithm W, and from an implementation perspective, much of the complexity of various type system extensions, like boxes or polymorphic weights, is in some sense artificial.

In this article we rephrase the HM type rules as *type inference under a prefix*, called HMQ. HMQ is sound and complete with respect to the HM type rules, but always derives principal types that correspond to the types inferred by algorithm W. The HMQ type rules are close to the clarity of the declarative HM type rules, but also specific enough to “read off” an inference algorithm, and can form an excellent basis to describe type system extensions in practice. We show in particular how to describe the FreezeML and HMF systems in terms of inference under a prefix, and how we no longer require complex side conditions. We also show a novel formalization of static overloading in HMQ as implemented in Koka language.

Additional Key Words and Phrases: Type Inference, Damas-Hindley-Milner, Static Overloading

1 Introduction

At the heart of Damas-Hindley-Milner style type inference [Damas and Milner 1982; Hindley 1969; Milner 1978] lies the abstraction rule which infers a type $\tau_1 \rightarrow \tau_2$ for a lambda expression $\lambda x. e$ under a type environment Γ :

$$\frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \text{FUN}$$

Interestingly, the type τ_1 of the parameter x occurs free and is “guessed” – it can be any type that fits the derivation. This encompasses both the *beauty*, but also the *bane*, of the Damas-Hindley-Milner (HM) type rules.

For example, for the identity function $\lambda x. x$ we can derive many types, like $\text{int} \rightarrow \text{int}$, or $\text{bool} \rightarrow \text{bool}$ etc. That seems a problem at first, but the beauty of the HM type rules is that there always exists a derivation with a most general type of which all other possible derivations are an instance – in the identity case the type $\forall \alpha. \alpha \rightarrow \alpha$. Moreover, there exist an algorithm W that always infers these most general types which is widely used in practice for HM style type inference.

Nevertheless, this rule is also the bane of HM type inference. In practice many languages extend HM typing with various extensions and it turns out that the inference rules need to be restricted in often complicated ways. For example, Leijen [2008] describes the HMF system that allows for impredicative higher-ranked types. He gives the following example:

let *wrapl* $x\ y = [y]$ in *wrapl* *ids* *id*

where *ids* has the impredicative type $[\forall\alpha.\alpha \rightarrow \alpha]$ (i.e. a list of polymorphic identity functions). If *wrapl* is given its most general type, namely $\forall\alpha\beta. \alpha \rightarrow \beta \rightarrow [\beta]$, we can derive the type $\forall\alpha.[\alpha \rightarrow \alpha]$ for the body. However, if we use the [FUN] rule to “guess” a less general type for *wrapl*, namely $\forall\alpha.[\alpha] \rightarrow \alpha \rightarrow [\alpha]$, then we can derive (in HMF) the type $[\forall\alpha.\alpha \rightarrow \alpha]$ for the body (as the shared α now matches with the polymorphic identity type). Unfortunately, these types are incomparable – neither is an instance of the other – and we lose principal type derivations. To fix this issue, the HMF system includes a side-condition on the let-rule to always assign most general types:

$$\frac{\Gamma \vdash e_1 : \sigma_1 \quad \Gamma, x : \sigma_1 \vdash e_2 : \sigma_2 \quad \forall\sigma. \Gamma \vdash e_1 : \sigma \Rightarrow \sigma_1 \sqsubseteq \sigma}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \sigma_2} \text{HMF-LET}$$

From a logical perspective this condition is quite unsatisfactory. It is no longer a natural deduction rule since the condition ranges negatively over *all* possible derivations, making it more difficult to reason about. In another more recent example, Emrich et al. [2022] describe the FreezeML system that also includes a side-condition on the let-rule that ranges over all possible derivations, and they write “the reader may be concerned about whether the typing judgement is well-defined given that it appears in a negative position in the definition of principal. [...] the definition is nevertheless well founded by indexing by untyped terms or the height of the derivation tree”. Vytiniotis et al [2006,§6] also introduce a similar let rule in the context of boxy type inference, and similar ideas were also used by Leroy and Mauny [1993] for the typing of dynamics in ML, and by Garrigue and Rémy [1999,§5] in their extension of ML with semi-explicit first-class polymorphism.

An example of a novel extension that we describe in this article is *static overloading*. With static overloading, we allow a function *f* to be defined in different modules, say *modi* and *modb*, where we can give their fully qualified names as *modi/f* and *modb/f*. Suppose they have the types:

modi/f : *int* → *int* *modb/f* : *bool* → *int*

The idea is now to use local type information to allow a programmer to write just *f* and have it be resolved to either definition. For example, *f* 1 would be elaborated to *modi/f* 1. Since the overloading is static, we reject expressions where the definition cannot be resolved uniquely. For example a bare *f* is rejected, but we would also like to reject $\lambda x. f\ x$. Unfortunately, the [FUN] rule again allows us to “guess” the type *int* for *x*, in which case we could elaborate to $\lambda x : \text{int}. \text{modi/f } x$ – but also we could guess the type *bool* for *x* and derive $\lambda x : \text{bool}. \text{modb/f } x$. Again, the flexibility of the [FUN] rule causes non-principal derivations.

The interesting part of all the previous examples is that it is only difficult to extend the declarative HM type rules with the new extensions – but for all of the example systems, the changes to the actual type inference *implementation*, based on algorithm W, are usually quite straightforward! For example, all HM based type inference algorithms already infer most general types for let-bindings – as required by HMF, FreezeML, or Boxy type inference; and they will already use a general type α for the *x* binding in the static overloading example (and not some arbitrary type *int* or *bool*). As such, most of the complexity that we see in the type rules of these systems are in some sense artificial, and are only needed to constrain the high level declarative rules enough to match the inference algorithm more closely! This leads one to ask if we can perhaps create a more restricted set of inference rules that match the inference algorithm more closely while still being close to the clarity of the HM type rules.

- In this article we rephrase the HM type rules as *type inference under a prefix*, called HMQ. These new rules always derive principal types that correspond to the types inferred by algorithm W [Damas and Milner 1982] (and we can use algorithm W unchanged to infer types for HMQ as well). In particular, if we can derive a type σ_1 in HM, then we can also derive a type σ_2 in HMQ such that σ_2 can be instantiated to σ_1 .

$e ::= x \mid f$	(variables)	$\tau ::= \alpha$	(type variable)
$ e e$	(application)	$ \tau \rightarrow \tau$	(function arrow)
$ \lambda x. e$	(function)	$ \text{int} \mid \text{bool} \mid \dots$	(type constants)
$ \text{let } x = e \text{ in } e$	(let binding)	$\sigma ::= \forall \alpha. \sigma$	(universal quant.)
		$ \tau$	(monomorphic type)
$\Gamma ::= x_1 : \sigma_1, \dots, x_n : \sigma_n$	(type environment)	$\frac{\bar{\beta} \notin \text{ftv}(\forall \bar{\alpha}. \tau)}{\forall \bar{\alpha}. \tau \sqsubseteq \forall \bar{\beta}. [\bar{\alpha} := \bar{\tau}] \tau} \text{INSTANCE}$	
$Q ::= \{\alpha_1 = \tau_1, \dots, \alpha_n = \tau_n\}$	(prefix)		

Fig. 1. Syntax of types and terms.

- Inspired by MLF [Le Botlan and Rémy 2003], we use a *prefix* Q to propagate type variable constraints in such a way that there is no need for complex side conditions and we retain natural deduction rules. As such, we believe the HMQ type rules are close to the clarity of the HM type rules and can form an excellent basis to describe type system extensions in practice.
- We provide evidence for this by rephrasing FreezeML and HMF in terms of HMQ (Section 4) where there is no need anymore for the elaborate side conditions of each system. The original type rules of each system differ significantly, but building from a common HMQ specification, we can now precisely characterize the relationship between the two systems.
- Finally, we show a novel formalization of static overloading (Section 6, as implemented in the Koka language). The extension is surprisingly straightforward – requiring no changes to the base system and only adding additional rules to resolve identifiers based on their context.

2 Inference under a Prefix

The goal of HMQ is two-fold: First of all, we’d like the rules to be closer to algorithm W so we are able to “read off” the algorithm from the declarative type rules. At the same time though, we’d like to retain the clarity of the original HM rules as much as possible. HMQ can serve as foundation to specify practical type systems in a declarative way that serves both purposes: users can easily reason about what programs are accepted by the type checker, while compiler writers can derive sound implementations from those same rules.


2.1 Syntax

Figure 1 shows the syntax of our standard lambda calculus expressions e , mono-types τ , and polymorphic type schemes σ . The [INSTANCE] rule gives the general instantiation where we write $\sigma_1 \sqsubseteq \sigma_2$ when a type σ_1 can be instantiated to a type σ_2 . For example, $\forall \alpha \beta. \alpha \rightarrow \beta \sqsubseteq \forall \beta. \text{int} \rightarrow \beta \sqsubseteq \text{int} \rightarrow \text{bool}$. Note that we can only instantiate *bound* type variables $\bar{\alpha}$, and not the *free* type variables in σ_1 (written as $\text{ftv}(\forall \bar{\alpha}. \tau)$), and in particular, $\forall \alpha. \alpha \rightarrow \beta \sqsubseteq \text{int} \rightarrow \text{bool}$ does not hold. A prefix Q is a set of type variable bindings, and we describe this in detail later in this section. A type environment Γ gives the types of variables bound by a lambda or let binding. We write $\Gamma, x : \sigma$ to extend a type environment with a new binding $x : \sigma$ (replacing any previous binding for x in Γ).

2.2 Type Rules

Figure 2 defines the HMQ type inference rules, where a judgment $Q \mid \Gamma \vdash e : \sigma$ states that under a prefix Q and type environment Γ , we can derive type σ for the expression e . The prefix Q and the type σ are synthesized (i.e. output) while Γ and e are inherited (i.e. input).

We will go through the rules one-by-one, explaining the design decisions as we go. We start with the [VAR] and [GEN] rules that closely match the corresponding HM type rules (see Figure 8 in Appendix A): the [VAR] rule gets the type of a variable from the type environment, while [GEN] generalizes over free type variables that no longer occur in Γ (and Q in our case).



with $\models Q$

$$\begin{array}{c}
\frac{x : \sigma \in \Gamma}{\emptyset \mid \Gamma \vdash x : \sigma} \text{VAR} \quad \frac{Q \mid \Gamma \vdash e : \forall \alpha. \sigma \quad \text{fresh } \alpha}{Q \mid \Gamma \vdash e : \sigma} \text{INST} \quad \frac{Q \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q, \Gamma)}{Q \mid \Gamma \vdash e : \forall \alpha. \sigma} \text{GEN} \\
\\
\frac{Q \mid \Gamma, x : \alpha \vdash e : \tau \quad \text{fresh } \alpha}{Q \mid \Gamma \vdash \lambda x. e : \alpha \rightarrow \tau} \text{FUN} \quad \frac{Q \cdot \alpha = \tau \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q, \Gamma)}{Q \mid \Gamma \vdash e : [\alpha := \tau] \sigma} \text{GENSUB} \\
\\
\frac{Q_1 \mid \Gamma \vdash e_1 : \tau_1 \quad Q_2 \mid \Gamma \vdash e_2 : \tau_2 \quad Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha \quad \text{fresh } \alpha}{Q_1, Q_2, Q_3 \mid \Gamma \vdash e_1 e_2 : \alpha} \text{APP} \\
\\
\frac{Q_1 \mid \Gamma \vdash e_1 : \sigma \quad Q_2 \mid \Gamma, x : \sigma \vdash e_2 : \tau \quad \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{Q_1, Q_2 \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \text{LET}
\end{array}$$

Fig. 2. Type rules under a prefix

2.3 Do Not Guess Types

As argued in the introduction, the guessing of types in the lambda rule is both problematic for describing type system extensions, but also for implementing an inference algorithm – what type to guess? In HMQ we follow algorithm W and always infer an *abstract* type α for a lambda-bound parameter. In particular, in the [FUN] rule the type is now always a fresh variable α – just as in algorithm W (see Figure 10 in Appendix A.2). For example, we can derive the type of the polymorphic identity function as:

$$\frac{\frac{x : \alpha \in (\Gamma, x : \alpha)}{\emptyset \mid \Gamma, x : \alpha \vdash x : \alpha} \text{VAR} \quad \frac{\text{fresh } \alpha}{\emptyset \mid \Gamma \vdash \lambda x. x : \alpha \rightarrow \alpha} \text{FUN} \quad \frac{\alpha \notin \text{ftv}(\emptyset, \Gamma)}{\emptyset \mid \Gamma \vdash \lambda x. x : \forall \alpha. \alpha \rightarrow \alpha} \text{GEN}$$

Unlike the HM type rules there is no choice here for the type of the binding and we can *only* derive the type of the polymorphic identity function and not for example $\text{int} \rightarrow \text{int}$.

The other rule where we prevent guessing a type is the [INST] rule where we always instantiate directly to a fresh type α (where we rely on α -renaming to match the quantifier). Again, this corresponds to how algorithm W always instantiates using fresh type variables.

Note that we use $\text{fresh } \alpha$ notation to create fresh names α , not only such that α is fresh in the local rule, but also to ensure there is no other occurrence of $\text{fresh } \alpha$ in the derivation. This is a convenient notation for a more explicit formalization where we pass a fresh name supply using disjoint union for multiple sub-derivations – see Figure 7 in Section 8 for the full rules of HMQ. However, adding an explicit name supply clutters the rules somewhat while not adding any essential insight, and we prefer the fresh notation when applicable (i.e. when not doing proofs).

2.4 The Prefix

Clearly, we cannot always keep a parameter type abstract. For example, we'd like to infer the type $\text{int} \rightarrow \text{int}$ for the expression $\lambda x. \text{inc } x$. This is where we need the *prefix* Q , which is a set of type variable bounds $\alpha = \tau$ (similar to the *rigid* bounds of MLF [Le Botlan and Rémy 2003]):

$$Q ::= \{\alpha_1 = \tau_1, \dots, \alpha_n = \tau_n\}$$

$$\begin{array}{c}
\boxed{Q \vdash \tau \approx \tau} \\
\downarrow \quad \uparrow \quad \uparrow \\
\text{out} \quad \text{in} \quad \text{in}
\end{array}
\qquad
\frac{}{\emptyset \vdash \tau \approx \tau} \text{EQ-ID}
\qquad
\frac{\alpha \notin \text{ftv}(\tau)}{\{\alpha=\tau\} \vdash \alpha \approx \tau} \text{EQ-VAR}$$

$$\frac{Q_1 \vdash \tau_1 \approx \tau'_1 \quad Q_2 \vdash \tau_2 \approx \tau'_2}{Q_1, Q_2 \vdash \tau_1 \rightarrow \tau_2 \approx \tau'_1 \rightarrow \tau'_2} \text{EQ-FUN}
\qquad
\frac{Q \vdash \tau_2 \approx \tau_1}{Q \vdash \tau_1 \approx \tau_2} \text{EQ-REFL}$$

Fig. 3. Type equivalence under a prefix.

The binders α form the domain of Q , and the types τ form the range. The codomain of Q consists of the free type variables of the range, and we define $\text{ftv}(Q)$ as all free type variables in Q , where $\text{ftv}(Q) = \text{dom}(Q) \cup \text{codom}(Q)$. Note that a general prefix is just a collection of type variable bounds, and can for example have duplicate bindings, like $\{\alpha=\beta \rightarrow \text{int}, \alpha=\text{int} \rightarrow \gamma\}$ or $\{\alpha=\text{bool}, \alpha=\text{int}\}$.

We write $\theta \models Q$ if a substitution θ is a *solution* to Q that satisfies all the constraints in Q where $\forall (\alpha=\tau) \in Q. \theta\alpha = \theta\tau$. If there exists any solution to Q , we say that Q is *consistent* or *solvable*, and we denote this by writing $\text{just} \models Q$. For example, $\{\beta=\text{int}, \alpha=\beta \rightarrow \text{int}\}$ or $\{\alpha=\beta \rightarrow \text{int}, \alpha=\text{int} \rightarrow \gamma\}$ are consistent prefixes. Examples of inconsistent prefixes that do not have a solution, are prefixes with incompatible bindings, like $\{\alpha=\text{int}, \alpha=\text{bool}\}$, or with cyclic bindings, like $\{\alpha=\beta, \beta=\alpha \rightarrow \alpha\}$.

We call a least solution of a prefix Q a *prefix solution*, written as $\langle Q \rangle$, such that for any other solution $\theta \models Q$, the prefix solution is more general¹: $\langle Q \rangle \sqsubseteq \theta$. We write $Q[\tau]$ as a shorthand for applying the prefix solution as $\langle Q \rangle(\tau)$. Also, we sometimes leave out the angled brackets when the prefix substitution is clear from the context, and for example write $Q \sqsubseteq \theta$ for $\langle Q \rangle \sqsubseteq \theta$.

Finally, we consider two prefixes *equivalent* whenever their solution substitutions are equivalent: $Q_1 \equiv Q_2 \Leftrightarrow \langle Q_1 \rangle \equiv \langle Q_2 \rangle$. For example, we have $\{\alpha=\beta \rightarrow \text{int}, \alpha=\gamma \rightarrow \gamma\} \equiv \{\gamma=\text{int}, \beta=\gamma, \alpha=\gamma \rightarrow \gamma\} \equiv \{\beta=\text{int}, \gamma=\text{int}, \alpha=\text{int} \rightarrow \text{int}\}$. Similar to α -renaming, we can always substitute equivalent prefixes in type derivations.

2.4.1 Type Equivalence under a Prefix. A *consistent union* is written as Q_1, Q_2 and denotes the union $Q_1 \cup Q_2$ where $Q_1 \cup Q_2$ is solvable. We use this in the conclusion of most type rules to ensure that we can only derive consistent prefixes. The consistent union allows for a better declarative specification than using substitutions, since we can *compose* prefixes from different sub-derivations as Q_1, Q_2 , and we do not need to thread substitutions statefully through the rules.

The elegance of composable prefixes is shown in the definition of equivalence between types under a prefix as shown in Figure 3 (corresponding closely to unification). A rule $Q \vdash \tau_1 \approx \tau_2$ states that type τ_1 is equal to τ_2 under a prefix Q . Note how in the [EQ-FUN] rule we can compose the prefixes Q_1 and Q_2 from each sub derivation without needing to thread a substitution linearly through the derivations. It is straightforward to show that our definition of type equivalence is sound and complete:

Theorem 2.1. (Type equivalence under a prefix is sound)

If $Q \vdash \tau_1 \approx \tau_2$ then $Q[\tau_1] = Q[\tau_2]$.

Theorem 2.2. (Type equivalence under prefix is complete)

If $\theta\tau_1 = \theta\tau_2$, then there exists a Q such that $Q \vdash \tau_1 \approx \tau_2$ and $Q \sqsubseteq \theta$.

Soundness states that if we can derive that τ_1 and τ_2 are equivalent under a prefix Q , then the types are syntactically equal under the prefix solution: $Q[\tau_1] = Q[\tau_2]$. Completeness shows that if there exists any substitution θ that makes two types equal, then we can also derive that these types are equivalent under a prefix Q , and that such a prefix is also the “best” (most-general) solution: $\langle Q \rangle \sqsubseteq \theta$.

¹Following Pierce [2002, §22.4.1] we write $\theta_1 \sqsubseteq \theta_2$ to denote that θ_1 is a more-general (or less-specific) substitution than θ_2 , which holds if there exists some substitution θ' such that $\theta_2 = \theta' \circ \theta_1$.

2.4.2 Application. We use type equivalence in the HMQ application rule [APP] in Figure 2 to match the function type τ_1 with the argument type τ_2 and a fresh result type α :

$$Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha$$

Similar to parameter types, we use a fresh type variable α to represent the result type of the application. The application rule now corresponds directly to the usual implementation in algorithm W where one unifies with the function type (see Figure 10 in Appendix A.2). We can now derive a type for the application $\text{inc } x$ as:

$$\frac{\emptyset \mid \Gamma, x:\alpha \vdash \text{inc} : \text{int} \rightarrow \text{int} \quad \emptyset \mid \Gamma, x:\alpha \vdash x : \alpha \quad \{\alpha=\text{int}, \beta=\text{int}\} \vdash \text{int} \rightarrow \text{int} \approx \alpha \rightarrow \beta}{\{\alpha=\text{int}, \beta=\text{int}\} \mid \Gamma, x:\alpha \vdash \text{inc } x : \beta} \text{APP}$$

Note that in the judgement $Q \mid \Gamma \vdash e : \sigma$, both the inferred type σ and the prefix Q are synthesized (i.e. output). Moreover, the rules are carefully set up to ensure that the resulting Q only contains constraints that are “induced” by the structure of the program and types, where the set of constraints is minimal. In particular, the only possible leaf nodes of a derivation are [VAR] and the type equivalence rules [EQ-ID] and [EQ-VAR]. Since both [VAR] and [EQ-ID] have an empty prefix \emptyset , the *only* way to create (or dismiss depending on your viewpoint!) prefix constraints is through the [EQ-VAR] rule. This property is crucial as it ensures that, unlike the HM type rules, we can never “make up” type constraints: *all constraints are induced by the structure of the program and types*.

2.4.3 Extracting Bounds to Substitute. We write $Q = Q' \cdot \alpha=\tau$ to *extract* a non-dependent bound $\alpha=\tau$ from a prefix Q such that $Q = Q' \cup \{\alpha=\tau\}$ with $\alpha \notin \text{ftv}(Q', \tau)$:

$$\frac{Q = Q' \cup \{\alpha=\tau\} \quad \alpha \notin \text{ftv}(Q', \tau)}{Q = Q' \cdot \alpha=\tau} \text{EXTRACT}$$

This allows us to split a prefix Q into a bound $\alpha=\tau$ and a remaining prefix Q' that does not depend on α . Using extraction, we can now *discharge* prefix bounds with the [GENSUB] rule. This is similar to generalization in the [GEN] rule, except that we substitute the inferred monomorphic type bound on α . With [GENSUB]², we can finally derive the type of $\lambda x. \text{inc } x$ as:

$$\frac{\frac{\frac{\emptyset \mid \Gamma, x:\alpha \vdash \text{inc} : \text{int} \rightarrow \text{int} \quad \emptyset \mid \Gamma, x:\alpha \vdash x : \alpha}{\{\alpha=\text{int}, \beta=\text{int}\} \vdash \text{int} \rightarrow \text{int} \approx \alpha \rightarrow \beta \quad \text{fresh } \beta} \text{APP}}{\{\alpha=\text{int}, \beta=\text{int}\} \mid \Gamma, x:\alpha \vdash \text{inc } x : \beta} \text{GENSUB}}{\frac{\{\alpha=\text{int}\} \mid \Gamma, x:\alpha \vdash \text{inc } x : \text{int} \quad \text{fresh } \alpha}{\{\alpha=\text{int}\} \mid \Gamma \vdash \lambda x. \text{inc } x : \alpha \rightarrow \text{int}} \text{FUN}}{\emptyset \mid \Gamma \vdash \lambda x. \text{inc } x : \text{int} \rightarrow \text{int}} \text{GENSUB}$$

2.5 Principal Derivations

In the [LET] rule we find a single side condition: $\text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$. This is to ensure that any free type variables in σ that do not occur in Γ are generalized by [GEN] or [GENSUB]. Since there are no longer “guessed” types, this condition is enough to guarantee that all let-bindings get a most general type. Let’s rephrase the [LET] rule to use a separate judgement (\Vdash) for the type scheme:

$$\frac{Q_1 \mid \Gamma \Vdash e_1 : \sigma \quad Q_2 \mid \Gamma, x:\sigma \vdash e_2 : \tau}{Q_1, Q_2 \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \text{LET} \quad \frac{Q \mid \Gamma \vdash e : \sigma \quad \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{Q \mid \Gamma \Vdash e : \sigma} \text{MGEN}$$

²We could simplify the condition $\alpha \notin \text{ftv}(Q, \Gamma)$ in [GENSUB] to just $\alpha \notin \text{ftv}(\Gamma)$ since $\alpha \notin \text{ftv}(Q)$ is implied by $Q \cdot \alpha=\tau$.

Since the prefix bounds are minimal, and always induced by either the structure of the program (by [APP]), or by the structure of the types (by [EQ-VAR]), the types that can be derived by the [MGEN] rule are always unique:

Theorem 2.3. (*Principal type scheme derivations*)

If $Q_1 \mid \Gamma \Vdash e : \sigma_1$ then for any other derivation $Q_2 \mid \Gamma \Vdash e : \sigma_2$, we have $\sigma_1 = \sigma_2$.

We also have that any derived mono-type is unique up to prefix substitution (due to [GENSUB]):

Theorem 2.4. (*Principal type derivations*)

If $Q_1 \mid \Gamma \vdash e : \tau_1$ then for any other derivation $Q_2 \mid \Gamma \vdash e : \tau_2$, we have $Q_1[\tau_1] = Q_2[\tau_2]$.

The proofs are given in Appendix D.9. There we also show that we can also change any derivation to use a stronger condition on [MGEN] where we require $(\text{dom}(Q) \cup \text{ftv}(\sigma)) \subseteq \text{ftv}(\Gamma)$. This condition forces any trivial substitutions (for any $\alpha = \tau \in Q$ with $\alpha \notin \text{ftv}(\sigma)$), which leads a stronger principality Lemma D.36 where for any $Q_1 \mid \Gamma \Vdash e : \sigma_1$ and $Q_2 \mid \Gamma \Vdash e : \sigma_2$ we also have $Q_1 = Q_2$ (besides $\sigma_1 = \sigma_2$) (see Appendix D.9.4). Moreover, we can also show that the HMQ rules are sound and complete with respect to the standard HM type rules:

Theorem 2.5. (*Soundness*)

If $Q \mid \Gamma \vdash e : \sigma$ then we also have $Q[\Gamma] \vdash_{\text{HM}} e : Q[\sigma]$.

Theorem 2.6. (*Completeness*)

If $\Gamma \vdash_{\text{HM}} e : \sigma$, then there exists a θ such that $\theta\Gamma' \sqsubseteq \Gamma$, with $Q \mid \Gamma' \vdash e : \sigma'$, $Q \sqsubseteq \theta$, and $\theta\sigma' \sqsubseteq \sigma$.

As a corollary, we have that [MGEN] always derives most-general types, and also that algorithm W is a valid type inference algorithm for HMQ (and since W is also complete it infers the same types as HMQ derives). The soundness theorem states that if we can derive a type σ under a prefix Q in HMQ, then we can also derive the type $Q[\sigma]$ in HM (see Figure 8 in Appendix A for the definition of \vdash_{HM}). We need to apply the prefix to σ since it can still contain bounds (that could be applied with [GENSUB]).

The completeness theorem is more involved. We may have expected to see a simpler statement like: if $\Gamma \vdash_{\text{HM}} e : \sigma$, then $Q \mid \Gamma' \vdash e : \sigma'$ with $Q[\sigma'] \sqsubseteq \sigma$. That does not hold though since derivations may contain abstract types in our system. In particular, any lambda bound parameter always has an “abstract” type α and there may be no bound yet.

For example,

$$\begin{array}{c} \frac{x : \alpha \in (x : \alpha)}{\emptyset \mid x : \alpha \vdash x : \alpha} \text{VAR} \\ \frac{\emptyset \mid \emptyset \vdash \lambda x. x : \alpha \rightarrow \alpha}{\emptyset \mid \emptyset \vdash \lambda x. x : \forall \alpha. \alpha \rightarrow \alpha} \text{FUN} \\ \frac{}{\emptyset \mid \emptyset \vdash \lambda x. x : \forall \alpha. \alpha \rightarrow \alpha} \text{GEN} \end{array} \quad , \text{ but in the HM type rules, we can also derive: } \begin{array}{c} \frac{x : \text{int} \in (x : \text{int})}{x : \text{int} \vdash_{\text{HM}} x : \text{int}} \text{VAR} \\ \frac{x : \text{int} \vdash_{\text{HM}} x : \text{int}}{\emptyset \vdash_{\text{HM}} \lambda x. x : \text{int} \rightarrow \text{int}} \text{FUN} \end{array}$$

For an inductive proof, it means that for the [VAR] case, we would need to show that if we derive $x : \text{int} \vdash_{\text{HM}} x : \text{int}$, we can also derive $\emptyset \mid x : \alpha \vdash x : \alpha$ with $\emptyset[\alpha] \sqsubseteq \text{int}$ which does not hold.

Therefore, the actual completeness theorem states that if there exists some substitution θ with $\theta\Gamma' \sqsubseteq \Gamma$, with $Q \sqsubseteq \theta$ with $\theta\sigma' \sqsubseteq \sigma$. In our example, when we use $\theta = [\alpha := \text{int}]$, we indeed have $\emptyset \sqsubseteq [\alpha := \text{int}]$ and $[\alpha := \text{int}]\alpha \sqsubseteq \text{int}$. There is one more subtlety in that we need to use $\theta\sigma' \sqsubseteq \sigma$ and cannot use equality as $\theta\sigma' = \sigma$. In particular, in the HM type rules we can also introduce more sharing for let-bindings than we can in HMQ. For let $\text{const} = \lambda x. \lambda y. x$ we always infer $\text{const} : \forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \alpha$ in HMQ but under the HM rules we can also derive the type $\forall \alpha. \alpha \rightarrow \alpha \rightarrow \alpha$. In such a case, for the [VAR] rule we still need to show $\forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \alpha \sqsubseteq \forall \alpha. \alpha \rightarrow \alpha \rightarrow \alpha$. (and thus we need the instance relation \sqsubseteq). For the inductive proof, the full required completeness theorem is still a bit more general than stated here (see Appendix D.7 for details).

