Chapter 11

Vindication of FOL Inference

We are now in a position to begin the main task of this monograph, by justifying and explaining mathematicians' ordinary use of first order logic in set theory, from a potentialist point of view. Note that my potentialist translations of a set theoretic sentence have a very different logical form from the original. Thus, it is not immediately obvious that whenever a set theoretic sentence β is a first order logical consequence of α, then $t(\beta)$ is a genuine logical consequence of $t(\alpha)$. And it is not obvious that set theorists' willingness to move from a sentence that looks like it asserts α to one that looks like it asserts β is justified.

In this section I will show that every first-order logical argument in the language of set theory can be transformed into an argument in the deduction system described above, which takes us from the translation of the premises for this argument to the translation of its conclusion. In particular, I will establish the following claim.

Proposition 11.0.1. *If $\gamma_1, \ldots, \gamma_m, \theta$ are sentences in the language of set theory and $\gamma_1 \ldots \gamma_m \vdash_{FOL} \theta$ then $t(\gamma_1) \ldots t(\gamma_m) \vdash (\theta)$ (where $\vdash_{FOL}$ represents provability in first order logic and $\vdash$ represents provability in the formal system just introduced).*

11.1 Translation Lemma

A key tool in making this argument will be the following Translation Lemma.

Translation Lemma: If $v_1 \ldots v_k$ are the only variables free in a set theoretic formula ϕ, then $\vdash V_n \geq V_m \wedge f_n = f_m(v_1) \wedge ... f_n(v_k) = f_m(v_k) \rightarrow (t_n(\phi) \leftrightarrow t_m(\phi))$.

Remember that $V_n \geq V_m$ means that V_n extends V_m when considered merely as a hierarchy of sets (unlike $\vdash \vec{V_n} \geq \vec{V_m}$ it does not require any agreement between the assignment functions f_n and f_m).

Intuitively, the Translation Lemma says that the way V_n, f_n assigns the free variables in a set theoretic formula ϕ completely determines whether $t_n(\phi)$ is true. Specifically, the truth-value of $t(\phi)$ relative to this V_n, f_n cannot be changed by considering it relative to any extending V_m, f_n which preserves these assignments for free variables.

Note: Hellman proves something analogous to this lemma in [3], assuming the axiom of inaccessibles. [add quote]

Proof. I will prove this claim by induction on complexity of formulas.

This principle is fairly obviously true for atomic sentences, which all take the form $x = y$ or $x \in y$. For, if $f_n(x) = f_m(x) \wedge f_n(y) = f_m(y) \rightarrow (f_n(x) = f_n(y) \leftrightarrow f_m(x) = f_m(y))$. Similarly, we get $f_