2.6 Idempotent Mappings

The reader may have noticed that $[\text{GENSUB}]$ may not always apply as we cannot always extract a binding α even if $\alpha \notin \text{ftv}(\Gamma)$. In particular, there might be multiple bounds for a type variable in the prefix, like $\{\alpha=\beta \rightarrow \text{int}, \alpha=\text{int} \rightarrow \gamma\}$, and in that case we cannot extract α directly for the $[\text{GENSUB}]$ rule (since $\alpha \in \text{ftv}(Q')$). However, for any consistent prefix, we can always simplify multiple bindings:

Theorem 2.7. (*Simplify*)

If $Q' \vdash \tau_1 \approx \tau_2$, then $Q \cup \{\alpha=\tau_1, \alpha=\tau_2\} \equiv Q \cup Q' \cup \{\alpha=\tau_1\}$

For example, $\{\alpha=\beta \rightarrow \text{int}, \alpha=\text{int} \rightarrow \gamma\} \equiv \{\beta=\text{int}, \gamma=\text{int}, \alpha=\beta \rightarrow \text{int}\}$. By using repeated simplification, we can always bring a consistent prefix into a form where all bindings are distinct (called a *mapping*).

Even with a mapping, there are still cases where we may have a dependency that prevents extraction. For example, consider extracting α from $\{\beta=\alpha, \alpha=\text{int} \rightarrow \text{int}\}$ (where $\alpha \in \text{ftv}(\{\beta=\alpha\})$). It turns out though that any consistent prefix is always equivalent to an *idempotent* mapping where $\text{dom}(Q) \not\cap \text{codom}(Q)$, e.g. $\{\beta=\alpha, \alpha=\text{int} \rightarrow \text{int}\} \equiv \{\beta=\text{int} \rightarrow \text{int}, \alpha=\text{int} \rightarrow \text{int}\}$.

Theorem 2.8. (*Any consistent prefix is equivalent to an idempotent mapping*)

If $\models Q$, then there exists an equivalent idempotent mapping Q' (where $Q \equiv Q'$, $|\text{dom}(Q')| = |Q'|$ and $\text{dom}(Q') \not\cap \text{codom}(Q')$).

This essentially allows us to always simplify a prefix enough to apply $[\text{GENSUB}]$ for any binding α in Q where $\alpha \notin \text{ftv}(\Gamma)$.

2.7 Flexible Bounds

The idea of inference under a prefix is inspired by the use of a prefix in the MLF type system [Le Botlan 2004; Le Botlan and Rémy 2003]. In MLF, the prefix does not just contain *rigid* bounds of the form $\alpha=\tau$, but also *flexible* bounds of the form $\alpha \geq \sigma$, which allows α to be any instance of σ . The flexible bound $\alpha \geq \perp$ allows α to be instantiated to any type. Finally, MLF uses quantification over a prefix, as in $\forall Q. \tau$ where $\forall \alpha. \sigma$ is a shorthand for $\forall \alpha \geq \perp. \sigma$. Moreover, rigid monomorphic bounds can be inlined, and $\forall \alpha=\tau. \sigma$ is equivalent to $[\alpha:=\tau]\sigma$. Using these richer bounds for a prefix, it is possible to use a single generalization rule instead of both $[\text{GEN}]$ and $[\text{GENSUB}]$. Let's extend our bounds to include $\alpha \geq \perp$ bounds as:

$$\alpha \diamond \rho ::= \alpha=\tau \mid \alpha \geq \perp \quad Q ::= \{ \alpha_1 \diamond \rho_1, \dots, \alpha_n \diamond \rho_n \}$$

We can then use a single generalization rule as:

$$\frac{Q \cdot (\alpha \diamond \rho) \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q, \Gamma)}{Q \mid \Gamma \vdash e : \forall(\alpha \diamond \rho). \sigma} \text{GENX}$$

This rule now concisely subsumes both $[\text{GEN}]$ and $[\text{GENSUB}]$ (and corresponds exactly to the $[\text{GEN}]$ rule of MLF [Le Botlan 2004, Fig. 5.2]). We would also extend simplification to merge flexible and rigid bounds where $Q \cup \{\alpha \geq \perp, \alpha=\tau\}$ simplifies to $Q \cup \{\alpha=\tau\}$.

We chose the current presentation in this paper for simplicity. Nevertheless, we believe that the use of an extended prefix with $\alpha \geq \perp$ bounds is perhaps more natural from a technical perspective and might also be better suited for example to extend HMQ to the MLF type rules.

2.8 Function Matching

The current $[\text{APP}]$ in Figure 2 has a drawback that it always creates a fresh type variable α for the result type. In practice, most implementations instead first match on the inferred type for e_1 to see if it is a function type $\tau' \rightarrow \tau$ already – and in that case directly use τ for the result type.

We can express this technique declaratively in HMQ as well. Figure 4 shows an improved $[\text{APP-MATCH}]$ rule that avoids creating a fresh result type variable by matching on the function type as $Q \vdash \tau_1 \approx \tau_2 \rightarrow \tau$, where τ_1 and τ_2 are given, and Q and the result type τ are synthesized. The (\approx) judgment has two rules. The $[\text{MFUN}]$ rule handles the case where it is already a function type

$$\begin{array}{c}
\boxed{
\begin{array}{c}
Q \vdash \tau \xrightarrow{\approx} \tau \rightarrow \tau \\
\downarrow \quad \uparrow \quad \uparrow \quad \downarrow \\
\text{out} \quad \text{in} \quad \text{in} \quad \text{out}
\end{array}
}
\end{array}
\quad
\frac{Q \vdash \tau_1 \approx \tau_2}{Q \vdash \tau_1 \rightarrow \tau \xrightarrow{\approx} \tau_2 \rightarrow \tau} \text{MFUN}
\quad
\frac{Q \vdash \alpha \approx \tau \rightarrow \beta \quad \text{fresh } \beta}{Q \vdash \alpha \xrightarrow{\approx} \tau \rightarrow \beta} \text{MVAR}$$

$$\frac{Q_1 \mid \Gamma \vdash e_1 : \tau_1 \quad Q_2 \mid \Gamma \vdash e_2 : \tau_2 \quad Q_3 \vdash \tau_1 \xrightarrow{\approx} \tau_2 \rightarrow \tau}{Q_1, Q_2, Q_3 \mid \Gamma \vdash e_1 e_2 : \tau} \text{APP-MATCH}$$

Fig. 4. Function matching.

and directly matches the expected parameter type with the inferred argument type. The only other possible case is that the type of e_1 is still an abstract α (for example, in $\lambda f. f \ 1$). The [MVAR] rule in that case applies and does create a fresh result type variable as before.

3 Implementing Inference under a Prefix

We believe the type rules in Figure 2 form a nice declarative specification of the type system where users can easily reason about what programs are accepted. At the same time though, it is possible to “read off” a type inference algorithm from the same rules. First we discuss how a direct implementation would look, and then consider a more standard implementation based on algorithm W.

3.1 Deriving a Direct Inference Algorithm

To directly derive an algorithm from the type rules, we first need to make the rules syntax-directed since the instantiation and generalization rules can be applied at any time. Following Damas and Milner [1982], we can make the rules syntax-directed by doing full instantiation at the leaves (in the [VAR] rule), and full generalization (with the [GEN] and [GENSUB] rules) at let-bindings. For example, the syntax directed rules for variables and let-bindings become:

$$\frac{x : \forall \bar{\alpha}. \tau \in \Gamma \quad \text{fresh } \bar{\alpha}}{\emptyset \mid \Gamma \vdash x : \tau}
\quad
\frac{Q_0 \mid \Gamma \vdash_s e_1 : \tau_1 \quad Q_2 \mid \Gamma, x : \sigma \vdash_s e_2 : \tau_2 \quad (Q_1, \sigma) = \text{gen}(Q_0, \Gamma, \tau_1)}{Q_1, Q_2 \mid \Gamma \vdash_s \text{let } x = e_1 \text{ in } e_2 : \tau_2}$$

where $\text{gen}(Q_0, \Gamma, \tau_1)$ generalizes a type τ_1 with respect to a given environment Γ and prefix Q_0 . See Figure 11 in Appendix B for the full syntax-directed rules. We can now almost implement each rule directly, except that we need to compute the consistent union between prefixes.

3.1.1 Computing the Prefix Solution. Any initial prefix at the leaves of a derivation is always either empty or a singleton $\{\alpha=\tau\}$ (with $\alpha \notin \text{ftv}(\tau)$). If we ensure that we always create an idempotent mapping from a consistent union Q_1, Q_2 then all our prefixes are always an idempotent mapping – and we can represent them in our implementation as regular substitutions; effectively representing Q as its minimal solution $\langle Q \rangle$. Using our notion of type equivalence, we can derive a straightforward algorithm to compute the prefix solution. In particular, we have that extraction corresponds to composition of prefix solutions:

Lemma 3.9. (*Extraction corresponds to composition of prefix solutions*)

If $\models Q$ and $Q = Q' \cdot \alpha=\tau$, then $\langle Q \rangle = \langle Q' \rangle \circ [\alpha:=\tau]$.

Using this lemma, we can write the initial cases of our algorithm as:

$$\text{solve}(\emptyset) = \text{id} \quad \text{and} \quad \text{solve}(Q \cup \{\alpha=\tau\}) = \text{solve}(Q) \circ [\alpha:=\tau] \quad \text{if } \alpha \notin \text{ftv}(Q, \tau)$$

If we cannot find any α with $\alpha \notin \text{ftv}(Q, \tau)$, that leaves two other cases to consider. If $\alpha \in \text{ftv}(\tau)$ or $\alpha \in \text{ftv}(\text{rng}(Q))$ there must be cyclic dependency and there is no solution. Otherwise, there must be duplicate bindings (with $\alpha \in \text{dom}(Q)$), and in such a case we can use Theorem 2.7 to simplify

the duplicate bindings³:

$solve : Q \rightarrow \theta$

$solve(\emptyset) = id$

$solve(Q \cup \{\alpha=\tau\}) = solve(Q) \circ [\alpha:=\tau] \quad \text{if } \alpha \notin \text{ftv}(Q, \tau)$

$solve(Q \cup \{\alpha=\tau_1, \alpha=\tau_2\}) = solve(Q \cup Q' \cup \{\alpha=\tau_1\}) \quad \text{if } Q' \vdash \tau_1 \approx \tau_2 \wedge \alpha \notin \text{ftv}(\tau_1, \tau_2, \text{rng}(Q))$

Essentially this algorithm picks non-dependent bindings and composes them recursively, while simplifying duplicate bindings away by unifying their types using the equivalence relation. But how can we compute $Q' \vdash \tau_1 \approx \tau_2$? To derive an implementation for the equivalence relation we need to make these syntax-directed as well. Similar to instantiation we can always apply [EQ-REFL] at the leaves of the derivation at the [EQ-VAR] rule and make them syntax-directed. We can then derive an implementation, representing prefixes again as substitutions, as:

$equiv(\alpha, \alpha) = id$

$equiv(\alpha, \tau) \text{ or } (\tau, \alpha) \mid \alpha \notin \text{ftv}(\tau) = [\alpha:=\tau]$

$equiv(\tau_1 \rightarrow \tau_2, \tau'_1 \rightarrow \tau'_2) = \text{let } \theta_1 = equiv(\tau_1, \tau'_1); \theta_2 = equiv(\tau_2, \tau'_2) \text{ in } solve(\theta_1 \cup \theta_2)$

This is recursive with $solve$ but we can show it is terminating since the number of free type variables is decreasing on each recursive invocation [Pierce 2002, §22.4.5].

Since we happen to represent the prefixes as a substitutions in our implementation, we can now compute prefix composition as $Q_1, Q_2 = solve(Q_1 \cup Q_2)$, and directly “read off” an inference algorithm from our type rules. For example, the inference case for the [APP_s] application rule becomes (using substitutions θ for idempotent mapping prefixes Q):

$inferD(\Gamma, e_1 e_2) : (\Gamma, e) \rightarrow (\theta, \tau) =$

let $(\theta_1, \tau_1) = inferD(\Gamma, e_1); (\theta_2, \tau_2) = inferD(\Gamma, e_2)$

let $\alpha = \text{fresh}; \theta_3 = equiv(\tau_1, \tau_2 \rightarrow \alpha)$

let $\theta = solve(solve(\theta_1 \cup \theta_2) \cup \theta_3)$

(θ, α)

3.1.2 Robinson Unification and Substitution Unification. Of course, we can also readily use standard Robinson unification [Robinson 1965] to compute $\langle Q \rangle$ as well. In particular, if we view Q as a set of type constraints C with constraints of the form $\tau_1 = \tau_2$, we can use the standard $\text{unify}(C)$ algorithm from Pierce [2002, §22.4] to compute the most general unifier of the constraints in Q – which is $\langle Q \rangle$ by definition (and therefore $solve(Q) = \text{unify}(Q)$). An approach that maps more directly to the idea of joining prefixes is the work by McAdam [1999] – describing an algorithm U_s for unifying substitutions which can be used directly to implement joining (idempotent mapping) prefixes where $Q_1, Q_2 \equiv U_s(Q_1, Q_2) \circ Q_1$ (and thus $solve(Q_1 \cup Q_2) = U_s(Q_1, Q_2) \circ Q_1$)

Nevertheless, we prefer $solve$ as that is parameterized by our type equivalence rules, $Q \vdash \tau_1 \approx \tau_2$, to determine equivalent types and least solutions. In contrast, the type equivalence is “built-in” in the unify and U_s algorithms. We believe that our approach lends itself better to type system extensions, like record types or impredicative types, where the equality between types can go beyond syntactical equality. In such cases, it is straightforward to extend our type equivalence relation with further rules.

3.2 Algorithm W

Even though we can implement the syntax-directed rules directly with $inferD$, it may not be the most efficient way to do this. However, since HMQ is sound and complete with respect to the HM type rules, we can also directly use algorithm W for HMQ inference *as-is*! That means also that any existing efficient implementation of HM, for example using in-place updating substitutions [Peyton Jones et al. 2007] or level-based generalization [Kiselyov 2022; Kuan and MacQueen 2007;

³The extra side condition $\alpha \notin \text{ftv}(\tau_1, \tau_2, \text{rng}(Q))$ is needed here to ensure that $solve$ terminates for any inconsistent Q with cyclic bindings – consider for example $solve(\{\beta=\alpha, \alpha=\beta, \alpha=int\})$.

Rémy 1992], is correct for HMQ as well.

There is a catch though – even though algorithm W is correct for the basic type rules of HMQ, it may not be correct anymore for some extensions and we need to be a bit more careful. In particular, in an application for example, algorithm W uses a single substitution that is threaded linearly through each sub derivation, resolving unification constraints eagerly. This linear traversal is what allows for an efficient in-place updating implementation of substitutions. For example, the case for applications in algorithm W is [Damas and Milner 1982]:

$$\begin{aligned} \text{inferW}(\Gamma, e_1 e_2) : (\Gamma, e) \rightarrow (\theta, \tau) = \\ \text{let } (\theta_1, \tau_1) = \text{inferW}(\Gamma, e_1); (\theta_2, \tau_2) = \text{inferW}(\theta_1 \Gamma, e_2) \\ \text{let } \alpha = \text{fresh}; \theta_3 = \text{unify}(\theta_2 \tau_1, \tau_2 \rightarrow \alpha) \\ (\theta_3 \circ \theta_2 \circ \theta_1, \theta_3 \alpha) \end{aligned}$$

where we see that the substitution θ_1 is applied to Γ when checking e_2 (see Figure 10 in Appendix A.2 for the full algorithm).

For the basic type rules this makes no difference, but if we were to inspect the types of λ -bound parameters the implementations start to differ. In algorithm W, since the substitution is updated eagerly, type information from an early variable occurrence may “leak” into another sub derivation – we call this spooky action at a distance. Consider for example

$\lambda x. (\text{inc } x, \text{show } x)$

At the first occurrence of x the type will be some fresh type α , and after checking the $\text{inc } x$ expression, we’ll have a substitution $[\alpha := \text{int}]$. When this substitution is propagated into the second derivation, the second occurrence of x in $\text{show } x$ now has the substituted type int in algorithm W! For static overloading (described in Section 6) this would mean that $\text{show } x$ can be resolved while it should be rejected according to the HMQ type rules where x always has an abstract type (and worse, it leaks the left-to-right bias of algorithm W, where $\lambda x. (\text{show } x, \text{inc } x)$ would be rejected).

3.3 Algorithm WQ

It turns out that a small change to algorithm W can prevent spooky action at a distance, and prevent leaking type information between separate sub derivations even when using efficient stateful substitutions. The core issue is that in algorithm W the fresh type variable for a λ -bound parameter is shared between sub derivations. What we can do instead is to use a *separate fresh type variables for each occurrence* and unify them all eventually. In particular, in an expression $\lambda x. e$ we rename all occurrences of a λ -bound parameter x in e sequentially to x_i (as e'), and rewrite the entire expression to $\lambda x. (\lambda x_1 \dots x_n. e') x \dots x$. After such transformation, algorithm W will effectively create a fresh type variable for each x_i and only unify them all afterwards (with the shared type of x) – thus preventing spooky action at a distance.

We can do this more efficiently by integrating this strategy directly in a revised algorithm WQ, where we generate a fresh “template” type variable α for a parameter, but expand that to a unique type variable α_i at each occurrence of a λ -bound parameter:

$$\begin{aligned} \text{inferWQ}(\Gamma, x_i) = & \text{inferWQ}(\Gamma, \lambda x. e) = \\ \text{let } \alpha = \Gamma(x) & \text{let } \alpha = \text{fresh} \\ (id, \alpha_i) & \text{let } (\theta_1, \tau) = \text{inferWQ}((\Gamma, x : \alpha), e) \\ & \text{let } \theta_2 = \text{unifies}(\alpha, \theta_1 \alpha_1, \dots, \theta_1 \alpha_n) \\ & (\theta_2, \theta_2(\alpha \rightarrow \tau)) \end{aligned}$$

The full algorithm WQ can be found in Figure 12 in Appendix B.1. For our standard rules, this change to algorithm W only changes when certain type errors happen (which are now sometimes delayed). However, when adding type propagation (Section 5) and overloading (Section 6), the new algorithm WQ prevents spooky action at a distance, and type information in separate sub derivations can no longer be accidentally shared.

$$\begin{array}{c}
\frac{}{\emptyset \vdash \tau \approx \tau} \text{EQF-ID} \qquad \frac{Q \vdash \sigma_2 \approx \sigma_1}{Q \vdash \sigma_1 \approx \sigma_2} \text{EQF-REFL} \qquad \frac{Q \vdash \sigma_1 \approx \sigma_2 \quad \alpha \notin \text{ftv}(Q)}{Q \vdash \forall \alpha. \sigma_1 \approx \forall \alpha. \sigma_2} \text{EQF-POLY} \\
\frac{\alpha \notin \text{ftv}(\sigma)}{\{\alpha = \sigma\} \vdash \alpha \approx \sigma} \text{EQF-VAR} \qquad \frac{Q \vdash \sigma_1 \approx \sigma_2}{Q \vdash [\sigma_1] \approx [\sigma_2]} \text{EQF-LIST} \qquad \frac{Q_1 \vdash \sigma_1 \approx \sigma_2 \quad Q_2 \vdash \sigma'_1 \approx \sigma'_2}{Q_1, Q_2 \vdash \sigma_1 \rightarrow \sigma_2 \approx \sigma'_1 \rightarrow \sigma'_2} \text{EQF-FUN}
\end{array}$$

Fig. 5. Equivalence of System F types under a prefix.

As such, algorithm WQ is a modest extension to algorithm W and we believe that it is straightforward to adapt for type system implementations in practice, while simultaneously benefitting from being able to use HMQ to specify the type rules and further extensions in a concise manner that matches the implementation closely.

4 Inference under a Prefix for FreezeML and HMF

We believe HMQ can be an excellent basis to describe common type system extensions in practice that are difficult to formalize directly in the HM type rules. In this section we look at some of the previous work on higher-rank and impredicative type inference, and consider how these systems could be viewed in terms of inference under a prefix. We consider in particular the recent FreezeML system [Emrich et al. 2020 2022] and HMF [Leijen 2007 2008]. Note that for the purposes of this article we restrict ourselves to highlight essential differences only – the goal of this section is to show how inference under a prefix may be a better way to formalize and compare such systems, and it is not meant as a general introduction to impredicative type inference.

Generally, these systems allow for higher-rank (i.e. nested quantifiers) and impredicative types (i.e. polymorphic types in a data structure), and extend the syntax of types essentially with:

$$\begin{array}{ll}
\sigma ::= \forall \alpha. \sigma & (\text{quantification}) \\
| \rho & (\text{no outer quantifier})
\end{array}
\qquad
\begin{array}{ll}
\rho ::= \sigma \rightarrow \sigma & (\text{higher-rank function}) \\
| \tau & (\text{monomorphic types}) \\
| [\sigma] & ((\text{impredicative}) \text{ list of } \sigma)
\end{array}$$

where we restrict ourselves to impredicative lists for example purposes. The instance relation can now instantiate polymorphic types as well:

$$\frac{\bar{\beta} \notin \text{ftv}(\forall \bar{\alpha}. \sigma_1)}{\forall \bar{\alpha}. \sigma_1 \sqsubseteq \forall \bar{\beta}. [\bar{\alpha} := \bar{\sigma}] \sigma_1} \text{INSTANCEF}$$

Generally, impredicative systems are *invariant* where we can only instantiate the outer quantifiers (as in Damas-Hindley-Milner), but not any inner quantifiers. For example, we can instantiate the identity function as $\forall \alpha. \alpha \rightarrow \alpha \sqsubseteq \text{int} \rightarrow \text{int}$, but we cannot instantiate a list of polymorphic identity functions as $[\forall \alpha. \alpha \rightarrow \alpha] \not\sqsubseteq [\text{int} \rightarrow \text{int}]$. In HMQ this shows up clearly when defining new type equivalence rules that take impredicative types into consideration as shown in Figure 5.

Note that we extended the prefix to include (rigid) polymorphic bounds $\alpha = \sigma$. Also, to prevent the bound α in the [EQF-POLY] rule from escaping into Q , we need the side-condition $\alpha \notin \text{ftv}(Q)$. The new equivalence rules again closely resemble common unification algorithms for impredicative types [Emrich et al. 2020, Fig.15; Leijen 2008, Fig.5]. To implement the [EQF-POLY] rule one usually instantiates both outer quantifiers with a fresh constant (often called a “skolem” constant) and afterwards check that the constant does not escape into Q .

There are two troublesome cases to consider with impredicative type inference. Generally, we cannot infer polymorphic types for lambda-bound parameters. Consider $\text{poly} = \lambda f. (f \ 1, f \ \text{True})$ – such definition would be rejected in HM systems since there is no monomorphic type for f that can be applied to both an *int* and a *bool*. We could assign a polymorphic type to f though – like $\forall \alpha. \alpha \rightarrow \alpha$. Unfortunately, there is no principal type for f and there are many other incomparable types possible (like $\forall \alpha. \alpha \rightarrow [\alpha]$ etc). The systems we discuss therefore never infer a polymorphic

type for a lambda-bound parameter and require a type annotation for polymorphic parameters.

The second issue occurs at applications where there is sometimes a choice between instantiations. Consider the application *single id* where *single* has type $\forall\alpha.\alpha \rightarrow [\alpha]$. If we instantiate *id* first, the result type is $\forall\alpha.[\alpha \rightarrow \alpha]$ after generalization – but if we keep *id* polymorphic, the result type is a list of polymorphic identity functions $[\forall\alpha.\alpha \rightarrow \alpha]$ instead. Unfortunately, neither type is an instance of the other. Generally, different proposed systems handle this case in very different ways. Here, we take a closer look at the FreezeML and HMF systems specifically, and consider how they could be simplified by using prefix based inference.

4.1 FreezeML

FreezeML [Emrich et al. 2020 2022] is an impredicative type inference system based on the idea of *freezing* the polymorphic type of a variable occurrence, written as $[x]$, and only allowing instantiation at regular variable occurrences x . Alas, that also means FreezeML is fundamentally syntax directed and we need to base a “FreezeHMQ” version on the syntax directed rules of HMQ (see Fig. 11 in Appendix B), where we instantiate at variable occurrences, and generalize ([GEN] and [GENSUB]) at let-bindings. The freezing rule of FreezeML becomes:

$$\frac{x:\sigma \in \Gamma}{\emptyset \mid \Gamma \vdash_s [x] : \sigma} \text{FREEZE}$$

Since we no longer can instantiate freely, a frozen type σ stays polymorphic while regular variable occurrences have instantiated ρ types. This resolves the *single id* ambiguity: *single id* has the HM type $\forall\alpha.[\alpha \rightarrow \alpha]$ since *id* is fully instantiated. If we wish to create a list of polymorphic identity functions we would write *single [id]* instead (which has type $[\forall\alpha.\alpha \rightarrow \alpha]$).

Even though FreezeML is syntax directed, it still requires let-bindings to have most-general types. Consider let $f = \lambda x.x$ in $[f]$ 42. We expect this to be rejected with the frozen type for f as $\forall\alpha.\alpha \rightarrow \alpha$ (which cannot be directly applied). However, if we allow less-general types for let-bindings we can also derive the type $\text{int} \rightarrow \text{int}$ for f , and in that case it *can* be typed. To resolve this, the [LET] rule in FreezeML adds the principal condition [Emrich et al. 2020, Fig. 7&8]:

$$\frac{\begin{array}{l} (_, \Delta') = \text{gen}(\Delta, \sigma_1, e) \quad (\Delta, \Delta', e, \sigma_1) \Downarrow \sigma \\ \Delta, \Delta' \mid \Gamma \vdash e_1 : \sigma_1 \quad \Delta \mid \Gamma, x:\sigma \vdash e_2 : \sigma_2 \\ \text{principal}(\Delta, \Gamma, e, \Delta', \sigma_1) \end{array}}{\Delta \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \sigma} \text{LET-FML}$$

$$\begin{array}{l} \text{principal}(\Delta, \Gamma, e, \Delta', \sigma') = \\ \Delta' = \text{ftv}(\sigma) - \Delta \text{ and } \Delta, \Delta' \mid \Gamma \vdash e : \sigma \text{ and} \\ (\forall \Delta'', \sigma''. \text{ if } \Delta'' = \text{ftv}(\sigma'') - \Delta \\ \text{and } \Delta, \Delta'' \mid \Gamma \vdash e : \sigma'' \text{ then} \\ \exists \delta. \Delta \vdash \delta : \Delta' \Rightarrow_\star \Delta'' \text{ and } \delta(\sigma') = \sigma'') \end{array}$$

We can disregard the \Downarrow rule as that is related to the value-restriction, and we can similarly ignore the Δ environment that tracks the free variables. The principal condition enforces that all let bindings are assigned most general types. Since it ranges over all possible derivations where the type inference judgment occurs negatively, it is not a natural deduction rule which makes it hard to reason about. Emrich et al [2020, §3.2] show though that it is still possible to stratify the relation to allow inductive reasoning.

However, in “FreezeHMQ” none of this complexity is required as we already always derive principal types, and we can keep the regular (syntax-directed) [LET] rule as is. We only need to extend the types and type equivalence as shown in the previous section together with [FUN-ANN] and [FREEZE] to model FreezeML. This also shows that an implementation of type inference for FreezeML only requires a modest extension of algorithm W – essentially just extending unification according to the rules in Figure 5 (as in [Emrich et al. 2020, Fig. 15]).

4.2 HMF

As another example, we take a close look at the HMF system [Leijen 2008], which is used for impredicative type inference in the Koka language [Leijen 2014 2021]. Unlike FreezeML, the HMF

rules are not required to be syntax-directed and one can freely instantiate and generalize. The HMF system, however, contains two inference rules with complex side-conditions:

$$\begin{array}{c}
\frac{\Gamma \vdash e_1 : \sigma_1 \quad \Gamma, x : \sigma_1 \vdash e_2 : \sigma_2 \quad \forall \sigma'_1. \Gamma \vdash e_1 : \sigma'_1 \Rightarrow \sigma_1 \sqsubseteq \sigma'_1}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2} \text{HMF-LET} \quad \frac{\Gamma \vdash e_1 : \sigma_2 \rightarrow \sigma \quad \Gamma \vdash e_2 : \sigma_2 \quad \forall \sigma' \sigma'_2. (\Gamma \vdash e_1 : \sigma'_2 \rightarrow \sigma' \wedge \Gamma \vdash e_2 : \sigma'_2) \Rightarrow \llbracket \sigma_2 \rightarrow \sigma \rrbracket \leq \llbracket \sigma'_2 \rightarrow \sigma' \rrbracket}{\Gamma \vdash e_1 e_2 : \sigma} \text{HMF-APP}
\end{array}$$

As before, the [HMF-LET] rule requires that we only assign most general types to let-bindings – but this comes for free in HMQ and we can again use our regular [LET] rule as is. In the [HMF-APP] rule, the condition requires that the inferred type must be the one with a minimal *polymorphic weight* (denoted as $\llbracket \sigma \rrbracket$), where the polymorphic weight of a type is defined as the number of nested quantifiers. This is how HMF disambiguates the *single id* example, which has the type $\forall \alpha. [\alpha \rightarrow \alpha]$ since that is the one with minimal nested polymorphism (and if a list of polymorphic identity functions is required one needs to use a type signature).

Just as in the [LET-HMF] rule, the side condition is stated over all derivations again – in this case this is needed as a polymorphic instantiation can be further up in the derivation. This issue is already avoided though in HMQ since the [INST] rule never guesses types and always instantiates with an abstract type variable. As a consequence, it is possible to locally extend the function matching in the application rule to explicitly disambiguate.

Leijen [2008] observes that the only ambiguity can arise when a function of the form $\alpha \rightarrow \dots$ is applied to a polymorphic argument σ . In such a case we need to instantiate the σ and not unify directly with α . We can extend the function match relation in Figure 4 to do this disambiguation:

$$\begin{array}{c}
\boxed{\begin{array}{c} Q \vdash \rho \approx \sigma \rightarrow \sigma \\ \downarrow \quad \uparrow \quad \uparrow \quad \downarrow \\ \text{out} \quad \text{in} \quad \text{in} \quad \text{out} \end{array}} \quad \frac{Q_1 \mid \Gamma \vdash e_1 : \rho \quad Q_2 \mid \Gamma \vdash e_2 : \sigma_2 \quad Q_3 \vdash \rho \approx \sigma_2 \rightarrow \sigma}{Q_1, Q_2, Q_3 \mid \Gamma \vdash e_1 e_2 : \sigma} \text{APP-HMF-MATCH} \\
\\
\frac{Q \vdash \rho_1 \approx \rho_2}{Q \vdash \rho_1 \rightarrow \sigma \approx \rho_2 \rightarrow \sigma} \text{MFUN} \quad \frac{Q \vdash \alpha \approx \rho \rightarrow \beta \quad \text{fresh } \beta}{Q \vdash \alpha \approx \rho \rightarrow \beta} \text{MVAR} \quad \frac{Q \vdash \sigma_1 \approx \sigma_2 \quad \sigma_1 \notin \rho}{Q \vdash \sigma_1 \rightarrow \sigma \approx \sigma_2 \rightarrow \sigma} \text{MQUNT}
\end{array}$$

The [MFUN] and [MVAR] rules are as before but extended to apply to impredicative ρ types. The [MQUNT] rule is added and matches an actual polymorphic parameter type σ_1 (where σ_1 cannot be an unquantified ρ -type). This ensures that in the *single id* case, we must use the regular [MFUN] rule which forces the argument type to be instantiated (as ρ_2) (and thus *single id* has type $\forall \alpha. [\alpha \rightarrow \alpha]$).

The new function match, together with the new type equivalence as shown in the previous section are the only changes needed to phrase HMF as inference under a prefix! This again also implies that only a modest extension to algorithm W is required to implement HMF under a prefix: indeed, the subsume and funmatch implementations as shown in the original HMF paper [Leijen 2008, Fig. 6&8] closely match the function match rules that we show here. As argued in the introduction, if we consider the complex polymorphic weight condition in the [APP-HMF] rule, it can be considered somewhat artificial and quite far removed from the relatively straightforward implementation based on local matching.

We believe that stating both FreezeML and HMF using common prefix inference rules also makes the relation between the two more clear – HMF disambiguates instantiation at applications by inspecting the expected parameter type, while FreezeML disambiguates syntactically at variable occurrences relying on syntax directed rules.

5 Bidirectional Inference under a Prefix

Almost all type inference systems in practice use a form of bidirectional type inference [Odersky et al. 2001; Pierce and Turner 2000] where type information is not only inferred, but also propagated up to the leaves of a derivation. One advantage is to improve type error messages, but often it is used

$$\boxed{
\begin{array}{c}
Q \mid \Gamma \vdash e \overset{\leftarrow}{:} \sigma \\
\downarrow \quad \uparrow \quad \uparrow \quad \uparrow \\
\text{out} \quad \text{in} \quad \text{in} \quad \text{in}
\end{array}
\text{ with } \models Q$$

$$\begin{array}{c}
\frac{Q \mid \Gamma \vdash e \overset{\leftarrow}{:} \sigma}{Q \mid \Gamma \vdash (e : \sigma) : \sigma} \text{ANN} \qquad \frac{Q_1 \mid \Gamma \vdash e : \tau_1 \quad Q_2 \vdash \tau_1 \approx \tau}{Q_1, Q_2 \mid \Gamma \vdash e \overset{\leftarrow}{:} \tau} \text{CHK} \\
\\
\frac{Q \mid \Gamma, x : \tau_1 \vdash e \overset{\leftarrow}{:} \tau_2}{Q \mid \Gamma \vdash \lambda x. e \overset{\leftarrow}{:} \tau_1 \rightarrow \tau_2} \text{FUNC} \qquad \frac{Q \mid \Gamma \vdash e \overset{\leftarrow}{:} \sigma \quad \alpha \notin \text{ftv}(Q, \Gamma, e)}{Q \mid \Gamma \vdash e \overset{\leftarrow}{:} \forall \alpha. \sigma} \text{GENC} \\
\\
\frac{Q_1 \mid \Gamma \vdash e_1 \overset{\leftarrow}{:} Q_2[\tau_2] \rightarrow \tau \quad Q_2 \mid \Gamma \vdash e_2 : \tau_2}{Q_1, Q_2 \mid \Gamma \vdash e_1 e_2 \overset{\leftarrow}{:} \tau} \text{APP-FUNC} \qquad \frac{Q \mid \Gamma \vdash e_1 e_2 \overset{\leftarrow}{:} \alpha \quad \text{fresh } \alpha}{Q \mid \Gamma \vdash e_1 e_2 : \alpha} \text{APP-CHK} \\
\\
\frac{Q_1 \mid \Gamma \vdash e_1 \overset{\leftarrow}{:} \beta \rightarrow \tau \quad Q_2 \mid \Gamma \vdash e_2 \overset{\leftarrow}{:} Q_1[\beta] \quad \text{fresh } \beta}{Q_1, Q_2 \mid \Gamma \vdash e_1 e_2 \overset{\leftarrow}{:} \tau} \text{APP-ARGC}
\end{array}$$

Fig. 6. Bidirectional type checking rules

to enable type system extensions. For example, this technique can be used to check higher-ranked types [Odersky and Läufer 1996; Peyton Jones et al. 2007]. It is straightforward to add bidirectional type rules to inference under a prefix as well as shown in Figure 6.

The checking judgement $Q \mid \Gamma \vdash e \overset{\leftarrow}{:} \sigma$ states that an expression e can be *checked* to have (the input) type σ under a given environment Γ and (output) prefix Q . The [ANN] rule switches from inference mode to checking mode with a given type annotation σ . Dually, we can always apply the [CHK] rule to switch from checking mode to inference mode where we use the type equivalence relation to ensure the inferred type τ_1 matches the checked type τ_2 . The [FUNC] rule splits a checked function type to bind the parameter type directly and propagate the result type to the body. The rule [GENC] instantiates propagated polymorphic types.

For checking applications $e_1 e_2$ there is a choice: we can either first infer the type of the argument and use that to check the function type ([APP-FUNC]), or we can first infer the type of the function and use that to check the type of the argument ([APP-ARGC]). The rule [APP-FUNC] is straightforward and just propagates the inferred type of the argument τ_2 into the function type. For [APP-ARGC] we use a fresh type β as a place holder for the argument type, and check if e_1 is a function $\beta \rightarrow \tau$. Here, we propagate just the information that e_1 must be a function with result type τ where we use β to be able to refer to the (inferred!) expected type of the argument. We propagate this type to check the argument as $Q_1[\beta]$. This is somewhat similar to boxy type inference [Vytiniotis et al. 2006] where one would check the function type as $\boxed{\tau_2} \rightarrow \tau$ where the boxed τ_2 represents inferred type information that cannot be used for checking – in our prefix based system such boxes are handled by abstract (fresh) type variables.

The new checking rules for applications can now be used to replace the inference rule for application with [APP-CHK] where we just propagate a fresh result type α . We have now neatly separated out different parts of the original [APP] rule: the creation of a fresh result type α in [APP-CHK], the inference of the argument in [APP-FUNC], and finally the equivalence of the function type τ_1 to $\tau_2 \rightarrow \alpha$ using [CHK] (combined with the use of the checking judgment in [APP-FUNC]).

The application checking rules are not syntax-directed though – which rule should we apply in practice? This choice is not so clear cut [Dunfield and Krishnaswami 2021]; usually it is considered best to use [APP-ARGC] to propagate type information into the argument expression [Peyton Jones et al. 2007; Pierce and Turner 2000] but this is not always the case, and it depends on intended usage (and we discuss this in more detail in Section 6.3). In particular, at the moment our checked

type rules do not *do* anything, and just propagate known type information. At this point, these rules can only improve type error messages in practice. In Section 6 though we look at a checking rule for variables that actually takes the propagated type information into account.

6 Static Overloading

After reconsidering existing systems like FreezeML and HMF in Section 4 in terms of inference under a prefix, we now take a look at a novel application where we rely on prefixes to disambiguate variables for *static overloading*. For example, we would like to write $\lambda x\ y.(x + 1, y + 1.0)$ and have the $(+)$ operations resolve to integer- and floating point addition respectively. One elegant solution to overloading is the use of type classes [Wadler and Blott 1989]. Even though type classes are very expressive and highly succesful in languages like Haskell and Lean, they are also a complex extension that changes the semantics of types, and require sophisticated constraint solving of type instance relations [Selsam et al. 2020; Vytiniotis et al. 2010] with many possible design choices from a language perspective [Jones and Diatchki 2008; Peyton Jones et al. 1997].

6.1 Overloading as Disambiguation

Instead, we consider a much simpler alternative here, and look at the most basic form of overloading where we only disambiguate statically between different known versions of an overloaded function f based on the local type context. This form of static overloading is quite common and for example used in the C language to overload various arithmetic operations to work over integers and floats.

For our purposes, we allow a function f to be defined with a qualified name, like `modi/f`, which allows multiple definitions for f in different modules or namespaces. For example, we could have:

```
modi/show : int → string = ...
modb/show : bool → string = ...
```

Generally, such qualified names can come from definitions in different imported modules, but we may also directly allow programmers to use qualified names when defining functions (as if the function is defined inside a mini-module). Note that these kinds of qualified names already occur naturally in any language with namespaces or modules, and languages already need some mechanism to deal with ambiguity: if one imports module `modi` and `modb` that both export the `show` definition, to which definition should an unqualified `show` refer to? In Haskell for example, one needs to use a fully qualified name to disambiguate.

Another advantage of using qualified names is that it does not require an upfront declaration of the variables which can be overloaded, and we can always refer directly to each definition by explicitly using their unique fully qualified name. As such we can view static overloading as a source-to-source translation that only disambiguates identifiers to their fully qualified name.

The idea is now to use static type information at a call site to allow a programmer to write an unqualified name, like `show`, and have it be disambiguated automatically to the full qualified name depending on the type context. For example, `show 1` is disambiguated to `modi/show 1` since it is used with an argument of type `int`. In contrast to type classes, static overloading rejects programs where a variable cannot be disambiguated uniquely, like $\lambda x. \text{show } x$ for example. This is of course a severe restriction as it prohibits abstraction over overloaded variables. However, Following Lewis et al. [2000], we believe such an abstraction should be a separate and orthogonal concept, where we use implicit parameters in combination with static overloading. We come back to this at the end of this Section where we discuss the implementation of static overloading in Koka language.

Even in this restricted form, static overloading can be quite useful in practice as it handles many common cases of first-order overloading. However, even though the idea is simple, it clearly does not work well with standard HM inference. If we consider $\lambda x. \text{show } x$ again, we can “guess” the type `int` for the lambda-bound x parameter, and in that case we can accept the expression and disambiguate to `modi/show` – or guess type `bool` and elaborate to `modb/show` instead. Similarly, if we allow non-principal types for let-bindings we can also derive different disambiguations.

6.2 Bidirectional Disambiguation

If we use inference under a prefix though, we avoid all these problems since parameter types are no longer guessed, and let-bindings have a principal type by construction. For example, the expression $\lambda x. \text{show } x$ is always rejected now since we cannot disambiguate on the abstract type variable that is assigned to x . It turns out that extending HMQ with static overloading is quite straightforward where we mainly need to extend the bidirectional rules of Figure 6 with a case for variables:

$$\frac{\text{unique } m/x : \sigma \in \Gamma \text{ with } Q \vdash \sigma \sqsubseteq \tau}{Q \mid \Gamma \vdash x \overset{\tau}{\dashv} \tau \rightsquigarrow m/x} \text{VARC} \qquad \frac{Q \vdash \tau_1 \approx \tau_2 \quad \text{fresh } \bar{\alpha}}{Q \vdash \forall \bar{\alpha}. \tau_1 \sqsubseteq \tau_2} \text{INSTANCEC}$$

We use unique notation in the [VARC] rule to mean: “there exists a unique $m/x : \sigma \in \Gamma$, where σ can be instantiated to τ (with $Q \vdash \sigma \sqsubseteq \tau$), and for all other $m'/x : \sigma' \in \Gamma$ with $m \neq m'$, $Q' \vdash \sigma' \sqsubseteq \tau$ does not hold”. The [INSTANCEC] rule checks if a type $\forall \bar{\alpha}. \tau_1$ can be instantiated to a given type τ_2 . This can be done directly by using the equivalence relation with fresh types for $\bar{\alpha}$ (as in the [INST] rule). The bidirectional type rules provide the type information τ required to disambiguate overloaded variables. For example, we can derive the type of *show 1* as:

$$\frac{\frac{\text{unique } \text{modi/show} : \text{int} \rightarrow \text{string} \in \Gamma \quad \text{with } \{\alpha = \text{string}\} \vdash \text{int} \rightarrow \text{string} \sqsubseteq \text{int} \rightarrow \alpha}{\{\alpha = \text{string}\} \mid \Gamma \vdash \text{show} \overset{\tau}{\dashv} \text{int} \rightarrow \alpha \rightsquigarrow \text{modi/show}} \text{VARC} \quad \frac{}{\emptyset \mid \Gamma \vdash 1 : \text{int}} \text{INT}}{\frac{\{\alpha = \text{string}\} \mid \Gamma \vdash \text{show } 1 \overset{\tau}{\dashv} \alpha \quad \text{fresh } \alpha}{\{\alpha = \text{string}\} \mid \Gamma \vdash \text{show } 1 : \alpha} \text{APP-FUNC}} \text{APP-FUNC}$$

$$\frac{\frac{\{\alpha = \text{string}\} \mid \Gamma \vdash \text{show } 1 \overset{\tau}{\dashv} \alpha \quad \text{fresh } \alpha}{\{\alpha = \text{string}\} \mid \Gamma \vdash \text{show } 1 : \alpha} \text{APPC}}{\emptyset \mid \Gamma \vdash \text{show } 1 : \text{string}} \text{GENSUB}$$

Extending HMQ with static overloading, essentially as just disambiguation over qualified names, is as simple as shown here – and the [VARC] and [INSTANCEC] rules are also straightforward to implement. However, to ensure that all overloaded variables are always resolved uniquely we need to use syntax-directed bidirectional type rules (as shown in Figure 13 in Appendix C), where we in particular disallow [CHK] (or otherwise we could always choose to switch to inference mode for a variable which cannot resolve overloading). Moreover, we need to reconsider how to check applications since both [APP-FUNC] and [APP-ARGC] could apply.

6.3 Arguments First versus Functions First

In the previous example *show 1* we used the [APP-FUNC] rule to push the type of the argument into the function derivation in order to resolve *show* using [VARC]. However, sometimes we need to do the opposite and push the function type into the argument in order to disambiguate. Suppose we have an overloaded definition of *neg* as:

modi/neg : $\text{int} \rightarrow \text{int}$ *modf/neg* : $\text{float} \rightarrow \text{float}$

with *sqr* : $\text{float} \rightarrow \text{float}$. If we now consider the expression $\lambda x. \text{sqr} (\text{neg } x)$ we can only accept this if we use [APP-ARGC] on the application *sqr* (*neg x*) to propagate the *float* result type into the argument expression *neg x*. Otherwise, if we use [APP-FUNC] we cannot disambiguate the *neg* variable (since x has an abstract type at that point).

The optimal choice between using [APP-ARGC] or [APP-FUNC] cannot be made locally at an application node and depends on the sub-expressions. A straightforward implementation that tries all combinations would be exponential in the number of nested application nodes. Instead, following the “Pfenning recipe” [Dunfield and Krishnaswami 2021], we propose a syntax-directed approach that can be decided locally and can be easily understood by the programmer. This is also the approach used in the Koka language [Leijen 2021].

In particular, we will disallow [APP-FUNC], and always use [APP-ARGC] where we propagate the expected argument types into the arguments. The only exception is for a direct n -ary application

to a variable of the form $f\ e_1 \dots e_n$. In that case, we infer the least amount of arguments i such that we can disambiguate f , and then propagate the remaining argument types into the remaining argument expressions:

$$\begin{array}{c}
\text{least } i \text{ with } 0 \leq i \leq n, \text{ fresh } \alpha_{i+1}, \dots, \alpha_n, Q = Q_1, \dots, Q_i \\
Q_1 \mid \Gamma \vdash e_1 : \tau_1 \quad \dots \quad Q_i \mid \Gamma \vdash e_i : \tau_i \\
\text{unique } m/f : \sigma \in \Gamma \text{ with } Q_0 \vdash \sigma \sqsubseteq Q[\tau_1] \rightarrow \dots \rightarrow Q[\tau_i] \rightarrow \alpha_{i+1} \rightarrow \dots \rightarrow \alpha_n \rightarrow \tau \\
Q_{i+1} \mid \Gamma \vdash e_{i+1} \overset{\leftarrow}{:} Q_0[\alpha_{i+1}] \quad \dots \quad Q_n \mid \Gamma \vdash e_n \overset{\leftarrow}{:} Q_0[\alpha_n] \\
\hline
Q_0, Q_1, \dots, Q_n \mid \Gamma \vdash f\ e_1 \dots e_i \dots e_n \overset{\leftarrow}{:} \tau
\end{array} \text{---APPN}$$

This strategy is straightforward to implement: first try to disambiguate f (without any inference of the arguments) and keep inferring one argument at a time until f can be disambiguated, and then use checking rules for the remaining arguments. There are two drawbacks to this approach: a left-to-right bias, and argument types are never propagated into a lambda expression. As an example of the left-to-right bias, consider the following definitions:

modi/add : $\text{int} \rightarrow \text{int} \rightarrow \text{int}$ *modf/add* : $\text{float} \rightarrow \text{float} \rightarrow \text{float}$

The expression $\lambda x. \text{add } 1 (\text{neg } x)$ can be accepted by [APPN] since after inferring the type of 1, *add* is resolved to *modi/add* and the *int* type is propagated into the *neg x* argument which can subsequently be disambiguated to *modi/neg x*. However, the expression $\lambda x. \text{add } (\text{neg } x) 1$ is not accepted since the type of *neg x* cannot be inferred (as *neg* cannot be uniquely disambiguated). Secondly, since we otherwise always prefer to propagate types into arguments, no argument type is propagated into a lambda expression. For example $(\lambda x. \text{show } x) 1$ is rejected.

We believe though that having an easy rule for type propagation is preferable to trying to maximise the accepted programs, and the current rule seems to work out well in practice within the Koka language. Nevertheless, further experience may be warranted and other design approaches may be valid as well. For example, following Serrano et al. [2020], instead of strictly inferring from left-to-right we may first take a “quick look” at all expressions and infer “easy” expressions first. Or following Xie and Oliveira [2018], if the function expression in an application is syntactically a lambda expression, we could choose to propagate the argument types into the function.

As discussed in Section 3.3, in an implementation of HMQ we need to be careful to not leak type information between separate sub derivations. The example given was $\lambda x. (\text{inc } x, \text{show } x)$, where *inc* has type $\text{int} \rightarrow \text{int}$. This expression should be rejected since x will have an abstract type in each derivation of *inc x* and *show x* and thus *show* cannot be disambiguated. However, if we naively use algorithm W, the type of $x : \alpha$ is substituted after checking *inc x* to $x : \text{int}$, and subsequently *show x* can be disambiguated! When using HMQ extended with static overloading, it is important to use algorithm WQ (or the direct algorithm *inferD*) which uses fresh type variables for each occurrence of a λ -bound parameter (and correctly rejects the example expression).

The full syntax-directed bidirectional rules, including [APPN], are given in Appendix C. There we also show that for any possible derivation, each overloaded identifier is always resolved uniquely. These are also the rules used to implement static overloading in the Koka language [Leijen 2021]. Koka allows *locally qualified* names where one can directly define both *fun int/show* and *fun bool/show* within a module and use a plain *show* to resolve to either one. This works especially well in combination with the *syntactic implicits* feature of Koka that can be used for abstraction. This can express many examples that are typically addressed by type classes but in an arguably simpler way. For example, one can define a *list/show* function as:

```

fun list/show( xs : list<a>, ?show : a -> string ) : string
match xs
  Cons(x,Nil) -> show(x) // ~> ?show(x)
  Cons(x,xx) -> show(x) ++ ", " ++ show(xx) // ~> ?show(x) ++ ", " ++ list/show(xx)
  Nil -> ""

```

An implicit parameter `?show` in Koka is just a parameter whose syntactic name is resolved at each lexical call site, where regular static overloading is used to disambiguate based on the type context. For example, `show([1,2])` elaborates to `list/show([1,2],int/show)`.

7 Related Work

Damas and Milner [1982] introduce the now common HM type rules and show that type inference with algorithm W is sound and complete. This work builds on earlier work by Hindley [1969], who shows principal types exist for objects in combinatory logic, and Milner [1978] who gives the first description of algorithm W.

Prefixes. As discussed in Section 2.7, the main idea of inference under a prefix comes from the work on impredicative type inference in MLF as described by Le Botlan and Rémy [2003]. The prefix in MLF is much richer though and contains both polymorphic rigid bounds, $\alpha=\sigma$, and polymorphic flexible bounds $\alpha \geq \sigma$, where α can be any instance of σ . We can also quantify over bounds as $\forall Q.\sigma$ which, as shown in Section 2.7, could be useful for HMQ as well – as it allows us to unify both generalization rules into a single one. MLF is still an HM style system though where the type of λ -bound parameters is “guessed”. We believe it should be possible to extend HMQ naturally to MLFQ by extending the prefix to contain rich MLF bounds, and the equivalence relation to MLF equivalence. Leijen [2009] describes a restriction of MLF to only use flexible bounds which would lend itself well to a HMQ extension as it simplifies unification between polymorphic types.

Gundry, McBride, and McKinna [2010] describe type inference under a *context* Θ defined as:

$$\Theta ::= \emptyset \mid \Theta, \alpha : * \mid \Theta, \alpha : \tau : * \mid \Theta, x : \sigma \mid \Theta ;$$

where $\alpha : *$ can be viewed as an MLF instance constraint $\alpha \geq \perp$, and where $\alpha : \tau : *$ corresponds to our $\alpha = \tau$ bindings. The other forms are environment bindings $x : \sigma$ and ordering constraints $;$. A context restricted to just $\alpha : *$ and $\alpha : \tau : *$ bindings is written as Ξ , which can be viewed as a dependency ordered prefix. Indeed, the generalization rule is defined as [Gundry 2013, Fig. 2.9]:

$$\frac{\Theta_0 ; \vdash e : \tau \mid \Theta_1 ; \Xi}{\Theta_0 \vdash e : \tau \mid \Xi, \tau \mid \Theta_1} \text{GEN-CTX} \quad \begin{array}{ll} \forall \emptyset. \tau & = \tau \\ \forall (\alpha : *, \Xi). \tau & = \forall \alpha. (\forall \Xi. \tau) \\ \forall (\alpha : \tau' : *, \Xi). \tau & = [\alpha : \tau'] (\forall \Xi. \tau) \end{array}$$

which corresponds closely to the [GENX] rule of Section 2.7 (and the corresponding MLF generalization rule) where we quantify over a prefix, written here as $\forall \Xi. \tau$, where all monomorphic bounds are substituted. The main idea of having dependency ordered contexts is to simplify generalization where there is no need in the [GEN-CTX] rule to compute the free type variables in the environment (similar to using level-based generalization [Kiselyov 2022; Kuan and MacQueen 2007; Rémy 1992]).

The (algorithmic) application rule also closely matches our [APP] rule:

$$\frac{\Theta_0 \vdash e_1 : \tau_1 \mid \Theta_1 \quad \Theta_1 \vdash e_2 : \tau_2 \mid \Theta_2 \quad \Theta_2 \vdash \tau_1 \equiv \tau_2 \rightarrow \alpha \mid \Theta_3 \quad \text{fresh } \alpha}{\Theta_0 \vdash e_1 e_2 : \alpha \mid \Theta_3} \text{APP-CTX}$$

Here the context is statefully threaded through the rules but we believe it should be possible to define consistent context composition similar to our prefix composition such that sub derivations can be composed independently.

Constraint Based Inference. Type inference based on constraint generation has many similarities to the prefix based approach. These systems generate sets of unification constraints of the form $\tau_1 \equiv \tau_2$. Pierce [2002, §22.3] describes constraint based inference for a monomorphic calculus with

essentially the following abstraction and application rules:

$$\frac{C \mid \Gamma, x : \tau_1 \vdash e : \tau_2}{C \mid \Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \text{FUN-CON} \quad \frac{C_1 \mid \Gamma \vdash e_1 : \tau_1 \quad C_2 \mid \Gamma \vdash e_2 : \tau_2 \quad \text{fresh } \alpha}{C_1 \cup C_2 \cup \{\tau_1 \equiv \tau_2 \rightarrow \alpha\} \mid \Gamma \vdash e_1 e_2 : \alpha} \text{APP-CON}$$

Their [APP-CON] is quite similar to our [APP] rule, except that the constraint $\{\tau_1 \equiv \tau_2 \rightarrow \alpha\}$ is directly included while in HMQ one derives a prefix Q_3 from the equivalence relation $Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha$. In that sense, a prefix is a restricted form of a general constraint set which makes it closer to an implementation based on (in-place) substitutions. Just like the standard HM type rules, the constraint based system of Pierce still “guesses” types for lambda bound parameters.

In contrast, Heeren, Hage, and Swierstra [2002] describe a bottom-up constraint based system which uses abstract fresh variables for lambda-bound parameters. As part of the bottom-up inference, there is no top-down Γ environment, but instead a bottom-up *assumption* environment A . The abstraction, variable, and application rules are [Heeren 2005, Fig. 4.5]:

$$\frac{M \cup \{\alpha\} \mid C \mid A \vdash e : \tau \quad \text{fresh } \alpha}{M \mid C \cup \{\alpha \equiv \tau' \mid x : \tau' \in A\} \mid A/x \vdash \lambda x. e : \alpha \rightarrow \tau} \text{FUN-BU} \quad \frac{\text{fresh } \alpha}{M \mid \emptyset \mid \{x : \alpha\} \vdash x : \alpha} \text{VAR-BU}$$

$$\frac{M \mid C_1 \mid A_1 \vdash e_1 : \tau_1 \quad M \mid C_2 \mid A_2 \vdash e_2 : \tau_2 \quad \text{fresh } \alpha}{M \mid C_1 \cup C_2 \cup \{\tau_1 \equiv \tau_2 \rightarrow \alpha\} \mid A_1 \cup A_2 \vdash e_1 e_2 : \alpha} \text{APP-BU}$$

(where M is the set of monomorphic type variables used for the generation of generalization constraints in the let-rule). These bottom-up algorithmic type rules are closer to HMQ as the types of the lambda-bound parameters are abstract. Moreover, the use of an assumption together with the [VAR-BU] and [FUN-BU] rules is also close to algorithm WQ (Section 3.3) where we use fresh type variables for each occurrence and unify them all eventually at the λ expression, corresponding to the $\{\alpha \equiv \tau' \mid x : \tau' \in A\}$ constraint set in [FUN-BU]. However, depending on how type constraints $\tau_1 \equiv \tau_2$ are resolved, further restrictions may be needed to ensure all let-bindings have a principal type. With the addition of those restrictions, we believe that Heeren’s bottom-up algorithm can be a valid implementation for HMQ.

Unifying Substitutions. McAdam [1999] describes a new inference algorithm W' which does not have a right-to-left bias by computing substitutions for each subderivation independently and afterwards unifying the substitutions.

This is very similar how HMQ uses the notion of a consistent union of prefixes. If we keep all prefixes as an idempotent mapping, we can use McAdam’s U_s algorithm directly to compute the consistent union of two prefixes (as shown in Section 3.1.2).

8 HMQ Type Rules with Explicit Fresh Names

Figure 7 gives full inductive type rules for HMQ using explicit fresh names Δ . The fresh α notation used in Figure 2 is essentially a convenient shorthand for the rules here with explicit names. We write Δ_1, Δ_2 for the disjoint union of Δ_1 and Δ_2 , where $\Delta_1, \Delta_2 \doteq \Delta_1 \cup \Delta_2$ with $\Delta_1 \not\cap \Delta_2$. Every time we used fresh α in the rules in Figure 2, we now pick a fresh α disjoint from existing fresh names Δ (as Δ, α) in [INST], [FUN], and [APP]. As in the original rules, we rely on α -renaming in the [INST] rule such that the quantifier matches the fresh name. Such introduced fresh names are all eventually consumed by the generalization rules [GEN] and [GENSUB] that have Δ, α in their premise. A well-formedness condition for the new rules is that Δ is disjoint from the free type variables in the environment, with $\Delta \not\cap \text{ftv}(\Gamma)$. This ensures that for every derivation the fresh names are indeed fresh and do not contain type variable names occurring in the environment.

Lemma 8.10. (*Output type variables are either fresh or occur free in the environment*)

If $\Delta \mid Q \mid \Gamma \vdash e : \sigma$, then $\text{ftv}(Q, \sigma) \subseteq \text{ftv}(\Delta, \Gamma)$.

$$\begin{array}{ccccc}
\Delta & | & Q & | & \Gamma \vdash e : \sigma \\
\downarrow & & \downarrow & & \uparrow \\
\text{out} & & \text{out} & & \text{in} \\
& & & & \uparrow \\
& & & & \text{in} \\
& & & & \downarrow \\
& & & & \text{out}
\end{array}$$

with $\models Q$ and $\Delta \not\cap \text{ftv}(\Gamma)$

$$\begin{array}{c}
\frac{x : \sigma \in \Gamma}{\emptyset \mid \emptyset \mid \Gamma \vdash x : \sigma} \text{VAR} \quad \frac{\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma} \text{INST} \quad \frac{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q)}{\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma} \text{GEN} \\
\\
\frac{\Delta \mid Q \mid \Gamma, x : \alpha \vdash e : \tau}{\Delta, \alpha \mid Q \mid \Gamma \vdash \lambda x. e : \alpha \rightarrow \tau} \text{FUN} \quad \frac{\Delta, \alpha \mid Q \cdot \alpha = \tau \mid \Gamma \vdash e : \sigma}{\Delta \mid Q \mid \Gamma \vdash e : [\alpha := \tau] \sigma} \text{GENSUB} \\
\\
\frac{\Delta_1 \mid Q_1 \mid \Gamma \Vdash e_1 : \sigma \quad \Delta_2 \mid Q_2 \mid \Gamma, x : \sigma \vdash e_2 : \tau}{\Delta_1, \Delta_2 \mid Q_1, Q_2 \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \text{LET} \quad \frac{\Delta \mid Q \mid \Gamma \vdash e : \sigma \quad \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{\Delta \mid Q \mid \Gamma \Vdash e : \sigma} \text{MGEN} \\
\\
\frac{\Delta_1 \mid Q_1 \mid \Gamma \vdash e_1 : \tau_1 \quad \Delta_2 \mid Q_2 \mid \Gamma \vdash e_2 : \tau_2 \quad Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha}{\Delta_1, \Delta_2, \alpha \mid Q_1, Q_2, Q_3 \mid \Gamma \vdash e_1 e_2 : \alpha} \text{APP}
\end{array}$$

Fig. 7. Full HMQ type rules under a prefix, using explicit fresh names Δ , where we write Δ_1, Δ_2 for the union $\Delta_1 \cup \Delta_2$ where the elements are disjoint ($\Delta_1 \not\cap \Delta_2$). By construction, we also have $\text{ftv}(Q, \sigma) \subseteq \text{ftv}(\Delta, \Gamma)$.

These rules also no longer need the full conditions on [GEN] and [GENSUB] as in Figure 2. For both rules, the $\alpha \notin \text{ftv}(\Gamma)$ requirement is implied by the Δ, α premise (and the well-formedness condition $\Delta, \alpha \not\cap \Gamma$), while for [GENSUB], $\alpha \notin \text{ftv}(Q)$ is implied by the $Q \cdot \alpha = \tau$ premise.

The need for fresh names in HMQ is not ideal, and one might have hoped to see a local condition instead, for example:

$$\frac{Q \mid \Gamma \vdash e : \forall \alpha. \sigma \quad \alpha \notin \text{ftv}(Q, \Gamma)}{Q \mid \Gamma \vdash e : \sigma} \text{INST-WRONG}$$

Unfortunately, such local constraints still allow the introduction of artificial sharing by using the same variable name in separate sub-derivations, which eventually leads to non-principal derivations again. Consider for example *const id 1* with *const* : $\forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \alpha$. If we instantiate *const* to $\alpha \rightarrow \beta \rightarrow \alpha$, we could instantiate the quantifier for *id* to also be β (if we can use [INST-WRONG]):

$$\frac{\emptyset \mid \Gamma \vdash \text{const} : \alpha \rightarrow \beta \rightarrow \alpha \quad \emptyset \mid \Gamma \vdash \text{id} : \beta \rightarrow \beta \quad \dots \vdash \alpha \rightarrow \beta \rightarrow \alpha \approx (\beta \rightarrow \beta) \rightarrow \gamma}{\frac{\{\alpha = \beta \rightarrow \beta, \gamma = \beta \rightarrow \alpha\} \mid \Gamma \vdash \text{const id} : \gamma}{\emptyset \mid \Gamma \vdash \text{const id} : \beta \rightarrow (\beta \rightarrow \beta)} \text{GENSUB}} \text{APP}$$

and thus *const id 1* gets type $\text{int} \rightarrow \text{int}$ instead of the expected $\forall \alpha. \alpha \rightarrow \alpha$. The rules in Figure 7 ensure that separate sub-derivations all use unique names by using a disjoint union of names used in each sub-derivation (and thus, we cannot instantiate *const* and *id* with a shared β as in our example).

9 Conclusion

Type inference under a prefix gives us declarative type rules that we believe are close to the clarity of the original HM rules. At the same time, we are able to “read off” the algorithm from the declarative type rules. HMQ can serve as foundation to specify practical type systems in a declarative way that serves both purposes: users can easily reason about what programs are accepted by the type checker, while compiler writers can derive sound implementations from those same rules.

Acknowledgements

We would like to thank Tim Whiting for his feedback and many discussions about the design of static overloading with implicit parameters in Koka.

References

- Damas, and Milner. 1982. Principal Type-Schemes for Functional Programs. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 207–212. POPL’82. ACM, Albuquerque, New Mexico. doi:<https://doi.org/10.1145/582153.582176>.
- Dunfield, and Krishnaswami. May 2021. Bidirectional Typing. *ACM Comput. Surv.* 54 (5). ACM. doi:<https://doi.org/10.1145/3450952>.
- Emrich, Lindley, Stolarek, Cheney, and Coates. 2020. FreezeML: Complete and Easy Type Inference for First-Class Polymorphism. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 423–437. PLDI 2020. ACM, London, UK. doi:<https://doi.org/10.1145/3385412.3386003>.
- Emrich, Stolarek, Cheney, and Lindley. Aug. 2022. Constraint-Based Type Inference for FreezeML. *Proc. ACM Program. Lang.* 6 (ICFP). ACM press. doi:<https://doi.org/10.1145/3547642>.
- Garrigue, and Rémy. 1999. Semi-Explicit First-Class Polymorphism for ML. *Information and Computation* 155 (1): 134–169. doi:<https://doi.org/10.1006/inco.1999.2830>.
- Gundry. 2013. Type Inference, Haskell and Dependent Types. Phdthesis, University of Strathclyde, Department of Computer and Information Sciences.
- Gundry, McBride, and McKinna. 2010. Type Inference in Context. In *Proceedings of the Third ACM SIGPLAN Workshop on Mathematically Structured Functional Programming*, 43–54. MSFP’10. ACM, Baltimore, Maryland, USA. doi:<https://doi.org/10.1145/1863597.1863608>.
- Heeren. Sep. 2005. Top Quality Type Error Messages. Phdthesis, Institute of Information and Computing Sciences, Utrecht University. <https://dspace.library.uu.nl/bitstream/handle/1874/7297/full.pdf>.
- Heeren, Hage, and Swierstra. 2002. Generalizing Hindley-Milner Type Inference Algorithms. UU-CS-2002-031. Institute of Information and Computing Sciences, Utrecht University. <https://ics-archive.science.uu.nl/research/techreps/repo/CS-2002/2002-031.pdf>.
- Heeren, Leijen, and Ijzendoorn. 2003. Helium, for Learning Haskell. In *Proceedings of the 2003 ACM SIGPLAN Workshop on Haskell*, 62–71. Haskell’03. ACM, Uppsala, Sweden. doi:<https://doi.org/10.1145/871895.871902>.
- Hindley. 1969. The Principal Type-Scheme of an Object in Combinatory Logic. *Transactions of the American Mathematical Society* 146. American Mathematical Society: 29–60.
- Jones, and Diatchki. 2008. Language and Program Design for Functional Dependencies. In *Proc. of the First ACM SIGPLAN Symp. on Haskell*, 87–98. Haskell’08. ACM, Victoria, BC, Canada. doi:<https://doi.org/10.1145/1411286.1411298>.
- Kiselyov. 2022. How OCaml Type Checker Works – or What Polymorphism and Garbage Collection Have in Common. <https://okmij.org/ftp/ML/generalization.html>. Blog post.
- Kuan, and MacQueen. 2007. Efficient Type Inference Using Ranked Type Variables. In *Proc. of the ML Workshop*, 3–14. Freiburg, Germany. doi:<https://doi.org/10.1145/1292535.1292538>.
- Le Botlan. 2004. MLF: Une Extension de ML Avec Polymorphisme de Second Ordre et Instanciation Implicite. Phdthesis, l’école polytechnique, Paris.
- Le Botlan, and Rémy. 2003. MLF: Raising ML to the Power of System F. In *Proc. of the 8th ACM SIGPLAN Int. Conf. on Functional Programming*, 27–38. ICFP’03. ACM press, Uppsala, Sweden. doi:<https://doi.org/10.1145/944705.944709>.
- Leijen. Sep. 2007. HMF: Simple Type Inference for First-Class Polymorphism. MSR-TR-2007-118. Microsoft Research.
- Leijen. Sep. 2008. HMF: Simple Type Inference for First-Class Polymorphism. In *Proc. of the 13th ACM Symp. of the Int. Conf. on Functional Programming*. ICFP’08. Victoria, Canada. doi:<https://doi.org/10.1145/1411204.1411245>.
- Leijen. 2009. Flexible Types: Robust Type Inference for First-Class Polymorphism. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 66–77. POPL’09. ACM, Savannah, GA, USA. doi:<https://doi.org/10.1145/1480881.1480891>.
- Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *MSFP’14, 5th Workshop on Mathematically Structured Functional Programming*. doi:<https://doi.org/10.4204/EPTCS.153.8>.
- Leijen. 2021. The Koka Language. <https://koka-lang.github.io>.
- Leroy, and Mauny. 1993. Dynamics in ML. *Journal of Functional Programming* 3 (4): 431–463. doi:<https://doi.org/10.1017/S0956796800000848>.
- Lewis, Launchbury, Meijer, and Shields. 2000. Implicit Parameters: Dynamic Scoping with Static Types. In *Proc. of the 27th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, 108–118. POPL’00. ACM, Boston, MA, USA. doi:<https://doi.org/10.1145/325694.325708>.
- McAdam. 1999. On the Unification of Substitutions in Type Inference. In *Impl. of Functional Languages*, edited by K. Hammond, T. Davie, and C. Clack, 137–152. Springer, Berlin, Heidelberg. doi:https://doi.org/10.1007/3-540-48515-5_9.

- Milner. 1978. A Theory of Type Polymorphism in Programming. *Journal of Computer and System Sciences* 17 (3): 348–375. doi:[https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4).
- Odersky, and Läufer. 1996. Putting Type Annotations to Work. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 54–57. POPL '96. ACM, St. Petersburg Beach, Florida, USA. doi:<https://doi.org/10.1145/237721.237729>.
- Odersky, Zenger, and Zenger. 2001. Colored Local Type Inference. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 41–53. POPL '01. ACM, London, United Kingdom. doi:<https://doi.org/10.1145/360204.360207>.
- Peyton Jones, Jones, and Meijer. Jan. 1997. Type Classes: An Exploration of the Design Space. In *Haskell Workshop*. <https://www.microsoft.com/en-us/research/publication/type-classes-an-exploration-of-the-design-space/>.
- Peyton Jones, Vytiniotis, Weirich, and Shields. 2007. Practical Type Inference for Arbitrary-Rank Types. *Journal of Functional Programming* 17 (1): 1–82. doi:<https://doi.org/10.1017/S0956796806006034>.
- Pierce. Feb. 2002. *Types and Programming Languages (TAPL)*. 1st edition. The MIT Press, Cambridge, Massachusetts 02142.
- Pierce, and Turner. Jan. 2000. Local Type Inference. *ACM Trans. Program. Lang. Syst.* 22 (1). ACM: 1–44. doi:<https://doi.org/10.1145/345099.345100>.
- Rémy. 1992. *Extending ML Type System with a Sorted Equational Theory*. Research Report 1766. Rocquencourt, BP 105, 78 153 Le Chesnay Cedex, France.
- Robinson. Jan. 1965. A Machine-Oriented Logic Based on the Resolution Principle. *J. ACM* 12 (1). ACM: 23–41. doi:<https://doi.org/10.1145/321250.321253>.
- Selsam, Ullrich, and Moura. 2020. Tabled Typeclass Resolution. doi:<https://doi.org/10.48550/arXiv.2001.04301>.
- Serrano, Hage, Peyton Jones, and Vytiniotis. Aug. 2020. A Quick Look at Impredicativity. *Proc. ACM Program. Lang.* 4 (ICFP). ACM. doi:<https://doi.org/10.1145/3408971>.
- Vytiniotis, Peyton Jones, and Schrijvers. 2010. Let Should Not Be Generalized. In *Proc. of the 5th ACM SIGPLAN Workshop on Types in Language Design and Impl.*, 39–50. TLDI '10. ACM, Madrid, Spain. doi:<https://doi.org/10.1145/1708016.1708023>.
- Vytiniotis, Weirich, and Peyton Jones. 2006. Boxy Types: Inference for Higher-Rank Types and Impredicativity. In *Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, 251–262. ICFP '06. ACM press, Portland, Oregon, USA. doi:<https://doi.org/10.1145/1159803.1159838>.
- Wadler, and Blott. 1989. How to Make Ad-Hoc Polymorphism Less Ad Hoc. In *Proc. of the 16th ACM SIGPLAN-SIGACT Symp. on Principles of Prog. Lang.*, 60–76. POPL'89. ACM, Austin, Texas, USA. doi:<https://doi.org/10.1145/75277.75283>.
- Wright, and Felleisen. Nov. 1994. A Syntactic Approach to Type Soundness. *Inf. Comput.* 115 (1): 38–94. doi:<https://doi.org/10.1006/inco.1994.1093>.
- Xie, and Oliveira. 2018. Let Arguments Go First. Edited by Amal Ahmed. *Programming Languages and Systems*, LNCS, 10801. Springer International Publishing: 272–299. doi:https://doi.org/10.1007/978-3-319-89884-1_10. ESOP'18.

$$\boxed{\begin{array}{c} \Gamma \vdash_{\text{HM}} e : \sigma \\ \uparrow \quad \uparrow \quad \downarrow \\ \text{in} \quad \text{in} \quad \text{out} \end{array}}$$

$$\frac{x : \sigma \in \Gamma}{\Gamma \vdash_{\text{HM}} x : \sigma} \text{VAR}_{\text{HM}} \qquad \frac{\Gamma \vdash_{\text{HM}} e_1 : \sigma \quad \Gamma, x : \sigma \vdash_{\text{HM}} e_2 : \tau}{\Gamma \vdash_{\text{HM}} \text{let } x = e_1 \text{ in } e_2 : \tau} \text{LET}_{\text{HM}}$$

$$\frac{\Gamma, x : \tau_1 \vdash_{\text{HM}} e : \tau_2}{\Gamma \vdash_{\text{HM}} \lambda x. e : \tau_1 \rightarrow \tau_2} \text{FUN}_{\text{HM}} \qquad \frac{\Gamma \vdash_{\text{HM}} e_1 : \tau_2 \rightarrow \tau \quad \Gamma \vdash_{\text{HM}} e_2 : \tau_2}{\Gamma \vdash_{\text{HM}} e_1 e_2 : \tau} \text{APP}_{\text{HM}}$$

$$\frac{\Gamma \vdash_{\text{HM}} e : \forall \alpha. \sigma}{\Gamma \vdash_{\text{HM}} e : [\alpha := \tau] \sigma} \text{INST}_{\text{HM}} \qquad \frac{\Gamma \vdash_{\text{HM}} e : \sigma \quad \alpha \notin \text{ftv}(\Gamma, e)}{\Gamma \vdash_{\text{HM}} e : \forall \alpha. \sigma} \text{GEN}_{\text{HM}}$$

Fig. 8. Damas-Hindley-Milner type rules.

$$\boxed{\begin{array}{c} \Gamma \vdash_{\text{HMS}} e : \tau \\ \uparrow \quad \uparrow \quad \downarrow \\ \text{in} \quad \text{in} \quad \text{out} \end{array}} \qquad \text{gen}(\Gamma, \tau) = \forall \bar{\alpha}. \tau \text{ with } \bar{\alpha} = \text{ftv}(\tau) - \text{ftv}(\Gamma)$$

$$\frac{x : \forall \bar{\alpha}. \tau \in \Gamma}{\Gamma \vdash_{\text{HMS}} x : [\bar{\alpha} := \bar{\tau}] \tau} \text{VAR}_{\text{HMS}} \qquad \frac{\Gamma \vdash_{\text{HMS}} e_1 : \tau_1 \quad \Gamma, x : \sigma \vdash_{\text{HMS}} e_2 : \tau_2 \quad \sigma = \text{gen}(\Gamma, \tau_1)}{\Gamma \vdash_{\text{HMS}} \text{let } x = e_1 \text{ in } e_2 : \tau_2} \text{LET}_{\text{HMS}}$$

$$\frac{\Gamma, x : \tau_1 \vdash_{\text{HMS}} e : \tau_2}{\Gamma \vdash_{\text{HMS}} \lambda x. e : \tau_1 \rightarrow \tau_2} \text{FUN}_{\text{HMS}} \qquad \frac{\Gamma \vdash_{\text{HMS}} e_1 : \tau_2 \rightarrow \tau \quad \Gamma \vdash_{\text{HMS}} e_2 : \tau_2}{\Gamma \vdash_{\text{HMS}} e_1 e_2 : \tau} \text{APP}_{\text{HMS}}$$

Fig. 9. Syntax directed HM type rules.

A The Damas-Hindley-Milner Type Rules

Figure 8 gives the standard HM type rules [Damas and Milner 1982]. A judgment $\Gamma \vdash_{\text{HM}} e : \sigma$ states that an expression can be given type σ under a type environment Γ . Γ and e are inherited while σ is synthesized. We write $\Gamma, x : \sigma$ to extend a type environment Γ with a fresh binding $x : \sigma$ where $x \notin \text{dom}(\Gamma)$ (which we can always ensure by appropriate renaming).

The $[\text{VAR}_{\text{HM}}]$ rule derives the type of a variable that is bound in the environment. This will always be a monomorphic type τ for lambda-bound variables but can be polymorphic for let-bound variables as the $[\text{LET}_{\text{HM}}]$ rule allows a σ type for the binding. As discussed in the introduction, the $[\text{FUN}_{\text{HM}}]$ rule allows “guessing” any τ_1 type for the parameter. The $[\text{INST}_{\text{HM}}]$ rule is another source of “guessing”, as we can freely instantiate a polymorphic binder to any τ' that fits the derivation.

A.1 Syntax Directed Type Rules for Damas-Hindley-Milner

As a step towards an inference algorithm, we can also give syntax directed rules for the HM rules as shown in Figure 9. Following Damas and Milner [1982], we always instantiate variables and generalize at let-bindings.

A.2 HM Type Inference: Algorithm W

Damas and Milner [1982] describe a type inference algorithm W (shown as *inferW* in Figure 10) which always infers a most-general type, and they show it is sound and complete with respect to the inference rules.

$ \begin{aligned} &unify : (\tau_1, \tau_2) \rightarrow \theta \\ &unify(\alpha, \alpha) = \\ &\quad id \\ &unify(\alpha, \tau) \text{ or } (\tau, \alpha) \mid \alpha \notin \text{ftv}(\tau) = \\ &\quad [\alpha := \tau] \\ &unify(\tau_1 \rightarrow \tau_2, \tau'_1 \rightarrow \tau'_2) = \\ &\quad \text{let } \theta_1 = unify(\tau_1, \tau'_1) \\ &\quad \text{let } \theta_2 = unify(\theta_1 \tau_2, \theta_1 \tau'_2) \\ &\quad (\theta_2 \circ \theta_1) \\ &gen : (\Gamma, \tau) \rightarrow \sigma \\ &gen(\Gamma, \tau) = \\ &\quad \text{let } \bar{\alpha} = \text{ftv}(\tau) - \text{ftv}(\Gamma) \\ &\quad \forall \bar{\alpha}. \tau \end{aligned} $	$ \begin{aligned} &inferW : (\Gamma, e) \rightarrow (\theta, \tau) \\ &inferW(\Gamma, x) = \\ &\quad \text{let } \forall \bar{\alpha}. \tau = \Gamma(x) \\ &\quad \text{let } \bar{\beta} = \text{fresh} \\ &\quad (id, [\bar{\alpha} := \bar{\beta}]\tau) \\ &inferW(\Gamma, e_1 e_2) = \\ &\quad \text{let } (\theta_1, \tau_1) = inferW(\Gamma, e_1) \\ &\quad \text{let } (\theta_2, \tau_2) = inferW(\theta_1 \Gamma, e_2) \\ &\quad \text{let } \alpha = \text{fresh} \\ &\quad \text{let } \theta_3 = unify(\theta_2 \tau_1, \tau_2 \rightarrow \alpha) \\ &\quad (\theta_3 \circ \theta_2 \circ \theta_1, \theta_3 \alpha) \\ &inferW(\Gamma, \lambda x. e) = \\ &\quad \text{let } \alpha = \text{fresh} \\ &\quad \text{let } (\theta, \tau) = inferW((\Gamma, x : \alpha), e) \\ &\quad (\theta, \theta \alpha \rightarrow \tau) \\ &inferW(\Gamma, \text{let } x = e_1 \text{ in } e_2) = \\ &\quad \text{let } (\theta_1, \tau_1) = inferW(\Gamma, e_1) \\ &\quad \text{let } \sigma = gen(\theta_1 \Gamma, \tau_1) \\ &\quad \text{let } (\theta_2, \tau_2) = inferW((\theta_1 \Gamma, x : \sigma), e_2) \\ &\quad (\theta_2 \circ \theta_1, \tau_2) \end{aligned} $
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 10. Algorithm W

$$\begin{array}{c}
\boxed{\begin{array}{ccccc} \Delta & | & Q & | & \Gamma \vdash_s e : \tau \\ \downarrow & & \downarrow & & \uparrow \\ \text{out} & & \text{out} & & \text{in} \end{array}}, \text{ and } \boxed{\begin{array}{ccccc} \Delta & | & Q & | & \Gamma \Vdash_s e : \sigma \\ \uparrow & & \downarrow & & \uparrow \\ \text{in} & & \text{out} & & \text{in} \end{array}} \quad \text{with } \vdash Q, \text{ and } \Delta \not\vdash \text{ftv}(\Gamma)
\end{array}$$

$$\begin{array}{c}
\frac{x : \forall \bar{\alpha}. \tau \in \Gamma}{\bar{\alpha} \mid \emptyset \mid \Gamma \vdash_s x : \tau} \text{VAR}_s \qquad \frac{\Delta_0 \mid Q_0 \mid \Gamma \vdash_s e : \tau \quad (\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \tau)}{\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma} \text{MGEN}_s \\
\\
\frac{\Delta \mid Q \mid \Gamma, x : \alpha \vdash_s e : \tau}{\Delta, \alpha \mid Q \mid \Gamma \vdash_s \lambda x. e : \alpha \rightarrow \tau} \text{FUN}_s \qquad \frac{\Delta_1 \mid Q_1 \mid \Gamma \Vdash_s e_1 : \sigma \quad \Delta_2 \mid Q_2 \mid \Gamma, x : \sigma \vdash_s e_2 : \tau_2}{\Delta_1, \Delta_2 \mid Q_1, Q_2 \mid \Gamma \vdash_s \text{let } x = e_1 \text{ in } e_2 : \tau_2} \text{LET}_s \\
\\
\frac{\Delta_1 \mid Q_1 \mid \Gamma \vdash_s e_1 : \tau_1 \quad \Delta_2 \mid Q_2 \mid \Gamma \vdash_s e_2 : \tau_2 \quad Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha}{\Delta_1, \Delta_2, \alpha \mid Q_1, Q_2, Q_3 \mid \Gamma \vdash_s e_1 e_2 : \alpha} \text{APP}_s
\end{array}$$

$$\begin{array}{l}
\text{gen} : (\Delta, Q, \Gamma, \sigma) \rightarrow (\Delta, Q, \sigma), \text{ with } \Delta \not\vdash \text{ftv}(\Gamma) \\
\text{gen}((\Delta, \alpha), Q \cdot \alpha = \tau', \Gamma, \tau) = \text{gen}(\Delta, Q, \Gamma, [\alpha := \tau']\tau) \\
\text{gen}((\Delta, \alpha), Q, \Gamma, \sigma) = \text{gen}(\Delta, Q, \Gamma, \forall \alpha. \sigma) \quad \text{if } \alpha \notin \text{ftv}(Q) \\
\text{gen}(\Delta, Q, \Gamma, \sigma) = (\Delta, Q, \sigma) \quad \text{if } (\text{dom}(Q) \cup \text{ftv}(\sigma)) \subseteq \text{ftv}(\Gamma)
\end{array}$$

Fig. 11. Syntax directed type rules under a prefix

B Syntax Directed Type Rules for HMQ

Figure 11 gives the syntax directed type rules for HMQ. Section D.9 shows how we can rewrite any derivation into a canonical form, where all instantiation is done at the leaves of a derivation at a variable occurrence, and all generalization is done at let bindings. The $[\text{VAR}_s]$ rule now instantiates the type of variable fully with fresh variables $\bar{\alpha}$ for the quantifiers – just like in the $[\text{INST}]$ rule this may require α -renaming of the quantifiers.

The $[\text{LET}_s]$ rules uses the (\Vdash_s) $[\text{MGEN}_s]$ rule to generalize the type. The generalization rule $[\text{MGEN}_s]$ uses the gen function that takes the fresh variables Δ , the prefix Q , the environment Γ , and a monotype τ , and returns a subset of the free variables Δ' , a new prefix Q' , and a generalized type scheme σ . The new prefix Q' only has bindings that still occur in Γ with $\text{dom}(Q') \subseteq \text{ftv}(\Gamma)$. The first case of generalization essentially applies $[\text{GENSUB}]$ for all binders α in Q that do not occur free in Γ , while the second case corresponds the $[\text{GEN}]$ rule, where we quantify over all free variables in τ that do not occur free in Γ .

In Section D.9.6 we show that the syntax directed rules are sound and complete to the HMQ type rules. In particular, if $\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma$, then also $\Delta \mid Q \mid \Gamma \vdash e : \sigma$, and dually, if $\Delta \mid Q \mid \Gamma \vdash e : \sigma$, then also $\Delta' \mid Q' \mid \Gamma \Vdash_s e : \sigma$ with $\Delta' \subseteq \Delta$ and $Q' \subseteq Q$.

B.1 Algorithm WQ

As discussed in Section 3.3, Figure 12 gives a type inference algorithm WQ for HMQ which is closely based on algorithm W. The algorithm assumes a pre-processing step where every lambda binding x is annotated (as x_n) with the total number of occurrences n in the body of the lambda, and where every lambda bound variable occurrence is annotated with its (unique) occurrence i (as x_i). This way, we can generate an initial fresh variable α for a lambda bound variable, and use that to generate a unique fresh type variable α_i per occurrence. Eventually, we unify all α_i occurrences again. This removes any accidental propagation of type inference unifications between different derivations which happens in plain algorithm W.

This use of a unique type variable per occurrence may delay unification errors until the moment all occurrences are unified. However, at the same time it may improve the precision of the type error since we can for example pick the most probable error instead of the the one that occurs

$\begin{aligned} & \text{unify} : (\tau_1, \tau_2) \rightarrow \theta \\ & \text{unify}(\alpha, \alpha) = \\ & \quad id \\ & \text{unify}(\alpha, \tau) \text{ or } (\tau, \alpha) \mid \alpha \notin \text{ftv}(\tau) = \\ & \quad [\alpha := \tau] \\ & \text{unify}(\tau_1 \rightarrow \tau_2, \tau'_1 \rightarrow \tau'_2) = \\ & \quad \text{let } \theta_1 = \text{unify}(\tau_1, \tau'_1) \\ & \quad \text{let } \theta_2 = \text{unify}(\theta_1 \tau_2, \theta_1 \tau'_2) \\ & \quad (\theta_2 \circ \theta_1) \\ & \text{unifies} : [\tau] \rightarrow \theta \\ & \text{unifies} [] \text{ or } [\tau] = id \\ & \text{unifies}(\tau_1 : \tau_2 : \bar{\tau}) = \\ & \quad \text{let } \theta_1 = \text{unify}(\tau_1, \tau_2) \\ & \quad \text{let } \theta_2 = \text{unifies}(\theta_1 \tau_2 : \theta_1 \bar{\tau}) \\ & \quad (\theta_2 \circ \theta_1) \\ & \text{gen} : (\Gamma, \tau) \rightarrow \sigma \\ & \text{gen}(\Gamma, \tau) = \\ & \quad \text{let } \bar{\alpha} = \text{ftv}(\tau) - \text{ftv}(\Gamma) \\ & \quad (\forall \bar{\alpha}. \tau) \end{aligned}$	$\begin{aligned} & \text{inferWQ} : (\Gamma, e) \rightarrow (\theta, \tau) \\ & \text{inferWQ}(\Gamma, x_i) = \\ & \quad \text{let } \alpha = \Gamma(x) \\ & \quad (id, \alpha_i) \\ & \text{inferWQ}(\Gamma, x) = \\ & \quad \text{let } \forall \bar{\alpha}. \tau = \Gamma(x) \\ & \quad \text{let } \bar{\beta} = \text{fresh} \\ & \quad (id, [\bar{\alpha} := \bar{\beta}] \tau) \\ & \text{inferWQ}(\Gamma, e_1 e_2) = \\ & \quad \text{let } (\theta_1, \tau_1) = \text{inferWQ}(\Gamma, e_1) \\ & \quad \text{let } (\theta_2, \tau_2) = \text{inferWQ}(\theta_1 \Gamma, e_2) \\ & \quad \text{let } \alpha = \text{fresh} \\ & \quad \text{let } \theta_3 = \text{unify}(\theta_2 \tau_1, \tau_2 \rightarrow \alpha) \\ & \quad (\theta_3 \circ \theta_2 \circ \theta_1, \theta_3 \alpha) \\ & \text{inferWQ}(\Gamma, \lambda x_n. e) = \\ & \quad \text{let } \alpha = \text{fresh} \\ & \quad \text{let } (\theta_1, \tau) = \text{inferWQ}((\Gamma, x : \alpha), e) \\ & \quad \text{let } \theta_2 = \text{unifies}(\alpha, \theta_1 \alpha_1, \dots, \theta_1 \alpha_n) \\ & \quad (\theta_2, \theta_2 \alpha \rightarrow \theta_2 \tau) \\ & \text{inferWQ}(\Gamma, \text{let } x = e_1 \text{ in } e_2) = \\ & \quad \text{let } (\theta_1, \tau_1) = \text{inferWQ}(\Gamma, e_1) \\ & \quad \text{let } \sigma = \text{gen}(\theta_1 \Gamma, \tau_1) \\ & \quad \text{let } (\theta_2, \tau_2) = \text{inferWQ}((\theta_1 \Gamma, x : \sigma), e_2) \\ & \quad (\theta_2 \circ \theta_1, \tau_2) \end{aligned}$
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 12. Algorithm WQ

first [Heeren et al. 2003], i.e. for λx . (*inc* x , *sqr* x , *not* x) we can give an error for the single *bool* occurrence since the two *int* occurrences are more common.

$$\boxed{\begin{array}{c} \Delta \mid Q \mid \Gamma \vdash_s e \vdash \tau \\ \downarrow \text{out} \quad \downarrow \text{out} \quad \uparrow \text{in} \quad \uparrow \text{in} \quad \uparrow \text{in} \end{array}}, \quad \boxed{\begin{array}{c} \Delta \mid Q \mid \Gamma \vdash_s e \vdash \tau \\ \downarrow \text{out} \quad \downarrow \text{out} \quad \uparrow \text{in} \quad \uparrow \text{in} \quad \downarrow \text{out} \end{array}}, \quad \boxed{\begin{array}{c} \Delta \mid Q \mid \Gamma \Vdash_s e \vdash \sigma \\ \downarrow \text{out} \quad \downarrow \text{out} \quad \uparrow \text{in} \quad \uparrow \text{in} \quad \downarrow \text{out} \end{array}} \quad \text{with } \models Q, \Delta \not\vdash \text{ftv}(\Gamma)$$

$$\begin{array}{c}
\frac{\text{unique } m/x : \sigma \in \Gamma \text{ with } \Delta \mid Q \vdash_s \sigma \sqsubseteq \tau}{\Delta \mid Q \mid \Gamma \vdash_s x \vdash \tau \rightsquigarrow m/x} \text{VARC}_s \quad \frac{Q \vdash \tau_1 \approx \tau_2}{\bar{\alpha} \mid Q \vdash_s \forall \bar{\alpha}. \tau_1 \sqsubseteq \tau_2} \text{INSTANCEC}_s \\
\\
\frac{\Delta \mid Q \mid \Gamma, x : \tau_1 \vdash_s e \vdash \tau_2}{\Delta \mid Q \mid \Gamma \vdash_s \lambda x. e \vdash \tau_1 \rightarrow \tau_2} \text{FUNC}_s \quad \frac{\Delta \mid Q_1 \mid \Gamma \vdash_s \lambda x. e \vdash \alpha_1 \rightarrow \alpha_2 \quad Q_2 \vdash \alpha_1 \rightarrow \alpha_2 \approx \alpha}{\Delta, \alpha_1, \alpha_2 \mid Q_1, Q_2 \mid \Gamma \vdash_s \lambda x. e \vdash \alpha} \text{IFUNC}_s \\
\\
\frac{\Delta \mid Q \mid \Gamma \vdash_s e \vdash \alpha}{\Delta, \alpha \mid Q \mid \Gamma \vdash_s e \vdash Q[\alpha]} \text{INFC}_s \quad \frac{\Delta_0 \mid Q_0 \mid \Gamma \vdash_s e \vdash \tau \quad (\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \tau)}{\Delta \mid Q \mid \Gamma \Vdash_s e \vdash \sigma} \text{GENC}_s \\
\\
\frac{\Delta_1 \mid Q_1 \mid \Gamma \Vdash_s e_1 \vdash \sigma \quad \Delta_2 \mid Q_2 \mid \Gamma, x : \sigma \vdash_s e_2 \vdash \tau}{\Delta_1, \Delta_2 \mid Q_1, Q_2 \mid \Gamma \vdash_s \text{let } x = e_1 \text{ in } e_2 \vdash \tau} \text{LET}_s \\
\\
\frac{\Delta_1 \mid Q_1 \mid \Gamma \vdash_s e_1 \vdash \beta \rightarrow \tau \quad \Delta_2 \mid Q_2 \mid \Gamma \vdash_s e_2 \vdash Q_1[\beta]}{\Delta_1, \Delta_2, \beta \mid Q_1, Q_2 \mid \Gamma \vdash_s e_1 e_2 \vdash \tau} \text{APP-ARGC}_s \\
\\
\begin{array}{l}
\text{least } i \text{ with } 0 \leq i \leq n \text{ and } Q = Q_1, \dots, Q_n \\
\text{unique } m/f : \sigma \in \Gamma \text{ with } \Delta \mid Q_0 \vdash_s \sigma \sqsubseteq Q[\tau_1] \rightarrow \dots \rightarrow Q[\tau_i] \rightarrow \alpha_{i+1} \rightarrow \dots \rightarrow \alpha_n \rightarrow \tau \\
\Delta_1 \mid Q_1 \mid \Gamma \vdash_s e_1 \vdash \tau_1 \quad \dots \quad \Delta_i \mid Q_i \mid \Gamma \vdash_s e_i \vdash \tau_i \\
\Delta_{i+1} \mid Q_{i+1} \mid \Gamma \vdash_s e_{i+1} \vdash Q_0[\alpha_{i+1}] \quad \dots \quad \Delta_n \mid Q_n \mid \Gamma \vdash_s e_n \vdash Q_0[\alpha_n]
\end{array} \\
\hline
\Delta, \Delta_1, \dots, \Delta_n, \alpha_{i+1}, \dots, \alpha_n \mid Q_0, Q_1, \dots, Q_n \mid \Gamma \vdash_s f e_1 \dots e_i \dots e_n \vdash \tau \quad \text{APPN}_s
\end{array}$$

Fig. 13. Syntax-directed bidirectional type checking rules for static overloading. We always prefer [APPN_s] over [APP-ARGC] when applicable.

C Static Overloading

For static overloading we cannot use the bidirectional rules of Figure 6 directly since these are not syntax-directed. In particular, the [CHK] rule can arbitrarily switch from checking mode to inference mode. Instead, we need to use syntax-directed bidirectional inference to deterministically resolve overloaded variables.

Figure 13 gives such rules based on the syntax-directed inference rules of Figure 11. An interesting aspect is that now all rules are checking rules with just two exceptions for [INFC_s] and [GENC_s]. In the [LET] rule, the type of the binding x must be inferred and we use [GENC_s] to infer a most general type, which uses in turn [INFC_s] to infer a mono type. The [INFC_s] rule infers a type by switching to checking if the expression can be typed with a fresh monotype α . This is also used in the [APPN_s] rule to infer the first i arguments needed to resolve f . These rules are all fully determined by the syntax (and the shape of the propagated type in [IFUNC_s]).

We write $e[x]$ for an expression containing an occurrence of a variable x . We can now state that for any type derivation, the overloaded variables are uniquely resolved:

Theorem C.11. (*Overloaded variables are always resolved uniquely*)

If $\Delta' \mid Q' \mid \Gamma' \Vdash_s e[x] \xrightarrow{\text{blue}} \sigma$, with $\emptyset \mid Q_1 \mid \Gamma \vdash_s x \xrightarrow{\text{red}} \tau_1 \rightsquigarrow m_1/x$ at the x location in $e[x]$, we have for any other derivation $\Delta'' \mid Q'' \mid \Gamma' \Vdash_s e[x] : \sigma$, with $\emptyset \mid Q_2 \vdash x \xrightarrow{\text{red}} \tau_2 \rightsquigarrow m_2/x$ at the x location in $e[x]$, that $\tau_1 = \tau_2$, $Q_1 = Q_2$, and $m_1 = m_2$.

Proof. (of Theorem C.11) Since the rules in Figure 13 are syntax directed, the structure of the two derivations must match exactly. Furthermore, by Theorem D.28, the premises $Q \vdash \tau_1 \approx \tau_2$ and $Q_2 \vdash \alpha_1 \rightarrow \alpha_2 \approx \alpha$ in $[\text{INSTANCEC}_s]$ and $[\text{IFUNC}_s]$ are also principal. Therefore, both derivations are exactly equal with no derivation choices. In particular, in a $[\text{VARC}_s]$ leaf, where we have $\Delta \mid Q \vdash x \xrightarrow{\text{red}} \tau \rightsquigarrow m/x$, τ is always the same for any derivation for $e[x]$, and $[\text{VARC}]$ always resolves to the same unique m/x for that particular τ . \square

D Proofs

D.1 Substitutions

A *substitution* θ is an idempotent function from type variables to types. The (finite) domain of θ is the set of type variables such that $\theta(\alpha) \neq \alpha$ for any $\alpha \in \text{dom}(\theta)$, while the codomain consists of the free type variables of its range.

We use the notation $[\alpha := \tau]$ for a singleton substitution θ with domain $\{\alpha\}$ and $\theta(\alpha) = \tau$. We usually write substitution application with explicit parenthesis as $\theta(\tau)$ but sometimes shorten to just $\theta\tau$ when appropriate⁵.

We write $\theta \sqsubseteq \theta'$ if θ is a *more general* (or less specific) substitution than θ' , such that $\theta' = \theta'' \circ \theta$ for some θ'' . We say that two substitutions are equivalent if each is as-general as the other, i.e. $\theta_1 \sqsubseteq \theta_2 \wedge \theta_2 \sqsubseteq \theta_1 \Leftrightarrow \theta_1 \equiv \theta_2$ (such substitutions are not always exactly equal since they can potentially differ on a renaming between type variables $\alpha_1 := \alpha_2$ where either direction is possible).

Properties D.12. (Substitution)

1. For any τ and well-formed θ , $\text{dom}(\theta) \not\subseteq \text{ftv}(\theta(\tau))$.
2. For any τ with $\text{dom}(\theta) \not\subseteq \text{ftv}(\tau)$, $\theta(\tau) = \tau$.
3. For a well-formed θ with $\alpha \notin \text{ftv}(\theta, \tau')$, $\theta([\alpha := \tau'](\tau)) = [\alpha := \theta(\tau')](\theta(\tau))$.

Proof. (of Property D.12.3) We have $\alpha \notin \text{ftv}(\theta, \tau')$ (1), and thus $\theta(\alpha) = \alpha$ (2). Induction on τ .

Case $[\tau = \alpha]$:

$$\begin{aligned} & \theta([\alpha := \tau'](\alpha)) \\ = & \theta(\tau') \quad \{ \text{def.} \} \\ = & [\alpha := \theta(\tau')](\alpha) \quad \{ \text{def.} \} \\ = & [\alpha := \theta(\tau')](\theta(\alpha)) \quad \{ (2) \} \end{aligned}$$

Case $[\tau = \tau_1]$: with $\alpha \notin \text{ftv}(\tau_1)$ (3):

$$\begin{aligned} & \theta([\alpha := \tau'](\tau_1)) \\ = & \theta(\tau_1) \quad \{ (3). \} \\ = & [\alpha := \theta(\tau')](\theta(\tau_1)) \quad \{ (3), (1), \text{prop D.12.2} \} \end{aligned}$$

Case $[\tau = \tau_1 \rightarrow \tau_2]$:

$$\begin{aligned} & \theta([\alpha := \tau'](\tau_1 \rightarrow \tau_2)) \\ = & \theta([\alpha := \tau']\tau_1 \rightarrow \theta([\alpha := \tau']\tau_2)) \quad \{ \text{def.} \} \\ = & [\alpha := \theta(\tau')](\theta(\tau_1)) \rightarrow [\alpha := \theta(\tau')](\theta(\tau_2)) \quad \{ \text{ind. hyp.} \} \\ = & [\alpha := \theta(\tau')](\theta(\tau_1 \rightarrow \tau_2)) \quad \{ \text{def.} \} \end{aligned}$$

□

Lemma D.13. (Substitution cancelation)

If $\alpha \notin \text{ftv}(\tau)$ and θ is well-formed with $\alpha \in \text{dom}(\theta)$, then $\theta = [\alpha := \tau'] \circ \theta$ (for any τ').

Proof. (of Lemma D.13) With $\alpha \in \text{dom}(\theta)$, we have for any τ , $\alpha \notin \text{ftv}(\theta(\tau))$ (1) (due to prop D.12.1). Therefore,

$$\begin{aligned} & ([\alpha := \tau'] \circ \theta)(\tau) \\ = & [\alpha := \tau'](\theta(\tau)) \quad \{ \text{def} \} \\ = & \theta(\tau) \quad \{ (1), \text{prop D.12.2} \} \end{aligned}$$

□

⁵Note that we do not use the common notation $\tau[\alpha := \tau']$ but always write this as $[\alpha := \tau'](\tau)$ (or $[\alpha := \tau']\tau$).

Lemma D.14. (*Substitution commutation*)

If $\alpha \notin \text{ftv}(\theta)$ (1), then $\theta \circ [\alpha := \tau] = [\alpha := \theta(\tau)] \circ \theta$

Proof. (of Lemma D.14) For any τ' , we have:

$$\begin{aligned}
& (\theta \circ [\alpha := \tau])(\tau') \\
= & \theta([\alpha := \tau](\tau')) \quad \{ \text{def.} \} \\
= & [\alpha := \theta(\tau)](\theta(\tau')) \quad \{ \text{prop D.12.3} \} \\
= & ([\alpha := \theta(\tau)] \circ \theta)(\tau') \quad \{ \text{def.} \}
\end{aligned}$$

□

Lemma D.15. (*Equivalence of Composed Substitutions*)

If $\theta_1 \circ [\alpha := \tau_1] = \theta_2 \circ [\alpha := \tau_2]$ with $\alpha \notin \text{ftv}(\theta_1, \tau_1, \theta_2, \tau_2)$, then $\theta_1 = \theta_2$ with $\theta_1(\tau_1) = \theta_1(\tau_2)$.

Proof. (of Lemma D.15) We have $\theta = \theta_1 \circ [\alpha := \tau_1] = \theta_2 \circ [\alpha := \tau_2]$ (1), with $\alpha \notin \text{ftv}(\theta_1, \tau_1, \theta_2, \tau_2)$ (2). For any τ with $\alpha \notin \text{ftv}(\tau)$, we have $\theta_1(\tau) = \theta_2(\tau)$ (3):

$$\begin{aligned}
& \theta_1(\tau) \\
= & (\theta_1 \circ [\alpha := \tau_1])(\tau) \quad \{ \alpha \notin \text{ftv}(\tau), (2) \} \\
= & \theta(\tau) \quad \{ (1) \} \\
= & (\theta_2 \circ [\alpha := \tau_2])(\tau) \quad \{ (1) \} \\
= & \theta_2(\tau) \quad \{ \alpha \notin \text{ftv}(\tau), (2) \}
\end{aligned}$$

and thus $\theta_1(\tau_1) = \theta_1(\tau_2)$ (4):

$$\begin{aligned}
& \theta_1(\tau_1) \\
= & (\theta_1 \circ [\alpha := \tau_1])(\alpha) \quad \{ (2) \} \\
= & (\theta_2 \circ [\alpha := \tau_2])(\alpha) \quad \{ (1) \} \\
= & \theta_2(\tau_2) \quad \{ (2) \} \\
= & \theta_1(\tau_2) \quad \{ (2), (3) \}
\end{aligned}$$

Therefore:

$$\begin{aligned}
& \theta_1 \circ [\alpha := \tau_2] \\
= & [\alpha := \theta_1(\tau_2)] \circ \theta_1 \quad \{ \text{Lemma D.14}, (2) \} \\
= & [\alpha := \theta_1(\tau_1)] \circ \theta_1 \quad \{ (4) \} \\
= & \theta_1 \circ [\alpha := \tau_1] \quad \{ \text{Lemma D.14}, (2) \} \\
= & \theta_2 \circ [\alpha := \tau_2] \quad \{ (1) \}
\end{aligned}$$

and thus $\theta_1 = \theta_2$.

□

D.2 More General Substitutions**Properties D.16.**

1. $\text{id} \sqsubseteq \theta$ (for any θ)
2. If $\theta_1 \sqsubseteq \theta_2$ then $\theta_1 \circ \theta \sqsubseteq \theta_2 \circ \theta$.
3. If $\theta(\alpha) = \theta(\tau)$ with $\alpha \notin \text{ftv}(\tau)$, then $[\alpha := \tau] \sqsubseteq \theta$.
4. If $\theta_1 \circ \theta_2 \sqsubseteq \theta$ then $\theta = \theta' \circ \theta_2$ with $\theta_1 \sqsubseteq \theta'$.
5. If $\theta \sqsubseteq \theta_1 \circ \theta_2$ and $\text{dom}(\theta_1) \cap \text{dom}(\theta) = \emptyset$, then $\theta \sqsubseteq \theta_2$.

Proof. (of Property D.16.1) For any θ ,

$$\begin{aligned}
& \theta \\
= & \theta \circ \text{id} \quad \{ \text{def.} \}
\end{aligned}$$

□

Proof. (of Property D.16.2) With $\theta_1 \sqsubseteq \theta_2$, we have $\theta_2 = \theta' \circ \theta_1$ (1), and thus:

$$\begin{aligned} & \theta_2 \circ \theta \\ = & (\theta' \circ \theta_1) \circ \theta \quad \{ (1) \} \\ = & \theta' \circ (\theta_1 \circ \theta) \quad \{ \text{assoc.} \} \end{aligned}$$

□

Proof. (of Property D.16.3) We have $\theta(\alpha) = \theta(\tau)$ (1).

$$\begin{aligned} & \theta(\alpha) = \theta(\tau) \\ \Rightarrow & \theta = \theta' \circ [\alpha := \tau] \end{aligned}$$

□

Proof. (of Property D.16.4) We have $\theta_1 \circ \theta_2 \sqsubseteq \theta$ (1).

$$\begin{aligned} & \theta \\ = & \theta'' \circ \theta_1 \circ \theta_2 \quad \{ (1), \text{ some } \theta'' \} \\ = & \theta' \circ \theta_2 \quad \{ \text{assume } \theta' = \theta'' \circ \theta_1 \text{ (2)} \} \\ \text{and} & \\ & \theta' \\ = & \theta'' \circ \theta_1 \quad \{ (2) \} \\ \Rightarrow & \theta_1 \sqsubseteq \theta' \end{aligned}$$

□

Proof. (of Property D.16.5) We have $\theta \sqsubseteq \theta_1 \circ \theta_2$ (1) and $\text{dom}(\theta_1) \not\cap \text{dom}(\theta)$ (2)

$$\begin{aligned} & \theta_1 \circ \theta_2 \\ = & \theta' \circ \theta \quad \{ (1) \} \\ = & \theta_3 \circ \theta_4 \circ \theta \quad \{ \text{for some } \theta_3, \theta_4 \text{ with } \text{dom}(\theta_4) \not\cap \text{dom}(\theta_1) \} \end{aligned}$$

since $\text{dom}(\theta_4 \circ \theta) \not\cap \text{dom}(\theta_1)$, it must be that $\theta_2 = \theta_4 \circ \theta$, and thus $\theta \sqsubseteq \theta_2$. □

D.3 Prefixes

The solution to a prefix is also a solution to any subset:

Lemma D.17. (Consistent weakening)

If $\theta \models Q_1 \cup Q_2$, then also $\theta \models Q_1$.

Proof. (of Lemma D.17) We have $\theta \models Q_1 \cup Q_2$ (1) and need to show $\theta \models Q_1$. From (1) and [SOLUTION], we have $\forall(\alpha=\tau) \in Q_1 \cup Q_2$. $\theta(\alpha) = \theta(\tau)$, and thus $\forall(\alpha=\tau) \in Q_1$. $\theta(\alpha) = \theta(\tau)$ with $\theta \models Q_1$. □

The least solution of a subset of a prefix is less specific than the prefix solution:

Lemma D.18. (Consistent union)

If $\models Q_1 \cup Q_2$, then $\langle Q_1 \rangle \sqsubseteq \langle Q_1 \cup Q_2 \rangle$.

Proof. (of Lemma D.18) By definition $\langle Q_1 \cup Q_2 \rangle \models Q_1 \cup Q_2$, and by lemma D.17 $\langle Q_1 \cup Q_2 \rangle$ is a solution of Q_1 , $\langle Q_1 \cup Q_2 \rangle \models Q_1$. By definition of prefix solution we now have $\langle Q_1 \rangle \sqsubseteq \langle Q_1 \cup Q_2 \rangle$.

An important property is that a prefix solution of a subset is a right identity:

Lemma D.19. (Prefix extension)

If $\models Q_1 \cup Q_2$, then $\langle Q_1 \cup Q_2 \rangle = \langle Q_1 \cup Q_2 \rangle \circ \langle Q_1 \rangle$.

Proof. (of Lemma D.19) We have $\models Q_1 \cup Q_2$ (1) and need to show $\langle Q_1 \cup Q_2 \rangle = \langle Q_1 \cup Q_2 \rangle \circ \langle Q_1 \rangle$. By (1) and Lemma D.18, $\langle Q_1 \rangle \sqsubseteq \langle Q_1 \cup Q_2 \rangle$, which implies $\langle Q_1 \cup Q_2 \rangle = \theta \circ \langle Q_1 \rangle$ for some θ . Therefore

$$\begin{aligned} & \langle Q_1 \cup Q_2 \rangle \\ = & \theta \circ \langle Q_1 \rangle & \{ (2) \} \\ = & \theta \circ \langle Q_1 \rangle \circ \langle Q_1 \rangle & \{ \text{subst. idem.} \} \\ = & \langle Q_1 \cup Q_2 \rangle \circ \langle Q_1 \rangle & \{ (2) \} \end{aligned}$$

□

A nice property of an extracted bound is that we get a stronger form of Lemma D.19 where we can write the prefix solution as a composition of each sub-solution:

Lemma 3.9. (Extraction corresponds to composition of prefix solutions)

If $\models Q$ and $Q = Q' \cdot \alpha = \tau$, then $\langle Q \rangle = \langle Q' \rangle \circ [\alpha := \tau]$.

Proof. (of Lemma 3.9) We have $\models Q$ (1), and $Q = Q' \cdot \alpha = \tau$, and thus $Q = Q' \uplus \{\alpha = \tau\}$ (2) with $\alpha \notin \text{ftv}(Q', \tau)$ (3) (and need to show $\langle Q \rangle = \langle Q' \rangle \circ [\alpha := \tau]$). From (1), we have $\forall \beta = \tau' \in Q. Q[\beta] = Q[\tau']$, and thus from (2,3), $\forall (\beta = \tau') \in Q'. Q[\beta] = Q[\tau']$ and $Q[\alpha] = Q[\tau]$ (4). From (3,4) and Lemma D.16.3, $[\alpha := \tau] \sqsubseteq \langle Q \rangle$ and thus $\langle Q \rangle = \theta \circ [\alpha := \tau]$ for some θ (5), and from (3,5) $\forall (\beta = \tau') \in Q'. \theta(\beta) = \theta(\tau')$. A minimal solution for θ is $\langle Q' \rangle$, and therefore $\langle Q \rangle = \langle Q' \rangle \circ [\alpha := \tau]$. □

Lemma D.21. (Prefixes with a common solution are consistent)

For any consistent Q_1 and Q_2 , if $Q_1 \sqsubseteq \theta$ and $Q_2 \sqsubseteq \theta$, then Q_1, Q_2 is consistent with $(Q_1, Q_2) \sqsubseteq \theta$.

Proof. (Of Lemma D.21) We have $\models Q_1$ (1a) with $Q_1 \sqsubseteq \theta$ (1b), and $\models Q_2$ with $Q_2 \sqsubseteq \theta$. By definition, from (1a), $\forall \alpha = \tau \in Q_1. \langle Q_1 \rangle \alpha = \langle Q_1 \rangle \tau$, and thus by (1b) we also have $\forall \alpha = \tau \in Q_1. \theta \alpha = \theta \tau$ with $\theta \models Q_1$ (2) and similarly, $\theta \models Q_2$ (3). It follows by definition that we also have $\theta \models (Q_1 \cup Q_2)$ and thus Q_1, Q_2 is consistent. Since $\langle Q_1, Q_2 \rangle$ is minimal by definition, we also have $(Q_1, Q_2) \sqsubseteq \theta$. □

Lemma D.22. (Prefix composition is an identity on substitution)

If $\models (Q_1, Q_2)$, then $(Q_1, Q_2) = (Q_1, Q_1[Q_2])$.

Proof. (of Lemma D.22) With $\theta \models (Q_1, Q_2)$, we have $\forall (\alpha = \tau) \in Q_1 \cup Q_2. \theta \alpha = \theta \tau$ (1). Now for any $\alpha = \tau' \in Q_1[Q_2]$, it must be $\alpha = \tau \in Q_2$, with $\tau' = Q_1[\tau]$. By Lemma D.18, $Q_1 \sqsubseteq \theta$, and we have $\theta = \theta' \cdot Q_1$, and thus $\theta \tau' = \theta(Q_1[\tau]) = (\theta' \cdot Q_1 \cdot Q_1)[\tau] = (\theta' \cdot Q_1)[\tau] = \theta \tau$ (2). Therefore, $\theta \alpha = \theta \tau$ (by (2)), and thus by (2), $\theta \alpha = \theta \tau'$, and $\theta \models (Q_1 \cup Q_1[Q_2])$. □

D.4 Consistent Prefixes are Substitutions

Proof. (Of Theorem 2.8) For any consistent Q we have a minimal solution $\langle Q \rangle (= \theta)$ which is a well-formed idempotent substitution of the form $[\alpha_1 := \tau_1, \dots, \alpha_n := \tau_n]$ (with α_i pairwise distinct and $\text{dom}(\theta) \not\cap \text{codom}(\theta)$) (1)). Define Q' as $\{\alpha_1 = \tau_1, \dots, \alpha_n = \tau_n\}$. The minimal solution for Q' is also θ and thus $Q \equiv Q'$ where Q' is an idempotent mapping (by (1)). □

Theorem D.23. (For a consistent prefix, any binder can be extracted)

If $\models Q$ and $(\alpha = \tau) \in Q$, then $Q \equiv Q' \cdot \alpha = \tau'$ (with $Q[\tau] = \tau'$).

Proof. (Of Lemma D.23) Since $\models Q$, by Theorem 2.8, we have a idempotent mapping Q_0 (1) with $Q_0 \equiv Q$ (2). Since $\alpha = \tau \in Q$, we must have $Q_0 = Q' \cup \{\alpha = \tau'\}$ for some τ' . Since Q_0 is an idempotent mapping (1), we have $\alpha \notin \text{ftv}(Q', \tau')$, and we can use [EXTRACT] to conclude $Q_0 = Q' \cdot \alpha = \tau'$ and by (2), $Q_0 \equiv Q' \cdot \alpha = \tau'$. Moreover, by (2), $Q[\tau] = Q_0[\tau']$, and since Q_0 is idempotent (1), $Q_0[\tau'] = \tau'$,

and therefore $Q[\tau] = \tau'$. \square

D.5 Type Equivalence

Type equivalence does not create fresh variables.

Lemma D.24.

If $Q \vdash \tau_1 \approx \tau_2$ then $\text{ftv}(Q) \subseteq \text{ftv}(\tau_1, \tau_2)$.

Proof. Follows directly by inspection of each equivalence rule.

Type equivalence is sound.

Proof. (of Theorem 2.1) We have $Q \vdash \tau \approx \tau'$ (1) (and need to show $Q[\tau] = Q[\tau']$). We proceed by induction over the rules of (\approx) .

Case [EQ-ID]: $\tau = \tau'$, and we have $\emptyset \vdash \tau \approx \tau$. It follows directly that $Q[\tau] = \emptyset[\tau] = \tau = \tau' = Q[\tau']$.

Case [EQ-VAR]: $\tau = \alpha$, and we have $\{\alpha = \tau'\} \vdash \alpha \approx \tau'$ with $\alpha \notin \text{ftv}(\tau')$ (2).

$$\begin{aligned}
 & Q[\alpha] \\
 = & \langle \{\alpha = \tau'\} \rangle(\alpha) \quad \{ \text{assumption} \} \\
 = & [\alpha := \tau'](\alpha) \quad \{ \text{def.} \} \\
 = & \tau' \quad \{ \text{def.} \} \\
 = & [\alpha := \tau'](\tau') \quad \{ (2) \} \\
 = & Q[\tau'] \quad \{ \text{def.} \}
 \end{aligned}$$

Case [EQ-REFL]: By the premise, $Q \vdash \tau' \approx \tau$ (2), and by induction we have $Q[\tau'] = Q[\tau]$.

Case [EQ-FUN]: We have $Q_1 \vdash \tau_1 \approx \tau'_1$, $Q_2 \vdash \tau_2 \approx \tau'_2$, and $\models Q_1 \cup Q_2$. By induction $Q_1[\tau_1] = Q_1[\tau'_1]$ and $Q_2[\tau_2] = Q_2[\tau'_2]$ (2). We can now derive:

$$\begin{aligned}
 & (Q_1, Q_2)[\tau_1 \rightarrow \tau_2] \\
 = & (Q_1, Q_2)[\tau_1] \rightarrow (Q_1, Q_2)[\tau_2] \quad \{ \text{def.} \} \\
 = & (Q_1, Q_2)[Q_1[\tau_1]] \rightarrow (Q_1, Q_2)[Q_2[\tau_2]] \quad \{ \text{lemma D.19} \} \\
 = & (Q_1, Q_2)[Q_1[\tau'_1]] \rightarrow (Q_1, Q_2)[Q_2[\tau'_2]] \quad \{ (2) \} \\
 = & (Q_1, Q_2)[\tau'_1] \rightarrow (Q_1, Q_2)[\tau'_2] \quad \{ \text{lemma D.19} \} \\
 = & (Q_1, Q_2)[\tau'_1 \rightarrow \tau'_2] \quad \{ \text{def.} \}
 \end{aligned}$$

\square

Lemma D.25. (Consistent solutions)

If $Q \sqsubseteq \theta$, then $\theta \models Q$.

Proof. (of Lemma D.25) Since $\langle Q \rangle \sqsubseteq \theta$, $\theta = \theta' \circ \langle Q \rangle$ for some θ' . Since $\langle Q \rangle \models Q$, we have $\forall(\alpha = \tau) \in Q. \langle Q \rangle(\alpha) = \langle Q \rangle(\tau)$. Therefore, $\forall(\alpha = \tau) \in Q. (\theta' \circ \langle Q \rangle)(\alpha) = (\theta' \circ \langle Q \rangle)(\tau)$, and $\theta \models Q$. \square

Lemma D.26. (Consistent strengthen)

If $Q_1 \sqsubseteq \theta$ and $Q_2 \sqsubseteq \theta$, then $\models Q_1 \cup Q_2$ and $Q_1 \cup Q_2 \sqsubseteq \theta$

Proof. (of Lemma D.26) We have $Q_1 \sqsubseteq \theta$ (1) and $Q_2 \sqsubseteq \theta$ (2). From (1) and Lemma D.25, we have $\theta \models Q_1$, and thus $\forall(\alpha = \tau) \in Q_1. \theta(\alpha) = \theta(\tau)$ (3). Similarly, from (2) we have $\forall(\alpha = \tau) \in Q_2. \theta(\alpha) = \theta(\tau)$ (4). Therefore, by (3,4) $\forall(\alpha = \tau) \in (Q_1 \cup Q_2). \theta(\alpha) = \theta(\tau)$, and thus $\theta \models Q_1 \cup Q_2$ (5) (and $\models Q_1 \cup Q_2$). By definition, (5) implies $\langle Q_1 \cup Q_2 \rangle \sqsubseteq \theta$. \square

Type equivalence is complete.

Proof. (of Theorem 2.2) We have $\theta(\tau_1) = \theta(\tau_2)$ (1). We proceed by induction over the shape of τ_1, τ_2 .

Case $[\theta(\tau) = \theta(\tau)]$: We have $\tau = \tau_1 = \tau_2$, and thus by [EQ-ID], $\emptyset \vdash \tau_1 \approx \tau_2$, and $\emptyset \sqsubseteq \theta$ (since $\theta = \theta \circ id$).

Case $[\theta(\alpha) = \theta(\tau_2)]$: If $\tau_2 = \alpha$, we have the previous case $\theta(\tau) = \theta(\tau)$. Since θ is idempotent, we otherwise have $\alpha \notin \text{ftv}(\tau_2)$ (2). By [EQ-VAR], $\{\alpha = \tau_2\} \vdash \alpha \approx \tau_2$. Moreover, by (2) and Property D.16.3, we also have $[\alpha = \tau_2] \sqsubseteq \theta$.

Case $[\theta(\tau_1) = \theta(\alpha)]$: We can apply [EQ-REFL] and by induction $\theta(\alpha) = \theta(\tau_1)$.

Case $[\theta(\tau_1 \rightarrow \tau_2) = \theta(\tau'_1 \rightarrow \tau'_2)]$: We have $\theta(\tau_1) = \theta(\tau'_1)$ and $\theta(\tau_2) = \theta(\tau'_2)$. By ind. hyp. $Q_1 \vdash \tau_1 \approx \tau'_1$ (2) with $Q_1 \sqsubseteq \theta$ and $Q_2 \vdash \tau_2 \approx \tau'_2$ (3) with $Q_2 \sqsubseteq \theta$. By lemma D.26, we have $\vdash Q_1 \cup Q_2$ (3) and $Q_1 \cup Q_2 \sqsubseteq \theta$, and thus by (2,3,4), $Q_1, Q_2 \vdash \tau_1 \rightarrow \tau_2 \approx \tau'_1 \rightarrow \tau'_2$.

□

The following lemma establishes the relation between substitution and equivalence, and is important to show that type derivations are principal (Theorem 2.3).

Lemma D.27. (Type equivalent substitution)

If $Q_1 \vdash Q[\tau_1] \approx Q[\tau_2]$ then also $Q_2 \vdash \tau_1 \approx \tau_2$ with $(Q, Q_1) = (Q, Q_2)$.

Proof. (of Lemma D.27)

Case [EQ-VAR]: We have $Q_1 = \{\alpha = Q[\tau_2]\} \vdash \alpha \approx Q[\tau_2]$ with $Q[\tau_1] = \alpha$ (and $\alpha \notin \text{ftv}(Q[\tau_2])$). From [EQ-VAR], we also have $Q_2 = \{\alpha = \tau_2\} \vdash \alpha \approx \tau_2$.

$$\begin{aligned}
& Q, Q_2 \\
= & Q, \{\alpha = \tau_2\} \quad \{ (1a) \} \\
= & Q, Q[\{\alpha = \tau_2\}] \quad \{ \text{Lemma D.22} \} \\
= & Q, \{\alpha = Q[\tau_2]\} \\
= & Q, Q_1
\end{aligned}$$

Case [EQ-FUN]: We have $Q_1 = (Q_3, Q_4) \vdash \tau_3 \rightarrow \tau_4 \approx \tau_5 \rightarrow \tau_6$ (1), where $Q[\tau_1] = \tau_3 \rightarrow \tau_4$ (1a) and $Q[\tau_2] = \tau_5 \rightarrow \tau_6$ (1b). From the premise of (1), we have $Q_3 \vdash \tau_3 \approx \tau_5$ (2a) and $Q_4 \vdash \tau_4 \approx \tau_6$ (2b). We proceed by case analysis on the shape of τ_1 and τ_2 .

subcase $\tau_1 = \tau'_3 \rightarrow \tau'_4$, $\tau_2 = \tau'_5 \rightarrow \tau'_6$, with $\tau_3 = Q[\tau'_3]$, $\tau_4 = Q[\tau'_4]$, etc. By induction on (2a), we also have $Q'_3 \vdash \tau'_3 \approx \tau'_5$ (3a) with $(Q, Q_3) = (Q, Q'_3)$ (3b). By induction on (2b), we also have $Q'_4 \vdash \tau'_4 \approx \tau'_6$ (4a) with $(Q, Q_4) = (Q, Q'_4)$ (4b). We now have $Q_2 = (Q'_3, Q'_4) \vdash \tau'_3 \rightarrow \tau'_4 \approx \tau'_5 \rightarrow \tau'_6$, by (3a,4a) with [EQ-FUN]. Therefore:

$$\begin{aligned}
& Q, Q_1 \\
= & Q, Q_3, Q_4 \\
= & Q, Q_3, Q, Q_4 \\
= & Q, Q'_3, Q, Q'_4 \quad \{ (3b, 4b) \} \\
= & Q, Q'_3, Q'_4 \\
= & Q, Q_2
\end{aligned}$$

subcase $\tau_1 = \tau'_3 \rightarrow \tau'_4$, $\tau_2 = \alpha$, $\alpha \notin \text{ftv}(\tau_1)$, with $\tau_3 = Q[\tau'_3]$, $\tau_4 = Q[\tau'_4]$, and $\tau_5 \rightarrow \tau_6 = Q[\alpha]$. We can use [EQ-VAR] to conclude $Q_2 = \{\alpha = \tau_1\} \vdash \tau_1 \approx \alpha$, and:

$$\begin{aligned}
& Q, Q_2 \\
= & (Q, \{\alpha = \tau_5 \rightarrow \tau_6\}), Q_2 \quad \{ (1b), Q[\alpha] = \tau_5 \rightarrow \tau_6 \} \\
= & Q, \{\alpha = \tau_5 \rightarrow \tau_6\}, \{\alpha = \tau_1\} \\
= & Q, \{\alpha = \tau_5 \rightarrow \tau_6\}, Q[\{\alpha = \tau_1\}] \quad \{ \text{Lemma D.22} \} \\
= & Q, \{\alpha = \tau_5 \rightarrow \tau_6\}, \{\alpha = Q[\tau_1]\} \\
= & Q, \{\alpha = \tau_5 \rightarrow \tau_6\}, \{\alpha = \tau_3 \rightarrow \tau_4\} \quad \{ (1a) \} \\
= & Q, \{\alpha = \tau_5 \rightarrow \tau_6\}, \{\alpha = \tau_3 \rightarrow \tau_4\} \quad \{ (1), \text{Theorem 2.7} \} \\
= & Q, (Q_3, Q_4), \{\alpha = \tau_5 \rightarrow \tau_6\} \quad \{ (1b), Q[\alpha] = \tau_5 \rightarrow \tau_6 \} \\
= & Q, (Q_3, Q_4)
\end{aligned}$$

subcase $\tau_1 = \alpha$, $\tau_2 = \tau'_5 \rightarrow \tau'_6$: as the previous case.

subcase $\tau_1 = \alpha$, $\tau_2 = \beta$, with $\alpha \neq \beta$. We can use [EQ-VAR] to conclude $Q_2 = \{\alpha = \beta\} \vdash \beta \approx \alpha$, and:

$$\begin{aligned}
& Q, Q_2 \\
= & (Q, \{\alpha = \tau_3 \rightarrow \tau_4, \beta = \tau_5 \rightarrow \tau_6\}), Q_2 \quad \{ (1a, 1b), Q[\alpha] = \tau_3 \rightarrow \tau_4, Q[\beta] = \tau_5 \rightarrow \tau_6 \} \\
= & Q, \{\alpha = \tau_3 \rightarrow \tau_4, \beta = \tau_5 \rightarrow \tau_6\}, \{\alpha = \beta\} \\
= & Q, \{\alpha = \tau_3 \rightarrow \tau_4, \beta = \tau_5 \rightarrow \tau_6, \alpha = \beta\} \\
= & Q, \{\alpha = \tau_3 \rightarrow \tau_4, \beta = \tau_3 \rightarrow \tau_4, \beta = \tau_5 \rightarrow \tau_6\} \quad \{ \text{Theorem 2.7} \} \\
= & Q, \{\beta = \tau_3 \rightarrow \tau_4, \beta = \tau_5 \rightarrow \tau_6\} \quad \{ (1a) \} \\
= & Q, (Q_3, Q_4), \{\beta = \tau_5 \rightarrow \tau_6\} \quad \{ (1), \text{Theorem 2.7} \} \\
= & Q, (Q_3, Q_4) \quad \{ (1b) \} \\
= & Q, Q_1
\end{aligned}$$

subcase $\tau_1 = \alpha$, $\tau_2 = \alpha$. We can use [EQ-REFL] to conclude $Q_2 = \emptyset \vdash \alpha \approx \alpha$, and:

$$\begin{aligned}
& Q, Q_2 \\
= & Q, \emptyset \\
= & Q \\
= & Q, \{\alpha = \tau_3 \rightarrow \tau_4, \alpha = \tau_5 \rightarrow \tau_6\} \quad \{ (1a, 1b), Q[\alpha] = \tau_3 \rightarrow \tau_4, Q[\alpha] = \tau_5 \rightarrow \tau_6 \} \\
= & Q, (Q_3, Q_4), \{\alpha = \tau_5 \rightarrow \tau_6\} \quad \{ (1), \text{Theorem 2.7} \} \\
= & Q, (Q_3, Q_4) \quad \{ (1b) \} \\
= & Q, Q_1
\end{aligned}$$

Case [EQ-REFL]: We have $Q_1 \vdash Q[\tau_1] \approx Q[\tau_2]$. From the premise $Q_1 \vdash Q[\tau_2] \approx Q[\tau_1]$, it follows by induction that $Q_2 \vdash \tau_2 \approx \tau_1$ (1a) with $(Q, Q_1) = (Q, Q_2)$ (1b). With [EQ-REFL] and (1a), we also have $Q_2 \vdash \tau_1 \approx \tau_2$ with $(Q, Q_1) = (Q, Q_2)$ (1b).

Case [EQ-ID]: We have $Q_1 \vdash Q[\tau_1] \approx Q[\tau_2]$ where $Q_1 = \emptyset$ (1). By Theorem 2.1 and (1), $Q[\tau_1] = Q[\tau_2]$, and by Theorem 2.2, there is a Q_2 with $Q_2 \vdash \tau_1 \approx \tau_2$ with $Q_2 \sqsubseteq Q$ (2), and:

$$\begin{aligned}
& (Q, Q_2) \\
= & Q \quad \{ (2) \} \\
= & (Q, \emptyset) \\
= & (Q, Q_1) \quad \{ (1) \}
\end{aligned}$$

□

Finally, we can show that type equivalence derivations are principal and always derive the same constraints.

Theorem D.28. (*Principal type equivalence*)

If $Q_1 \vdash \tau_1 \approx \tau_2$, then for any other derivation $Q_2 \vdash \tau_1 \approx \tau_2$, we have $Q_1 = Q_2$.

Proof. (of Theorem D.28) The only non-syntax directed rule is [EQ-REFL]. We observe that we can change any derivation $Q \vdash \tau_1 \approx \tau_2$ to a canonical form with the same Q by always “floating” all [EQ-REFL] applications to the leaves of the derivation. We consider all rules followed by [EQ-REFL] and rewrite those to float up the [EQ-REFL] application:

Case [EQ-FUN]: We have:

$$\frac{\frac{Q_1 \vdash \tau_1 \approx \tau'_1 \quad Q_2 \vdash \tau_2 \approx \tau'_2}{(Q_1, Q_2) \vdash \tau_1 \rightarrow \tau_2 \approx \tau'_1 \rightarrow \tau'_2} \text{EQ-FUN}}{(Q_1, Q_2) \vdash \tau'_1 \rightarrow \tau'_2 \approx \tau_1 \rightarrow \tau_2}$$

From the premise, we can rewrite this as:

$$\frac{\frac{Q_1 \vdash \tau_1 \approx \tau'_1}{Q_1 \vdash \tau'_1 \approx \tau_1} \quad \frac{Q_2 \vdash \tau_2 \approx \tau'_2}{Q_2 \vdash \tau'_2 \approx \tau_2}}{(Q_1, Q_2) \vdash \tau'_1 \rightarrow \tau'_2 \approx \tau_1 \rightarrow \tau_2} \text{EQ-FUN}$$

Case [EQ-REFL]: We have

$$\frac{\frac{Q \vdash \tau_1 \approx \tau_2}{Q \vdash \tau_2 \approx \tau_1} \text{EQ-REFL}}{Q \vdash \tau_1 \approx \tau_2}$$

These cancel out and it is equivalent to: $Q \vdash \tau_1 \approx \tau_2$.

Case [EQ-ID]: We have

$$\frac{}{\emptyset \vdash \tau \approx \tau} \text{EQ-ID}$$

which is equivalent to just [EQ-ID]

Case [EQ-VAR]: We have

$$\frac{\frac{\alpha \notin \text{ftv}(\tau)}{\{\alpha=\tau\} \vdash \alpha \approx \tau} \text{EQ-VAR}}{\{\alpha=\tau\} \vdash \tau \approx \alpha}$$

We need to consider two sub-cases:

- (1) $\tau \neq \beta$ (τ is not a type variable): In that case this is the only possible derivation.
- (2) $\tau = \beta$. We have $\beta \neq \alpha$ since $\alpha \notin \text{ftv}(\tau)$. That means we can *also* derive $\{\beta=\alpha\} \vdash \beta \approx \alpha$ by using [EQ-VAR] directly. However, these prefixes are actually equivalent, with $\{\alpha=\beta\} = \{\beta=\alpha\}$, since each minimal substitution is as general as the other (see Section D.1).

Note that we could consider an alternative approach to avoid this, where we make the equivalence of two type variables fully deterministic. Since we can write any derivation by floating all [EQ-REFL] application to the [EQ-VAR] leaves, we can specialize for this. We remove [EQ-REFL] and replace it with:

$$\frac{Q \vdash \alpha \approx \tau \quad \tau \neq \beta}{Q \vdash \tau \approx \alpha} \text{EQ-VARR}$$

This would make the variable equivalence effectively left-biased.

□

D.6 Computing Prefixes

We can replace sub-prefixes with equivalent constraints:

Lemma D.29. (*Prefix replacement*)

If $Q_1 \equiv Q_2$, then $Q \cup Q_1 \equiv Q \cup Q_2$.

Proof. (of Lemma D.29) We have $Q_1 \equiv Q_2$ and by definition $\langle Q_1 \rangle = \langle Q_2 \rangle$ (1) (and need to show $Q \cup Q_1 \equiv Q \cup Q_2$). Since $Q_1 \sqsubseteq Q \cup Q_1$, we have from Lemma D.18, $\langle Q \cup Q_1 \rangle = \theta_1 \circ \langle Q_1 \rangle$ (2) and similarly, $\langle Q \cup Q_2 \rangle = \theta_2 \circ \langle Q_2 \rangle$ (3) for some θ_1, θ_2 .

By definition, θ_1 is a least substitution such that $\forall(\alpha=\tau) \in Q. (\theta_1 \circ \langle Q_1 \rangle)(\alpha) = (\theta_1 \circ \langle Q_1 \rangle)(\tau)$ (4) $\wedge \forall(\alpha=\tau) \in Q_1. (\theta_1 \circ \langle Q_1 \rangle)(\alpha) = (\theta_1 \circ \langle Q_1 \rangle)(\tau)$. The second part induces no constraints on θ_1 since it holds for any substitution, so θ_1 only needs to be the least substitution such that (4) holds. Similarly, θ_2 is the least substitution such that $\forall(\alpha=\tau) \in Q. (\theta_2 \circ \langle Q_2 \rangle)(\alpha) = (\theta_2 \circ \langle Q_2 \rangle)(\tau)$ (5) holds. With (1,5), θ_2 is also the least substitution for $\forall(\alpha=\tau) \in Q. (\theta_2 \circ \langle Q_1 \rangle)(\alpha) = (\theta_2 \circ \langle Q_1 \rangle)(\tau)$ and with (4) it follows that $\theta_1 \equiv \theta_2$ (6). We can now derive:

$$\begin{aligned} & \langle Q \cup Q_1 \rangle \\ = & \theta_1 \circ \langle Q_1 \rangle \{ (2) \} \\ = & \theta_1 \circ \langle Q_2 \rangle \{ (1) \} \\ \equiv & \theta_2 \circ \langle Q_2 \rangle \{ (6) \} \\ = & \langle Q \cup Q_2 \rangle \{ (3) \} \end{aligned}$$

□

Using the replacement Lemma, we can show that we can simplify duplicate bounds:

Proof. (of Theorem 2.7) We have $Q' \vdash \tau_1 \approx \tau_2$ (1) By (1) and Lemma 2.1 and 2.2, we have that $\langle Q' \rangle$ is the least substitution such that $Q'[\tau_1] = Q'[\tau_2]$. Since $\alpha \notin \text{ftv}(Q', \tau_1, \tau_2)$ and Lemma 3.9, we have that $\theta_1 = \langle Q' \cup \{\alpha=\tau_1\} \rangle = \langle Q' \cdot \alpha=\tau_1 \rangle = \langle Q' \rangle \circ [\alpha:=\tau_1]$, and thus θ_1 is a least solution such that $\theta_1(\alpha) = \theta_1(\tau_1) = \theta_1(\tau_2)$. Writing $Q_0 = \{\alpha=\tau_1, \alpha=\tau_2\}$, $\langle Q_0 \rangle$ is also a least solution such that $Q_0[\alpha] = Q_0[\tau_1] = Q_0[\tau_2]$, and we have $\langle Q_0 \rangle \equiv \theta_1$, and thus $Q_0 \equiv Q' \cup \{\alpha=\tau_1\}$. It follows from Lemma D.29 that $Q \cup Q_0 \equiv Q \cup Q' \cup \{\alpha=\tau_1\}$. □

Lemma D.30.

If $\models Q$, we can simplify Q to an equivalent mapping $\lfloor Q \rfloor$.

Proof. (of Lemma D.30) Since $\models Q$, for any duplicate binding $\{\alpha=\tau_1, \alpha=\tau_2\} \subseteq Q$, we have $\langle Q \rangle[\tau_1] = \langle Q \rangle[\tau_2]$ and by Lemma 2.2, $Q' \vdash \tau_1 \approx \tau_2$ (for some $Q' \sqsubseteq Q$), and by Theorem 2.7 we can simplify the duplicate binding to $Q' \cup \{\alpha=\tau_1\}$ (with $\alpha \notin \text{ftv}(Q')$). Repeated application reduces Q to a mapping $\lfloor Q \rfloor$.

Lemma D.31.

If $\models Q$, we can order all bindings in $\lfloor Q \rfloor$ as $\alpha_1=\tau_1 \cdot \dots \cdot \alpha_n=\tau_n$.

Proof. (of Lemma D.31) If there is no possible order in the bounds of a mapping $\lfloor Q \rfloor$, there must a subset $Q' \subseteq \lfloor Q \rfloor$ where for all bounds $\alpha=\tau \in Q'$, $\alpha \in \text{ftv}(Q', \tau)$. Since this is a cyclic dependency, such Q' has no solution (with an idempotent substitution). This contradicts our assumption that Q is consistent.

D.7 Completeness of the Type Rules

We use x_λ for lambda bound variables, and x_{let} for let-bound variables. We can decompose a HM type environment into a substitution and HMQ environment where all lambda-bound variables are

bound to a (fresh) variable:

$$\begin{aligned} \emptyset &\cong id(\emptyset) \\ \Gamma, x_\lambda : \tau &\cong (\theta \circ [\alpha := \tau])(\Gamma', x_\lambda : \alpha) \text{ where } \Gamma \cong \theta\Gamma', \text{ fresh } \alpha \quad [\text{SPLIT-MONO}] \\ \Gamma, x_{\text{let}} : \theta\sigma &\cong \theta(\Gamma', x_{\text{let}} : \sigma) \quad \text{where } \Gamma \cong \theta\Gamma' \quad [\text{SPLIT-POLY}] \end{aligned}$$

Note the $\theta\sigma$ term in the let-case is to allow the type σ in HM Q to have specific sharing that was substituted in the HM system. Consider for example $\lambda x. \text{let } y = (x, 1) \text{ in } ..$ where the HM environment may be instantiated to $x : \text{int}, y : (\text{int}, \text{int})$ while in HM Q we have $x : \alpha, y : (\alpha, \text{int})$. This precision is needed in the inductive let-case of the completeness proof.

We can now state completeness as:

Theorem D.32. (*Completeness & most general typings*)

If $\Gamma \vdash_{\text{HM}} e : \sigma$ with $\Gamma \cong \theta_0\Gamma'$, then we also have $\Delta \mid Q \mid \Gamma' \vdash e : \sigma'$ for some $\Delta \not\vdash \text{ftv}(\Gamma')$, and $\theta = \theta' \circ \theta_0$ with $\text{dom}(\theta') \subseteq \Delta$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma')$, such that $Q \sqsubseteq \theta$ and $\theta\sigma' \sqsubseteq \sigma$.

For the inductive proof we have strengthened the earlier Theorem 2.6 to decompose Γ into θ_0 and Γ' , where θ' only substitutes fresh variables. The invariant $\text{dom}(\theta') \subseteq \Delta$ is used in the [APP] and [LET] cases to combine the substitutions from both sub derivations where we can then show that $\theta'_1 \circ \theta'_2 = \theta'_2 \circ \theta'_1$ which shows that each sub derivation is independent (and there is no sharing of constraints between them).

We use the more general substitution Lemma from HM in the completeness proof [Wright and Felleisen 1994, Lemma 4.6]:

Lemma D.33. (*More general substitution of HM*)

If $\Gamma, x : \sigma \vdash_{\text{HM}} e : \tau$ and $\sigma' \sqsubseteq \sigma$ then $\Gamma, x : \sigma' \vdash_{\text{HM}} e : \tau$.

Proof. (of Theorem D.32) We have $\Gamma \vdash_{\text{HM}} e : \sigma$ (1) with $\Gamma \cong \theta_0\Gamma'$ (1a), and we need to show $\Delta \mid Q \mid \Gamma' \vdash e : \sigma'$ (I) for some $\theta = \theta' \circ \theta_0$ with $Q \sqsubseteq \theta$ (II), and $\theta\sigma' \sqsubseteq \sigma$ (III) (and $\text{dom}(\theta') \subseteq \Delta$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma')$ (IV)).

We proceed by induction on the typing rules $\Gamma \vdash_{\text{HM}} e : \sigma$ and the size of the expression e :

Case [VAR_{HM}]: We have $x : \sigma \in \Gamma$ (2). With $\theta = \theta_0$ (3) (and $\theta' = id$ (4)), we consider two cases: $\sigma = \tau$ with $x_\lambda : \tau \in \Gamma$, and $\sigma = \theta_0\sigma'$ with $x_{\text{let}} : \theta_0\sigma' \in \Gamma$. For a let-bound x_{let} :

$$\begin{aligned} &x : \sigma' \in \Gamma' \quad \{ [\text{SPLIT-POLY}] \} \\ \Rightarrow &\emptyset \mid \emptyset \mid \Gamma' \vdash x : \sigma' \quad \{ [\text{VAR}], \text{ (I)} \} \end{aligned}$$

where $\emptyset \sqsubseteq \theta$ (II) holds by definition, and $\theta\sigma' = \theta_0\sigma' = \sigma$ (3, [SPLIT-POLY]) and thus $\theta\sigma' \sqsubseteq \sigma$ (III). Otherwise, a lambda-bound $x_\lambda : \tau \in \Gamma$, and thus $x_\lambda : \alpha \in \Gamma'$ with $\theta\alpha = \tau$ (5) [SPLIT-MONO], and we can derive:

$$\begin{aligned} &x : \alpha \in \Gamma' \quad \{ [\text{SPLIT-MONO}] \} \\ \Rightarrow &\emptyset \mid \Gamma' \vdash x : \alpha \quad \{ [\text{VAR}], \text{ (I)} \} \end{aligned}$$

where $\emptyset \sqsubseteq \theta$ (II) always holds, and by (5) $\theta\alpha \sqsubseteq \tau$ (III). For both cases, by (4), $\text{dom}(\theta') = \emptyset \subseteq \Delta$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma)$ (IV).

Case [FUN_{HM}]: We have $\Gamma, x : \tau_1 \vdash_{\text{HM}} e : \tau_2$ with $\Gamma, x : \tau_1 \cong (\theta_0 \circ [\alpha := \tau_1])(\Gamma', x : \alpha)$ (2) (with a fresh $\alpha \notin (\Delta \cup \text{ftv}(\Gamma, \Gamma', \tau_2, \tau'_2))$ (2a)), $(\theta_0 \circ [\alpha := \tau_1]) \alpha = \tau_1$ (2b), and by induction $\Delta \mid Q \mid \Gamma', x : \alpha \vdash e : \tau'_2$ (3) with $\theta = \theta'' \circ (\theta_0 \circ [\alpha := \tau_1])$ where $\theta' = \theta'' \circ [\alpha := \tau_1]$ (3a), $\text{dom}(\theta'') \subseteq \Delta$ and $\text{codom}(\theta'') \subseteq \text{ftv}(Q, \tau'_2)$ (3b), such that $Q \sqsubseteq \theta$ (3c) and $\theta\tau'_2 \sqsubseteq \tau_2$ (3d). We can now derive:

$$\begin{aligned} &\Delta \mid Q \mid \Gamma', x : \alpha \vdash e : \tau'_2 \quad \{ (3) \} \\ \Rightarrow &\Delta, \alpha \mid Q \mid \Gamma' \vdash \lambda x. e : \alpha \rightarrow \tau'_2 \quad \{ [\text{FUN}], (2a), \text{ (I)} \} \end{aligned}$$

From (3c), we have directly $Q \sqsubseteq \theta$ (II).

$$\begin{aligned}
& \theta(\alpha \rightarrow \tau'_2) \\
= & \theta\alpha \rightarrow \theta\tau'_2 & \{ \text{subst.} \} \\
\sqsubseteq & \theta\alpha \rightarrow \tau_2 & \{ (3d) \} \\
= & \theta''((\theta_0 \circ [\alpha := \tau_1])\alpha) \rightarrow \tau_2 & \{ (3a) \} \\
= & \theta''\tau_1 \rightarrow \tau_2 & \{ (2b) \} \\
= & \tau_1 \rightarrow \tau_2 & \{ (3b), \text{ (III)} \}
\end{aligned}$$

From (3b), it follows $\text{dom}(\theta') \subseteq (\Delta, \alpha)$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \alpha \rightarrow \tau'_2)$ (IV)

Case [APP_{HM}]: By (1), we have $\Gamma \vdash_{\text{HM}} e_1 : \tau_2 \rightarrow \tau$ (2a) and $\Gamma \vdash e_2 : \tau_2$ (2b), and by induction $\Delta_1 \mid Q_1 \mid \Gamma' \vdash e_1 : \tau_3$ (3a) and $\Delta_2 \mid Q_2 \mid \Gamma' \vdash e_2 : \tau_4$ (4a), where we pick $\Delta_1 \not\sqsubseteq \Delta_2$ (2c). For the first derivation, we have $\theta_1 = \theta'_1 \circ \theta_0$, $Q_1 \sqsubseteq \theta_1$, $\theta_1\tau_3 \sqsubseteq \tau_2 \rightarrow \tau$, $\text{dom}(\theta'_1) \subseteq \Delta_1$, and $\text{codom}(\theta'_1) \subseteq \text{ftv}(Q_1, \tau_3)$ (3b). For the second derivation, $\theta_2 = \theta'_2 \circ \theta_0$, $Q_2 \sqsubseteq \theta_2$, $\theta_2\tau_4 \sqsubseteq \tau_2$, $\text{dom}(\theta'_2) \subseteq \Delta_2$, and $\text{codom}(\theta'_2) \subseteq \text{ftv}(Q_2, \tau_4)$ (4b).

We can now derive:

$$\begin{aligned}
& \text{codom}(\theta'_1) \subseteq \text{ftv}(Q_1, \tau_3) \quad \{ (3b) \} \\
\Rightarrow & \text{codom}(\theta'_1) \subseteq \text{ftv}(\Delta_1, \Gamma') \quad \{ \text{Lemma 8.10} \} \\
\Rightarrow & \text{codom}(\theta'_1) \not\sqsubseteq \Delta_2 \quad \{ (2c), \Delta_2 \not\sqsubseteq \text{ftv}(\Gamma'). \}
\end{aligned}$$

and similarly $\text{codom}(\theta'_2) \not\sqsubseteq \Delta_1$. Moreover $\text{dom}(\theta'_1) \subseteq \Delta_1$ (3b) implies with (2c) that $\text{dom}(\theta'_1) \not\sqsubseteq \Delta_2$. A similar argument holds for θ'_2 and therefore the (co)domains of θ'_1 and θ'_2 are disjoint, and thus $\theta'_1 \circ \theta'_2 = \theta'_2 \circ \theta'_1$ (5). As an aside, this property is why we need fresh names as otherwise sharing between the two sub-derivations can occur and (5) would not hold.

We can now define $\theta' = \theta'_1 \circ \theta'_2 \circ \theta_0$ (2d), and derive:

$$\begin{aligned}
& Q_2 \sqsubseteq \theta_2 \quad \{ (4b) \} \\
= & Q_2 \sqsubseteq \theta'_2 \circ \theta_0 \quad \{ (4b) \} \\
\Rightarrow & Q_2 \sqsubseteq \theta'_1 \circ \theta'_2 \circ \theta_0 \quad \{ \text{def.} \} \\
= & Q_2 \sqsubseteq \theta' \quad \{ (2d), (4c) \}
\end{aligned}$$

Since $\text{dom}(\theta'_1) \not\sqsubseteq \text{codom}(\theta_2)$ (4b), we have $\theta'\tau_4 = \theta_2\tau_4$ (4b) and thus $\theta'\tau_4 \sqsubseteq \tau_2$ (4d). By (5), we also have $\theta' = \theta'_2 \circ \theta'_1 \circ \theta_0$, and we can use the same reasoning for the left derivation to conclude $Q_1 \sqsubseteq \theta'$ (3c) and $\theta'\tau_3 \sqsubseteq \tau_2 \rightarrow \tau$ (3d).

From (3d,4d) it follows that $\theta'\tau_3 = \tau_2 \rightarrow \tau$ and $\theta'\tau_4 = \tau_2$ since these are monotypes. Therefore, $\theta'\tau_3 = \theta'\tau_4 \rightarrow \tau$. We now define $\theta = [\alpha := \tau] \circ \theta'$ for some fresh $\alpha \notin \Delta_1, \Delta_2$ (6). Since α is fresh and Lemma 8.10, we have $\theta\tau_3 = \theta(\tau_4 \rightarrow \alpha)$, and from Theorem 2.2 it follows $Q_3 \vdash \tau_3 \approx \tau_4 \rightarrow \alpha$ (7), with $Q_3 \sqsubseteq \theta$ (7a).

From (3c,6), we have $Q_1 \sqsubseteq \theta$ and by (4c,6) $Q_2 \sqsubseteq \theta$. With (7a), and Lemma D.21 we have Q_1, Q_2, Q_3 is consistent with $Q_1, Q_2, Q_3 \sqsubseteq \theta$ (II). Together with (3a,4a,7), we can now use [APP] to conclude $\Delta_1, \Delta_2, \alpha \mid Q_1, Q_2, Q_3 \mid \Gamma' \vdash e_1 e_2 : \alpha$ (I) and by (6), $\theta\alpha \sqsubseteq \tau$ (III). Finally, by (3b,4b,6), we also have $\text{dom}(\theta) \subseteq \Delta_1, \Delta_2, \alpha$ with $\text{codom}(\theta) \subseteq \text{ftv}((Q_1, Q_2, \alpha), \alpha)$ (IV).

Case [INST_{HM}]: We have $\Gamma \vdash_{\text{HM}} e : \forall \alpha. \sigma$ with $\Gamma \cong \theta_0\Gamma'$ (2), and thus by induction $\Delta \mid Q \mid \Gamma' \vdash e : \sigma'$ (3) with $\theta = \theta' \circ \theta_0$ (3a) such that $Q \sqsubseteq \theta$ (3b) and $\theta\sigma' \sqsubseteq \forall \alpha. \sigma$ (3c), and $\text{dom}(\theta') \subseteq \Delta$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma')$ (3d). We can assume a fresh $\alpha \notin \Delta$ by α -renaming (4).

From (3c) and [INSTANCE], we must have $\sigma' = \forall \alpha. \sigma_0$ (5) (for some σ_0). We can thus derive:

$$\begin{aligned}
& \Delta \mid Q \mid \Gamma' \vdash e : \forall \alpha. \sigma_0 \quad \{ (3, 5) \} \\
\Rightarrow & \Delta, \alpha \mid Q \mid \Gamma' \vdash e : \sigma_0 \quad \{ (4), \text{ (I)} \}
\end{aligned}$$

From (3b), we directly have $Q \sqsubseteq \theta$ (II). Moreover, by (3c,5) $\theta(\forall \alpha. \sigma_0) \sqsubseteq \forall \alpha. \sigma$. Since α is fresh, we must have $\theta\sigma_0 \sqsubseteq \sigma$ (III). Finally, from (3d) and (4), it follows directly that $\text{dom}(\theta') \subseteq \Delta, \alpha$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma')$ (IV).

Case [GEN_{HM}]: We have $\Gamma \vdash_{\text{HM}} e : \sigma$ with $\Gamma \cong \theta_0 \Gamma'$ (2), and $\alpha \notin \text{ftv}(\Gamma)$ (3), and thus by induction $\Delta \mid Q \mid \Gamma' \vdash e : \sigma'$ (3a) with $\theta = \theta' \circ \theta_0$ (3b) such that $Q \sqsubseteq \theta$ (3c) and $\theta \sigma' \sqsubseteq \sigma$ (3d) (and $\Delta \not\vdash \text{ftv}(\Gamma')$ (3e)), $\text{dom}(\theta') \subseteq \Delta$ and $\text{codom}(\theta') \subseteq \text{ftv}(Q, \sigma')$ (3f). From (3) and the definition of (\cong) , we have $\alpha \notin \text{ftv}(\Gamma')$ (4).

With $\alpha \in \text{ftv}(\sigma)$, then from (3d) we have $\alpha \in \text{ftv}(\theta \sigma')$ (5a) and thus $\alpha \notin \text{dom}(\theta)$ (5c). We now have three cases to consider:

- A. Suppose $\alpha \notin \text{ftv}(Q)$, in that case we can apply [GEN] with (3a,4) and derive $\Delta \mid Q \mid \Gamma' \vdash e : \forall \alpha. \sigma'$ (I) with (3c) $Q \sqsubseteq \theta$ (II) and (3d,5c) $\theta(\forall \alpha. \sigma') \sqsubseteq \forall \alpha. \sigma$ (III). By (3f,2), $\text{dom}(\theta') \subseteq \Delta$ and also we have that $\text{codom}(\theta') \subseteq \text{ftv}(Q, \forall \alpha. \sigma')$ (IV).
- B. Suppose we have $\alpha \in \text{dom}(Q)$, in that case by (3c), $\alpha \in \text{dom}(\theta)$ but that contradicts (5c).
- C. Otherwise, we must have $\alpha \in \text{codom}(Q)$. With Theorem D.23, this implies $Q \equiv Q_1 \cdot \beta = \tau$ (6a) with $\alpha \in \text{ftv}(\tau)$ (6b) and $\beta \notin \text{ftv}(Q_1, \tau)$ (6c).

Suppose $\beta \in \text{ftv}(\Gamma')$. If $x_\lambda : \beta \in \Gamma'$, then with (3c) we have $x_\lambda : \tau \in \Gamma$ with $\alpha \in \text{ftv}(\Gamma)$ (6b) which contradicts (3). Otherwise, $x_{\text{let}} : \sigma_0 \in \Gamma'$ with $\beta \in \text{ftv}(\sigma_0)$ which implies by (3b,3c) that $\beta \in \text{dom}(\theta')$, and thus $\beta \in \Delta$ – but that contradicts (3e). Therefore, we must have $\beta \notin \text{ftv}(\Gamma')$ (7). With (6c,7) we can now use [GENSUB] to derive $\Delta \mid Q_1 \mid \Gamma' \vdash e : [\beta := \tau] \sigma'$, where (3c,6a) $Q_1 \sqsubseteq \theta$, and $\theta \circ [\beta := \tau] = \theta$, and thus $\theta[\beta := \tau] \sigma' \sqsubseteq \sigma$. We can repeatedly apply [GENSUB] until case (A) applies.

Case [LET_{HM}]: We have $\Gamma \vdash_{\text{HM}} e_1 : \sigma$ with $\Gamma \cong \theta_0 \Gamma'$ (2), and $\Gamma, x : \sigma \vdash_{\text{HM}} e_2 : \tau$ (3). By induction, we have $\Delta_1 \mid Q_1 \mid \Gamma' \vdash e_1 : \sigma_1$ (4a), and for the left derivation, we have $\theta_1 = \theta'_1 \circ \theta_0$, $Q_1 \sqsubseteq \theta_1$, $\theta_1 \sigma_1 \sqsubseteq \sigma$, $\text{dom}(\theta'_1) \subseteq \Delta_1$, and $\text{codom}(\theta'_1) \subseteq \text{ftv}(Q_1, \sigma_1)$ (4b).

To apply [LET], we need to satisfy the side condition $\text{ftv}(\sigma_1) \subseteq \text{ftv}(\Gamma')$. Suppose there is a $\alpha \in \text{ftv}(\sigma_1)$ with $\alpha \notin \text{ftv}(\Gamma')$. In that case we can apply either [GEN] or [GENSUB] depending on whether $\alpha \in \text{ftv}(Q_1)$. If $\alpha \in \text{ftv}(Q_1)$, we have by Lemma 2.8, $Q_1 \equiv Q_3 \cdot \alpha = \tau_3$ (5), and we can apply [GENSUB] to derive $\Delta_1 \mid Q_3 \mid \Gamma' \vdash e_1 : \sigma_3$ with $\sigma_3 = [\alpha := \tau_3] \sigma_1$. From (5) and Lemma 3.9, we have $\langle Q_1 \rangle = \langle Q_3 \rangle \circ [\alpha := \tau_3]$. Since $\langle Q_1 \rangle \sqsubseteq \theta_1$ (4b), we also have $\theta_1 = \theta_1 \circ [\alpha := \tau_3]$ (5a). We can now derive:

$$\begin{aligned} & \theta_1 \sigma_1 \sqsubseteq \sigma & \{ (4b) \} \\ = & \theta_1 ([\alpha := \tau_3] \sigma_1) \sqsubseteq \sigma & \{ (5a) \} \\ = & \theta_1 \sigma_3 \sqsubseteq \sigma & \{ (5b) \} \end{aligned}$$

Similarly, if $\alpha \notin \text{ftv}(Q_1)$ we can apply the [GEN] rule to derive $\Delta_1 \mid Q_1 \mid \Gamma' \vdash e_1 : \sigma_3$ with $\sigma_3 = \forall \alpha. \sigma_1$. By (4b), we have $\theta_1 \sigma_3 \sqsubseteq \sigma$ (5b).

Therefore, after repeated application we have $\Delta_1 \mid Q_n \mid \Gamma' \vdash e_1 : \sigma_n$ with $Q_n \sqsubseteq \theta_1$ and $\theta_1 \sigma_n \sqsubseteq \sigma$ (5c). Furthermore, since now $\text{ftv}(\sigma_n) \subseteq \text{ftv}(\Gamma')$, we have $\theta_1 \sigma_n = \theta_0 \sigma_n \sqsubseteq \sigma$ (5d).

By (3), $\Gamma, x : \sigma \vdash_{\text{HM}} e_2 : \tau$. With (5d), and Lemma D.33, we also have $\Gamma, x : \theta_0 \sigma_n \vdash_{\text{HM}} e_2 : \tau$, and by induction, $\Delta_2 \mid Q_2 \mid \Gamma', x : \sigma_n \vdash e_2 : \tau'$ (6a) with $\theta_2 = \theta'_2 \circ \theta_0$, $Q_2 \sqsubseteq \theta_2$, $\theta_2 \tau' \sqsubseteq \tau$, $\text{dom}(\theta'_2) \subseteq \Delta_2$, and $\text{codom}(\theta'_2) \subseteq \text{ftv}(Q_2, \tau')$ (6b).

We now proceed as in the [APP_{HM}] case, and define $\theta = \theta'_1 \circ \theta'_2 \circ \theta_0$ (7) with $\theta'_1 \circ \theta'_2 = \theta'_2 \circ \theta'_1$ (7a). Just as in the [APP_{HM}] case, it follows that $Q_n \sqsubseteq \theta$ (7c), $Q_2 \sqsubseteq \theta$ (7d), and and, since $\text{dom}(\theta'_1) \subseteq \Delta_1$ and (6b), $\theta \tau' \sqsubseteq \tau$ (III). From (7c,7d) it follows from Lemma D.21 that Q_n, Q_2 is consistent with $Q_n, Q_2 \sqsubseteq \theta$ (II).

We can now finally apply [LET] to derive $\Delta_1, \Delta_2 \mid Q_n, Q_2 \mid \Gamma' \vdash \text{let } x = e_1 \text{ in } e_2 : \tau'$ (I), (and like the [APP_{HM}] case, $\text{dom}(\theta) \subseteq \text{ftv}(\Delta_1, \Delta_2)$ and $\text{codom}(\theta) \subseteq \text{ftv}((Q_n, Q_2), \Gamma')$ (IV)).

□

D.8 Soundness of the Type Rules

Theorem 2.5. (Soundness)

If $\Delta \mid Q \mid \Gamma \vdash e : \sigma$, then we can also derive $Q[\Gamma] \vdash_{\text{HM}} e : Q[\sigma]$.

We use the substitution Lemma from HM in the soundness proof:

Lemma D.35. (Weakening of HM)

If $\Gamma \vdash_{\text{HM}} e : \sigma$ then also $\theta\Gamma \vdash_{\text{HM}} e : \theta\sigma$.

We often use this Lemma when $Q_1 \sqsubseteq Q$, and $Q_1[\Gamma] \vdash_{\text{HM}} e : Q_1[\sigma]$, then we also have $Q[\Gamma] \vdash_{\text{HM}} e : Q[\sigma]$.

Proof. (of Theorem 2.5) By induction on the typing rules of $\Delta \mid Q \mid \Gamma \vdash e : \sigma$:

Case [VAR]: We have $x : \sigma \in \Gamma$ (1) and $Q = \emptyset$ (2), and thus $Q[\Gamma] = \Gamma$ (3) and $Q[\sigma] = \sigma$ (4).

$$\begin{aligned} & x : \sigma \in \Gamma && \{ (1) \} \\ \Rightarrow & \Gamma \vdash_{\text{HM}} x : \sigma && \{ [\text{VAR}_{\text{HM}}] \} \\ = & Q[\Gamma] \vdash_{\text{HM}} x : \Gamma[\sigma] && \{ (3, 4) \} \end{aligned}$$

Case [FUN]: We have $\Delta \mid Q \mid \Gamma, x : \alpha \vdash e : \tau$ (1) (with $\alpha \notin \Delta$). We can now derive:

$$\begin{aligned} & Q[\Gamma, x : \alpha] \vdash_{\text{HM}} e : Q[\tau] && \{ \text{induction over (1)} \} \\ = & Q[\Gamma], x : Q[\alpha] \vdash_{\text{HM}} e : Q[\tau] && \{ \text{def.} \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} \lambda x. e : Q[\alpha] \rightarrow Q[\tau] && \{ [\text{FUN}_{\text{HM}}] \} \\ = & Q[\Gamma] \vdash_{\text{HM}} \lambda x. e : Q[\alpha \rightarrow \tau] && \{ \text{def.} \} \end{aligned}$$

Case [APP]: We have $\Delta_1 \mid Q_1 \mid \Gamma \vdash e_1 : \tau_1$ (1) and $\Delta_2 \mid Q_2 \mid \Gamma \vdash e_2 : \tau_2$ (2) with $Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha$ (3) and $\vdash Q_1, Q_2, Q_3$ (4). Writing $Q = Q_1, Q_2, Q_3$, we have by Lemma D.18, $Q_1 \sqsubseteq Q$, $Q_2 \sqsubseteq Q$, and $Q_3 \sqsubseteq Q$ (6).

$$\begin{aligned} & Q_1[\Gamma] \vdash_{\text{HM}} e_1 : Q_1[\tau_1] && \{ \text{induction over (1)} \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} e_1 : Q[\tau_1] && \{ \text{Lemma D.35, (6), (7)} \} \end{aligned}$$

and also:

$$\begin{aligned} & Q_2[\Gamma] \vdash_{\text{HM}} e_2 : Q_2[\tau_2] && \{ \text{induction over (2)} \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} e_2 : Q[\tau_2] && \{ \text{Lemma D.35, (6), (8)} \} \end{aligned}$$

Furthermore:

$$\begin{aligned} & Q_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha && \{ (3) \} \\ \Rightarrow & Q_3[\tau_1] = Q_3[\tau_2 \rightarrow \alpha] && \{ \text{Theorem 2.1} \} \\ \Rightarrow & Q[\tau_1] = Q[\tau_2 \rightarrow \alpha] && \{ (6) \} \\ \Rightarrow & Q[\tau_1] = Q[\tau_2] \rightarrow Q[\alpha] && \{ \text{def.} \} \end{aligned}$$

and by (7), $Q[\Gamma] \vdash_{\text{HM}} e_1 : Q[\tau_2] \rightarrow Q[\alpha]$, and with $[\text{APP}_{\text{HM}}]$ and (8), we have $Q[\Gamma] \vdash_{\text{HM}} e_1 e_2 : Q[\alpha]$.

Case [GENSUB]: We have $\Delta \mid Q \cdot \alpha = \tau \mid \Gamma \vdash e : \sigma$ (1) with $\alpha \notin \text{ftv}(Q, \Gamma)$ (2). Writing Q' for $Q \cdot \alpha = \tau$, we can derive:

$$\begin{aligned} & Q'[\Gamma] \vdash_{\text{HM}} e : Q'[\sigma] && \{ \text{induction over (1)} \} \\ = & Q[[\alpha := \tau]\Gamma] \vdash_{\text{HM}} e : Q[[\alpha := \tau]\sigma] && \{ \text{Lemma 3.9} \} \\ = & Q[\Gamma] \vdash_{\text{HM}} e : Q[[\alpha := \tau]\sigma] && \{ (2) \} \end{aligned}$$

Case [GEN]: We have $\Delta \mid Q \mid \Gamma \vdash e : \sigma$ (1) with $\alpha \notin \text{ftv}(Q, \Gamma)$ (2). We can derive:

$$\begin{aligned} & Q[\Gamma] \vdash_{\text{HM}} e : Q[\sigma] && \{ \text{induction over (1)} \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} e : \forall \alpha. Q[\sigma] && \{ [\text{GEN}_{\text{HM}}], (2) \} \\ = & Q[\Gamma] \vdash_{\text{HM}} e : Q[\forall \alpha. \sigma] && \{ (2) \} \end{aligned}$$

Case [INST]: We have $\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma$ (1) with $\alpha \notin \Delta$ (2), and $\Delta, \alpha \not\vdash \text{ftv}(Q, \Gamma)$ and thus $\alpha \notin \text{ftv}(Q, \Gamma)$ (3). We can derive:

$$\begin{aligned} & Q[\Gamma] \vdash_{\text{HM}} e : Q[\forall \alpha. \sigma] \quad \{ \text{induction over (1)} \} \\ = & Q[\Gamma] \vdash_{\text{HM}} e : \forall \alpha. Q[\sigma] \quad \{ (3) \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} e : [\alpha := \alpha](Q[\sigma]) \quad \{ [\text{INST}_{\text{HM}}], (3) \} \\ = & Q[\Gamma] \vdash_{\text{HM}} e : Q[\sigma] \quad \{ \text{def.} \} \end{aligned}$$

Case [LET]: We have $\Delta_1 \mid Q_1 \mid \Gamma \vdash e_1 : \sigma$ (1) and $\Delta_2 \mid Q_2 \mid \Gamma, x : \sigma \vdash e_2 : \tau$ (2) with $\text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$ (3) and $\models Q_1, Q_2$ (4). Writing $Q = Q_1, Q_2$, by Lemma D.18, we also have $Q_1 \sqsubseteq Q$ and $Q_2 \sqsubseteq Q$ (5). We can derive:

$$\begin{aligned} & Q_1[\Gamma] \vdash_{\text{HM}} e_1 : Q_1[\sigma] \quad \{ \text{ind. over (1)} \} \\ \Rightarrow & Q[\Gamma] \vdash_{\text{HM}} e_1 : Q[\sigma] \quad \{ \text{Lemma D.35, (5), (6)} \} \end{aligned}$$

and similarly:

$$\begin{aligned} & Q_2[\Gamma, x : \sigma] \vdash_{\text{HM}} e_2 : Q_2[\tau] \quad \{ \text{ind. over (2)} \} \\ \Rightarrow & Q[\Gamma, x : \sigma] \vdash_{\text{HM}} e_2 : Q[\tau] \quad \{ \text{Lemma D.35 (5), (7)} \} \end{aligned}$$

We can now apply [LET_{HM}] to derive $Q[\Gamma] \vdash_{\text{HM}} \text{let } x = e_1 \text{ in } e_2 : Q[\tau]$.

□

D.9 Principal Type Derivations

We can prove Theorem 2.3 and 2.4 also by showing that we can rewrite every derivation into an equivalent (syntax-directed) canonical derivation.

D.9.1 Float Up Instantiations. We now show that we can always float any [INST] rules in a derivation up to the ([VAR]) leaves of a derivation.

Proof. For [INST] to apply, it must follow a derivation $\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma$. This must be [VAR], [INST], [GENSUB], or [GEN] (since [INST] applies to type schemes).

Case [GEN]: We have:

$$\frac{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q)}{\frac{\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma} \text{INST}} \text{GEN}$$

which cancel out, and is directly equivalent to $\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma$.

Case [GENSUB]: We have:

$$\frac{\frac{\Delta, \beta \mid Q \cdot \beta = \tau \mid \Gamma \vdash e : \forall \alpha. \sigma}{\Delta \mid Q \mid \Gamma \vdash e : [\beta := \tau](\forall \alpha. \sigma) = \forall \alpha. [\beta := \tau]\sigma} \text{GENSUB}}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : [\beta := \tau]\sigma} \text{INST}$$

where we use α -renaming to ensure $\alpha \notin \text{ftv}(\tau, \beta)$. This means we can rewrite this to the equivalent:

$$\frac{\frac{\Delta, \beta \mid Q \cdot \beta = \tau \mid \Gamma \vdash e : \forall \alpha. \sigma}{\Delta, \beta, \alpha \mid Q \cdot \beta = \tau \mid \Gamma \vdash e : \sigma} \text{INST}}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : [\beta := \tau]\sigma} \text{GENSUB}$$

Case [INST]: we have

$$\frac{\frac{\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha \beta. \sigma}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \forall \beta. \sigma} \text{INST}}{\Delta, \alpha, \beta \mid Q \mid \Gamma \vdash e : \sigma} \text{INST}$$

Since $\forall \alpha \beta. \sigma = \forall \beta \alpha. \sigma$, the order does not matter and we also have:

$$\frac{\frac{\Delta \mid Q \mid \Gamma \vdash e : \forall \beta \alpha. \sigma}{\Delta, \beta \mid Q \mid \Gamma \vdash e : \forall \alpha. \sigma} \text{INST}}{\Delta, \beta, \alpha \mid Q \mid \Gamma \vdash e : \sigma} \text{INST}$$

Case [VAR]: we have:

$$\frac{\frac{x : \forall \alpha. \sigma \in \Gamma}{\emptyset \mid \emptyset \mid \Gamma \vdash x : \forall \alpha. \sigma} \text{VAR}}{\{\alpha\} \mid \emptyset \mid \Gamma \vdash e : \sigma} \text{INST}$$

We cannot float up any further since [VAR] is a leaf rule.

□

This means that we can rewrite any derivation into one where all instantiations directly follow a [VAR] rule (and where the order of those instantiations does not matter).

Therefore, we can also create a syntax directed rule for variables as:

$$\frac{x : \forall \bar{\alpha}. \tau \in \Gamma}{\bar{\alpha} \mid \emptyset \mid \Gamma \vdash x : \tau} \text{VAR}_s$$

and replace any [VAR] followed by a sequence of [INST] applications by [VAR]_s

Proof. Suppose we have an application of [VAR] followed by n [INST] applications:

$$\frac{\frac{x : \forall \bar{\alpha} \bar{\beta}. \tau \in \Gamma}{\vdots} \text{VAR}}{\bar{\alpha} \mid \emptyset \mid \Gamma \vdash x : \forall \bar{\beta}. \tau} \text{INST}$$

If $\bar{\beta} = \emptyset$ we can use [VAR]_s directly. Otherwise, it must be followed by a [MGEN], [GENSUB], or [GEN] rule (as it is a type scheme). Since $Q = \emptyset$, it cannot be [GENSUB]. If it is followed by [MGEN] or [GEN], we instead apply a sequence of [INST] on all $\bar{\beta}$ such that we end in a monotype and generalize that again. We can thus rewrite to [VAR]_s followed by a sequence of [GEN] to restore the $\bar{\beta}$ bindings. □

For any derivation, we can thus float up all [INST] rules to [VAR], and subsequently replace them with [VAR]_s (and [GEN]) rules such that there are no more [INST] rules left.

D.9.2 Push Down Generalizations. Similarly to instantiations, we can show that we can always push down generalizations to the end of a derivation.

Proof. If [GEN] is at the end of a derivation we are done. Otherwise, the [GEN] rule must be followed by an [INST], [GENSUB] or [GEN] rule (as it results in a type scheme).

Case [GEN]: we have:

$$\frac{\frac{\Delta, \alpha, \beta \mid Q \mid \Gamma \vdash e : \sigma \quad \beta \notin \text{ftv}(Q)}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \forall \beta. \sigma \quad \alpha \notin \text{ftv}(Q)} \text{GEN}}{\Delta \mid Q \mid \Gamma \vdash e : \forall \alpha \beta. \sigma} \text{GEN}$$

Since $\forall\alpha\beta.\sigma = \forall\beta\alpha.\sigma$, the order does not matter and we can also derive:

$$\frac{\frac{\Delta, \alpha, \beta \mid Q \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q)}{\Delta, \beta \mid Q \mid \Gamma \vdash e : \forall\alpha.\sigma} \text{GEN}}{\Delta \mid Q \mid \Gamma \vdash e : \forall\beta\alpha.\sigma} \text{GEN}$$

Case [INST]: we have:

$$\frac{\frac{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q)}{\Delta \mid Q \mid \Gamma \vdash e : \forall\alpha.\sigma} \text{GEN}}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma} \text{INST}$$

These cancel out and we have directly $\Delta, \alpha \mid Q \mid \Gamma \vdash e : \sigma$.

Case [GENSUB]: we have:

$$\frac{\frac{\Delta, \beta, \alpha \mid Q \cdot \beta=\tau \mid \Gamma \vdash e : \sigma \quad \alpha \notin \text{ftv}(Q \cdot \beta=\tau)}{\Delta, \beta \mid Q \cdot \beta=\tau \mid \Gamma \vdash e : \forall\alpha.\sigma} \text{GEN}}{\Delta \mid Q \mid \Gamma \vdash e : [\beta=\tau](\forall\alpha.\sigma)} \text{GENSUB}$$

By the pre-condition on [GEN], we have $\alpha \notin \text{ftv}(\tau, \beta)$, and thus $[\beta=\tau](\forall\alpha.\sigma) = \forall\alpha. [\beta=\tau]\sigma$, and we can rewrite this as:

$$\frac{\frac{\Delta, \beta, \alpha \mid Q \cdot \beta=\tau \mid \Gamma \vdash e : \sigma}{\Delta, \alpha \mid Q \mid \Gamma \vdash e : [\beta=\tau]\sigma} \text{GENSUB}}{\Delta \mid Q \mid \Gamma \vdash e : \forall\alpha. [\beta=\tau]\sigma} \text{GEN}$$

□

This means we can always rewrite a derivation into one where all generalizations happen at the end of a derivation (and the order of those does not matter).

D.9.3 Pushing Down Substitutions. The [GENSUB] rule takes a type scheme in the premise, meaning it can be preceded by [VAR], [INST], [GEN], or [GENSUB]. However, it cannot be [VAR] since $Q = \emptyset$ in such case. Moreover, in the previous section we also showed that we can always rewrite a derivation to use [VAR_s] without any needing [INST] rules. Finally, since [GEN] can always be pushed down to the end of a derivation this means that any remaining [GENSUB] rules always apply to mono-types, and we only need to consider:

$$\frac{\Delta, \alpha \mid Q \cdot \alpha=\tau' \mid \Gamma \vdash e : \tau}{\Delta \mid Q \mid \Gamma \vdash e : [\alpha:=\tau']\tau} \text{GENSUB-MONO}$$

Just like generalizations, we can now show we can push down these remaining [GENSUB-MONO] rules to the end of a derivation (just before any [GEN] rules).

Proof. If [GENSUB-MONO] is followed by a [GEN], [GENSUB-MONO], or at the end of a derivation already, we are done. Otherwise, since [GENSUB-MONO] results in a mono-type, it must be followed by [FUN], [APP], or (right) [LET].

Case [FUN]: we have

$$\frac{\frac{\Delta, \beta \mid Q \cdot \beta=\tau' \mid \Gamma, x:\alpha \vdash e : \tau}{\Delta \mid Q \mid \Gamma, x:\alpha \vdash e : [\beta=\tau']\tau} \text{GENSUB-MONO}}{\Delta, \alpha \mid Q \mid \Gamma \vdash \lambda x.e : \alpha \rightarrow [\beta=\tau']\tau} \text{FUN}$$

Since $\beta \notin \text{ftv}(Q, (\Gamma, x : \alpha))$, we have $\beta \neq \alpha$, and thus $\alpha \rightarrow [\beta=\tau']\tau = [\beta=\tau'](\alpha \rightarrow \tau)$, and we can also derive:

$$\frac{\frac{\Delta, \beta \mid Q \cdot \beta=\tau' \mid \Gamma, x : \alpha \vdash e : \tau}{\Delta, \beta, \alpha \mid Q \cdot \beta=\tau' \mid \Gamma \vdash \lambda x. e : \alpha \rightarrow \tau} \text{FUN}}{\Delta, \alpha \mid Q \mid \Gamma \vdash \lambda x. e : [\beta=\tau'](\alpha \rightarrow \tau)} \text{GENSUB-MONO}$$

Case [LET]: We have:

$$\frac{\frac{\Delta_1 \mid Q_1 \mid \Gamma \Vdash e_1 : \sigma}{\Delta_1, \alpha \mid Q_2 \cdot \alpha=\tau' \mid \Gamma, x : \sigma \vdash e_2 : \tau} \text{GENSUB-MONO}}{\Delta_1, \Delta_2 \mid Q_1, Q_2 \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : [\alpha=\tau']\tau} \text{LET}$$

We can directly rewrite this to:

$$\frac{\frac{\Delta_1 \mid Q_1 \Vdash e_1 : \sigma \quad \Delta_2, \alpha \mid Q_2 \cdot \alpha=\tau' \mid \Gamma, x : \sigma \vdash e_2 : \tau}{\Delta_1, \Delta_2, \alpha \mid Q_1, Q_2 \cdot \alpha=\tau' \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau} \text{LET}}{\Delta_1, \Delta_2 \mid Q_1, Q_2 \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : [\alpha=\tau']\tau} \text{GENSUB-MONO}$$

Case [APP] (left): We have,

$$\frac{\frac{\Delta_1, \beta \mid Q_1 \cdot \beta=\tau \mid \Gamma \vdash e_1 : \tau_1}{\Delta_1 \mid Q_1 \mid \Gamma \vdash e_1 : [\beta=\tau]\tau_1} \text{GENSUB-MONO}}{\Delta_2 \mid Q_2 \mid \Gamma \vdash e_2 : \tau_2 \quad Q_3 \vdash [\beta=\tau]\tau_1 \approx \tau_2 \rightarrow \alpha} \text{APP}$$

Since $\text{ftv}(Q_2, \tau_2) \subseteq \text{ftv}(\Delta_1, \Gamma)$, $\beta \notin \text{ftv}(Q_1, \Gamma)$, and $\Delta_1 \not\sqcap \Delta_2$, we must have $\beta \notin \text{ftv}(\tau_2)$. Therefore $\tau_2 \rightarrow \alpha = [\beta=\tau](\tau_2 \rightarrow \alpha)$, and by Lemma D.27, we have $Q'_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha$ with $(Q_3, \{\beta=\tau\}) = (Q'_3, \{\beta=\tau\})$ (1). We can then derive:

$$\begin{aligned} & (Q_1 \cdot \beta=\tau), Q_2, Q'_3 \\ = & (Q_1, \beta=\tau), Q_2, Q'_3 \\ = & Q_1, Q_2, Q'_3, \beta=\tau \\ = & Q_1, Q_2, Q_3, \beta=\tau \quad \{ (1) \} \\ = & Q_1, Q_2, (Q_3 \cdot \beta=\tau') \quad \{ \text{Theorem D.23, for some } \tau' \} \\ = & (Q_1, Q_2, Q_3) \cdot \beta=\tau' \quad \{ (2), \beta \notin \text{ftv}(Q_1, Q_2, Q_3) \} \end{aligned}$$

With this we can push down the [GENSUB-MONO] as:

$$\frac{\frac{\Delta_1, \beta \mid Q_1 \cdot \beta=\tau \mid \Gamma \vdash e_1 : \tau_1 \quad \Delta_2 \mid Q_2 \mid \Gamma \vdash e_2 : \tau_2 \quad Q'_3 \vdash \tau_1 \approx \tau_2 \rightarrow \alpha}{\Delta_1, \Delta_2, \beta, \alpha \mid Q_1 \cdot \beta=\tau, Q_2, Q'_3 \mid \Gamma \vdash e_1 e_2 : \alpha} \text{APP}}{= \Delta_1, \Delta_2, \alpha, \beta \mid (Q_1, Q_2, Q_3) \cdot \beta=\tau' \mid \Gamma \vdash e_1 e_2 : \alpha \quad (\beta \notin \text{ftv}(Q_1, Q_2, Q_3, \Gamma))} \text{GENSUB-MONO}$$

Case [APP] (right): same as the previous case.

D.9.4 Trivial Generalizations. Before we can prove the main Theorem of principality, we are first going to use a stronger theorem that is independent of *trivial* generalizations. We call a [GENSUB] rule trivial if $\alpha \notin \text{ftv}(\sigma)$ since such $[\alpha=\tau]\sigma$ leaves the derived type σ unchanged. These present an issue for the proof of principality though (Theorem 2.3). As stated, it shows for any $\Delta_1 \mid Q_1 \mid \Gamma \Vdash e : \sigma_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \Vdash e : \sigma_2$ we have $\sigma_1 = \sigma_2$. However, we would like to also show $\Delta_1 = \Delta_2$ and $Q_1 = Q_2$. However, due to such trivial generalizations, we can still have $\Delta_1 \neq \Delta_2$

or $Q_1 \neq Q_2$ depending on how many trivial generalizations were applied.

Without loss of generality, we can resolve this by using a stronger premise on $[\text{MGEN}]$ to always force all such trivial generalizations. In particular, we can strengthen $[\text{MGEN}]$ to:

$$\frac{\Delta \mid Q \mid \Gamma \vdash e : \sigma \quad \text{dom}(Q) \cup \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{\Delta \mid Q \mid \Gamma \Vdash_x e : \sigma} \text{MGENX}$$

which ensures all trivial substitutions have been applied while keeping the derived types equal. When using $[\text{MGENX}]$, we can prove a stronger principality lemma:

Lemma D.36. (*Type derivations with $[\text{MGENX}]$ are principal*)

For any $\Delta_1 \mid Q_1 \mid \Gamma \Vdash_x e : \sigma_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \Vdash_x e : \sigma_2$, we have $\Delta_1 = \Delta_2$, $Q_1 = Q_2$, and $\sigma_1 = \sigma_2$.

Moreover, we can show we can always rewrite any derivation using the plain $[\text{MGEN}]$ to a derivation using $[\text{MGENX}]$ without changing the derived types.

Lemma D.37. (*Any derivation can be rewriting using $[\text{MGENX}]$*)

For any $\Delta \mid Q \mid \Gamma \vdash e : \sigma$ using $[\text{MGEN}]$, we can rewrite it to a derivation $\Delta' \mid Q' \mid \Gamma \vdash e : \sigma$ using just $[\text{MGENX}]$ where $\Delta' \subseteq \Delta$ and $Q' \subseteq Q$.

Proof. (of Lemma D.37) Consider any application of $[\text{MGEN}]$:

$$\frac{\Delta \mid Q \mid \Gamma \vdash e : \sigma \quad \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{\Delta \mid Q \mid \Gamma \Vdash e : \sigma} \text{MGEN}$$

If $\text{dom}(Q) \subseteq \text{ftv}(\Gamma)$ we can immediately rewrite to $[\text{MGENX}]$. Otherwise, we must have $Q = Q' \cdot \alpha=\tau$ (by Lemma D.23) with $\alpha \notin \text{ftv}(\Gamma)$, and by $\text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$, also $\alpha \notin \text{ftv}(\sigma)$. Since $\text{ftv}(Q, \sigma) \subseteq \text{ftv}(\Delta, \Gamma)$, we also have $\Delta = \Delta', \alpha$. We can therefore use $[\text{GENSUB}]$ to derive:

$$\frac{\frac{\Delta', \alpha \mid Q' \cdot \alpha=\tau \mid \Gamma \vdash e : \sigma}{\Delta' \mid Q' \mid \Gamma \vdash e : \sigma} \text{GENSUB} \quad \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)}{\Delta' \mid Q' \mid \Gamma \Vdash e : \sigma} \text{MGEN}$$

The original derivation originating at $\Delta \mid Q \mid \Gamma \Vdash e : \sigma$ (which equals $\Delta', \alpha \mid Q' \cdot \alpha=\tau \mid \Gamma \Vdash e : \sigma$), still holds for $\Delta' \mid Q' \mid \Gamma \Vdash e : \sigma$. In particular, since $\alpha \in \Delta$ and $\alpha \notin \text{ftv}(\Gamma, \sigma)$, it is always independent of any combined constraints (in the $[\text{APP}]$ or $[\text{LET}]$ rules) and cannot change the derived types. We only need to consider a possible $[\text{GENSUB}]$ application that substitutes α in the original derivation:

$$\frac{\Delta, \alpha \mid Q \cdot \alpha=\tau \mid \Gamma \vdash e : \sigma}{\Delta \mid Q \mid \Gamma \vdash e : [\alpha:=\tau]\sigma} \text{GENSUB}$$

where we now no longer have Δ, α or $Q \cdot \alpha=\tau$ (but just Δ and Q). In such case though, this must be a trivial generalization since $\alpha \notin \text{ftv}(\Gamma, \sigma)$, and we can thus remove it, deriving $\Delta \mid Q \mid \Gamma \vdash e : \sigma$ directly. \square

Theorem 2.3 now follows directly from the previous lemmas.

Proof. (of Theorem 2.3) For any $\Delta_1 \mid Q_1 \mid \Gamma \Vdash e : \sigma_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \Vdash e : \sigma_2$, we can rewrite each into derivations $\Delta'_1 \mid Q'_1 \mid \Gamma \Vdash_x e : \sigma_1$ and $\Delta'_2 \mid Q'_2 \mid \Gamma \Vdash_x e : \sigma_2$ replacing all $[\text{MGEN}]$ by $[\text{MGENX}]$ (by Lemma D.37). From Lemma D.36 it now follows that $\sigma_1 = \sigma_2$ (and $\Delta'_1 = \Delta'_2$ with $Q'_1 = Q'_2$). \square

D.9.5 Derivations are Principal. At this point we have shown that any derivation can be rewritten such that (1) all leaves are $[\text{VAR}_S]$ rules (with no $[\text{INST}]$ rules left), (2) they end in a sequence of

[GENSUB-MONO] rules followed by [GEN] rules, and (3) we only use [MGEX] (instead of [MGEN]).

$$\begin{array}{c}
\frac{x_1 : \forall \bar{\alpha}. \tau_1 \in \Gamma_1}{\bar{\alpha} \mid \emptyset \mid \Gamma_1 \vdash x_1 : \tau_1} \text{VAR}_S \quad \dots \quad \frac{x_n : \forall \bar{\beta}. \tau_n \in \Gamma_n}{\bar{\beta} \mid \emptyset \mid \Gamma_n \vdash x_n : \tau_n} \text{VAR}_S \\
\vdots \qquad \qquad \qquad \vdots \\
\hline
\Delta' \mid Q' \mid \Gamma \vdash_0 e : \tau \\
\vdots \text{ [GENSUB-MONO]} \\
\vdots \text{ [GEN]} \\
\hline
\Delta \mid Q \mid \Gamma \vdash_n e : \sigma
\end{array}$$

We call these *canonical* derivations. We write (\rightsquigarrow) for such rewrite of a derivation tree, and for any derivation $\Delta \mid Q \mid \Gamma \vdash e : \sigma \rightsquigarrow \Delta' \mid Q' \mid \Gamma \vdash e : \sigma'$ we have $\Delta = \Delta'$, $Q = Q'$, and $\sigma = \sigma'$.

We can now prove Lemma D.36 and 2.4 that derivations are principal. We prove these recursively below but they are inductive on the number of (\Vdash_x) derivations (i.e. the number of let-bindings) and thus the proofs are still well-founded.

We write \vdash_n for a canonical derivation that ends in n applications of [GENSUB-MONO] and [GEN], originating in \vdash_0 on a syntax-directed rule ([FUN], [VAR_S], [APP], or [LET]).

Lemma D.38. (*(\vdash_0) derivations are principal*)

If $\Delta_1 \mid Q_1 \mid \Gamma \vdash_0 e : \tau_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \vdash_0 e : \tau_2$, then we have $\Delta_1 = \Delta_2$, $Q_1 = Q_2$, and $\tau_1 = \tau_2$.

Proof. Given that (\vdash_0) derivations are fully syntax directed, with no derivation choices, and since (\approx) is principal (Theorem D.28), and (\Vdash_x) is principal (inductively on Lemma D.36), the derivation is deterministic and it must be that $\Delta_1 = \Delta_2$, $Q_1 = Q_2$, and $\tau_1 = \tau_2$.

Definition 1. (*Canonical mono-type derivations*)

Consider a canonical $\Delta' \mid Q' \mid \Gamma \vdash_n e : \tau'$ derivation that ends in a monotype τ , and originated at a syntax directed rule $\Delta \mid Q \mid \Gamma \vdash_0 e : \tau$. Since we end in a monotype, it must end in a sequence of n [GENSUB-MONO] applications, such that $Q = Q_1 \cdot Q_2 \cdot Q_3$ for some Q_1, Q_2, Q_3 , with $\text{dom}(Q_1) \subseteq \text{ftv}(\Gamma)$ and $\text{dom}(Q_2, Q_3) \not\cap \text{ftv}(\Gamma)$, where $Q' = Q_1 \cdot Q_2$, $\tau' = Q_3[\tau]$ (where [GENSUB-MONO] has been applied over Q_3), and $\Delta' = \Delta$, $\text{dom}(Q_3)$.

As such, we write $\Delta \mid Q_1 \cdot Q_2 \mid \Gamma \vdash_n e : Q_3 \mid \tau$ for such canonical monotype derivations. We also often write just $\Delta \mid Q \mid \Gamma \vdash_n e : Q' \mid \tau$ if the extra precision is not required.

With this definition, we can give a more precise version of Theorem 2.4:

Lemma D.39. (*Mono-type derivations are principal – extended*)

If $\Delta_1 \mid Q_1 \cdot Q'_1 \mid \Gamma \vdash_n e : Q'_1 \mid \tau_1$ and $\Delta_1 \mid Q_2 \cdot Q'_2 \mid \Gamma \vdash_n e : Q'_2 \mid \tau_2$, then we also have $Q_1 = Q_2$, with $Q'_1 \cdot Q'_1 = Q'_2 \cdot Q'_2$, and $\tau_1 = \tau_2$ (and $\Delta_1, \text{dom}(Q'_1) = \Delta_2, \text{dom}(Q'_2)$).

Proof. Each derivation must originate in $\Delta_3 \mid Q_3 \mid \Gamma \vdash_0 e : \tau_3$ and $\Delta_4 \mid Q_4 \mid \Gamma \vdash_0 e : \tau_4$ and by Lemma D.38, we have $\Delta_3 = \Delta_4$, $Q_3 = Q_4$, and $\tau_3 = \tau_4$ – we use Δ, Q, τ for those respectively. Since each derivation ends in a sequence of [GENSUB-MONO] rules, we also have $Q = Q_1 \cdot Q'_1 \cdot Q'_1$ and $Q = Q_2 \cdot Q'_2 \cdot Q'_2$ with $\text{dom}(Q'_1, Q'_1, Q'_2, Q'_2) \not\cap \text{ftv}(\Gamma)$, and $\text{dom}(Q_1, Q_2) \subseteq \text{ftv}(\Gamma)$, and therefore $Q_1 = Q_2$, and $Q'_1 \cdot Q'_1 = Q'_2 \cdot Q'_2$. Moreover, we have that $\tau = \tau_1 = \tau_2$, and also $\Delta_1, \text{dom}(Q'_1) = \Delta = \Delta_2, \text{dom}(Q'_2)$. \square

As a corollary, we also have the simplified lemma:

Lemma D.40. (*Mono-type derivations are principal – simplified*)

If $\Delta_1 \mid Q_1 \mid \Gamma \vdash_n e : Q'_1 \mid \tau_1$ and $\Delta_1 \mid Q_2 \mid \Gamma \vdash_n e : Q'_2 \mid \tau_2$, then $Q_1 \cdot Q'_1 = Q_2 \cdot Q'_2$, and $\tau_1 = \tau_2$ (and $\Delta_1, \text{dom}(Q'_1) = \Delta_2, \text{dom}(Q'_2)$).

This follows directly from Lemma D.39. From this, we can prove our main Theorem 2.4:

Proof. (of Theorem 2.4, *principal mono-type derivations*) From the assumption, we have the derivations $\Delta_1 \mid Q_1 \mid \Gamma \vdash e : \tau_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \vdash e : \tau_2$ and thus we have the equivalent canonical derivations $\Delta_1 \mid Q_1 \mid \Gamma \vdash_n e : Q'_1 \mid \tau'_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \vdash_n e : Q'_2 \mid \tau'_2$ with $\tau_1 = Q'_1[\tau'_1]$ and $\tau_2 = Q'_2[\tau'_2]$ (1). From Lemma D.40, we have $Q_1 \cdot Q'_1 = Q_2 \cdot Q'_2$ (2), and $\tau'_1 = \tau'_2$ (3), and thus:

$$\begin{aligned}
& Q_1[\tau_1] \\
= & Q_1[Q'_1[\tau'_1]] \quad \{ (1) \} \\
= & (Q_1 \cdot Q'_1)[\tau'_1] \\
= & (Q_2 \cdot Q'_2)[\tau'_1] \quad \{ (2) \} \\
= & (Q_2 \cdot Q'_2)[\tau'_2] \quad \{ (3) \} \\
= & Q_2[Q'_2[\tau'_2]] \quad \{ (2) \} \\
= & Q_2[\tau_2] \quad \{ (1) \}
\end{aligned}$$

□

We can finally prove our main Lemma of principality.

Proof. (of Lemma D.36, *principal type derivations*) From the assumption, we have the canonical form of the premises $\Delta_1 \mid Q_1 \mid \Gamma \vdash_n e : \sigma_1$ and $\Delta_2 \mid Q_2 \mid \Gamma \vdash_m e : \sigma_2$ (1) with $\text{dom}(Q_1) \cup \text{ftv}(\sigma_1) \subseteq \text{ftv}(\Gamma)$ and $\text{dom}(Q_2) \cup \text{ftv}(\sigma_2) \subseteq \text{ftv}(\Gamma)$ (2). Moreover, each originates at a syntax rule as $\Delta'_1 \mid Q'_1 \mid \Gamma \vdash_0 e : \tau'_1$ and $\Delta'_2 \mid Q'_2 \mid \Gamma \vdash_0 e : \tau'_2$. By (induction on) Lemma D.38, we have $\Delta'_1 = \Delta'_2$, $Q'_1 = Q'_2$, and $\tau'_1 = \tau'_2$ – we will write Δ' , Q' , τ' for those respectively.

Each derivation first applies [GENSUB-MONO] rules, ending in a monotype $\Delta'_1 \mid Q_1 \mid \Gamma \vdash_i e : \tau_1$ and $\Delta'_2 \mid Q_2 \mid \Gamma \vdash_j e : \tau_2$ (with Q_1 and Q_2 as the subsequent [GEN] rules leave the prefix unchanged). Since [GENSUB-MONO] only substitutes, we must have $Q' = Q_1 \cdot Q'_1$ and $Q' = Q_2 \cdot Q'_2$ (3a) with $\tau_1 = Q'_1[\tau']$ and $\tau_2 = Q'_2[\tau']$ (3b). By (2) $\text{dom}(Q_1) \subseteq \text{ftv}(\Gamma)$. Consider any $\alpha \in \text{dom}(Q'_1)$. It cannot be that $\alpha \in \text{ftv}(\Gamma)$ or otherwise [GENSUB] would not apply, and similarly for Q'_2 . Therefore, by (3a), we must have $Q_1 = Q_2$, i.e. we have maximally substituted. It follows we also have $Q'_1 = Q'_2$, and thus $\tau_1 = \tau_2$ and $\Delta'_1 = \Delta'_2$ – we write Q, τ, Δ , for those respectively.

Finally, we have a sequence of [GEN] rules from $\Delta \mid Q \mid \Gamma \vdash_i e : \tau$ and $\Delta \mid Q \mid \Gamma \vdash_j e : \tau$. For every $\alpha \in \text{ftv}(\tau)$ with $\alpha \notin \text{ftv}(\Gamma)$ (we have $\alpha \notin \text{ftv}(Q)$ by (4)), and there must be a corresponding [GEN] application or otherwise the side condition (2) cannot hold. (and any other [GEN] applications with $\alpha \notin \text{ftv}(\tau, \Gamma)$ are trivial). With $\bar{\alpha} = \text{ftv}(\tau)$, we therefore have $\Delta = \Delta_1, \bar{\alpha} = \Delta_2, \bar{\alpha}$, and thus $\Delta_1 = \Delta_2$, and $\sigma_1 = \forall \bar{\alpha}. \tau = \sigma_2$. □

D.9.6 Syntax Directed Type Rules. As shown, in a canonical derivation, all leaves are [VAR_s] rules (with no [INST] rules left), and they end in a sequence of [GENSUB-MONO] rules followed by [GEN] rules. This suggests a deterministic procedure for (\Vdash_x) where we first apply all [GENSUB-MONO] rules followed by the [GEN] rules:

$$\begin{aligned}
& \text{gen} : (\Delta, Q, \Gamma, \sigma) \rightarrow (\Delta, Q, \sigma), \text{ with } \Delta \not\cap \text{ftv}(\Gamma) \\
& \text{gen}((\Delta, \alpha), Q \cdot \alpha = \tau', \Gamma, \tau) = \text{gen}(\Delta, Q, \Gamma, [\alpha := \tau']\tau) \\
& \text{gen}((\Delta, \alpha), Q, \Gamma, \sigma) = \text{gen}(\Delta, Q, \Gamma, \forall \alpha. \sigma) \quad \text{if } \alpha \notin \text{ftv}(Q) \\
& \text{gen}(\Delta, Q, \Gamma, \sigma) = (\Delta, Q, \sigma) \quad \text{if } \text{dom}(Q) \cup \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)
\end{aligned}$$

We can then make the [MGENX] rule syntax directed by applying all [GENSUB-MONO] and [GEN] rules at once using gen:

$$\frac{\Delta_0 \mid Q_0 \mid \Gamma \vdash_s e : \tau \quad (\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \tau)}{\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma} \text{MGEN}_s$$

Together with $[\text{VAR}_s]$, we can now make all type rules syntax-directed as shown in Figure 11 in Appendix B. We can show the syntax directed rules are sound and complete to the HMQ rules:

Theorem D.41. (*The syntax-directed type rules are sound*)

If $\Delta \mid Q \mid \Gamma \Vdash_x e : \sigma$, then also $\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma$.

Note that by Lemma D.37, it follows that if $\Delta \mid Q \mid \Gamma \Vdash e : \sigma$, then also $\Delta' \mid Q' \mid \Gamma \Vdash_s e : \sigma$ with $\Delta' \subseteq \Delta$ and $Q' \subseteq Q$, i.e. since the syntax directed system always generalizes all trivial substitutions eagerly it will give the shortest Δ', Q' .

Theorem D.42. (*The syntax-directed type rules are sound for mono-types*)

If $\Delta \mid Q \mid \Gamma \vdash_0 e : \tau$, then also $\Delta \mid Q \mid \Gamma \vdash_s e : \tau$.

Theorem D.43. (*The syntax-directed type rules are complete*)

If $\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma$, then also $\Delta \mid Q \mid \Gamma \vdash e : \sigma$.

Theorem D.44. (*The syntax-directed type rules are complete for mono-types*)

If $\Delta \mid Q \mid \Gamma \vdash_s e : \tau$, then also $\Delta \mid Q \mid \Gamma \vdash_0 e : \tau$.

The proofs of soundness and completeness are recursive, but we again use induction on the number of let bindings to ensure these are well-founded.

Proof. (of Theorem D.44) We do case analysis on the derivation $\Delta \mid Q \mid \Gamma \vdash_s e : \tau$. The $[\text{VAR}_s]$ rules correspond directly, as do the $[\text{APP}_s]/[\text{APP}]$, and $[\text{FUN}_s]/[\text{FUN}]$ rules. For the $[\text{LET}_s]$ rule, we have the premise $\Delta_1 \mid Q_1 \mid \Gamma \Vdash_s e : \sigma$ and by (induction on) Theorem D.43, we also have $\Delta_1 \mid Q_1 \mid \Gamma \vdash e : \sigma$, and $[\text{LET}]$ holds as well. \square

Lemma D.45. (*gen corresponds to a sequence of $[\text{GENSUB-MONO}]$ and $[\text{GEN}]$ applications*)

If $\Delta_0 \mid Q_0 \mid \Gamma \vdash e : \tau$ with $(\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \tau)$, then also $\Delta \mid Q \mid \Gamma \vdash_n e : \sigma$ (using only the $[\text{GENSUB-MONO}]$ and $[\text{GEN}]$ rules) with $\text{dom}(Q) \cup \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$.

Proof. (of Lemma D.45) We have $(\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \sigma_0)$. By inspecting the rules of gen, we must have $\Delta_0 = \Delta, \bar{\alpha}, \bar{\beta}$, $Q_0 = Q \cdot Q'$, $\bar{\beta} = \text{dom}(Q')$, and $\sigma = \forall \bar{\alpha}. Q'[\tau]$ where (by rule gen.1) $\text{ftv}(Q') \not\cap \text{ftv}(Q, \Gamma)$ (1a), and (by rule gen.2) as $\not\cap \text{ftv}(Q, \Gamma)$ (1b), finally (by rule gen.3) $\text{dom}(Q) \cup \text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$. By (1a), we can first apply a sequence of $[\text{GENSUB-MONO}]$ to derive

$$\frac{\Delta, \bar{\alpha}, \bar{\beta} \mid Q \cdot Q' \mid \Gamma \vdash e : \tau \quad \text{ftv}(Q') \not\cap \text{ftv}(Q, \Gamma)}{\vdots} \text{GENSUB-MONO}$$

$$\frac{\vdots}{\Delta, \bar{\alpha} \mid Q \mid \Gamma \vdash e : Q'[\tau]} \text{GENSUB-MONO}$$

and by (1b), we can then apply a sequence of $[\text{GEN}]$ rules:

$$\frac{\Delta, \bar{\alpha} \mid Q \mid \Gamma \vdash e : Q'[\tau] \quad \bar{\alpha} \not\cap \text{ftv}(Q, \Gamma)}{\vdots} \text{GEN}$$

$$\frac{\vdots}{\Delta \mid Q \mid \Gamma \vdash e : \forall \bar{\alpha}. Q'[\tau]} \text{GEN}$$

\square

Proof. (of Theorem D.43) We have a derivation $\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma$, and the premises $\Delta_0 \mid Q_0 \mid \Gamma \vdash_s e : \tau$ (1a) with $(\Delta, Q, \sigma) = \text{gen}(\Delta_0, Q_0, \Gamma, \tau)$ (1b). From (1b) and the final clause of gen , we also have $\text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$ (1c). By (induction on) Theorem D.44 and (1a), we have $\Delta_0 \mid Q_0 \mid \Gamma \vdash_0 e : \tau$ (2). By (1b,2) and Lemma D.45, we have $\Delta \mid Q \mid \Gamma \vdash_n e : \sigma$ (3). Together with (1c), we can now apply $[\text{MGEN}]$ to derive $\Delta \mid Q \mid \Gamma \Vdash e : \sigma$. \square

The proofs of soundness are dual to the completeness proofs:

Proof. (of Theorem D.42) We do case analysis on the derivation $\Delta \mid Q \mid \Gamma \vdash_0 e : \tau$. The $[\text{VAR}_S]$ rules correspond directly, as do the $[\text{APP}]/[\text{APP}_S]$, and $[\text{FUN}]/[\text{FUN}_S]$ rules. For the $[\text{LET}]$ rule, we have the premise $\Delta_1 \mid Q_1 \mid \Gamma \Vdash e : \sigma$ and (by induction on Theorem D.41), we also have $\Delta_1 \mid Q_1 \mid \Gamma \vdash_s e : \sigma$, and $[\text{LET}_S]$ holds as well. \square

Proof. (of Theorem D.41) We have a derivation $\Delta \mid Q \mid \Gamma \Vdash e : \sigma$, and the premises $\Delta \mid Q \mid \Gamma \vdash_n e : \sigma$ (1a) with $\text{ftv}(\sigma) \subseteq \text{ftv}(\Gamma)$ (1b). The derivation (1a) must originate in a $\Delta_0 \mid Q_0 \mid \Gamma \vdash_0 e : \tau$ (2a). By (induction on) Theorem D.42 and (2a), we have $\Delta_0 \mid Q_0 \mid \Gamma \vdash_s e : \tau$ (3a). Since (1a) is followed by a sequence of n applications of $[\text{GENSUB}]$ and $[\text{GEN}]$, we also have $\Delta' = \Delta, \bar{\alpha}, Q_0 = Q \cdot Q'$, and $\sigma = \forall \bar{\alpha}. Q'[\tau]$ with $\text{dom}(Q') \not\cap \text{ftv}(Q, \Gamma)$ and $\bar{\alpha} \not\cap \text{ftv}(Q, \Gamma)$ (by 1b). It now follows that $\text{gen}(\Delta_0, Q_0, \Gamma, \tau) = \text{gen}((\Delta, \bar{\alpha}), Q \cdot Q', \Gamma, \tau) = (\Delta, Q, \Gamma, \forall \bar{\alpha}. Q'[\tau]) = (\Delta, Q, \Gamma, \sigma)$ (3b). With (3a,3b), we can apply $[\text{MGEN}_S]$ to derive $\Delta \mid Q \mid \Gamma \Vdash_s e : \sigma$. \square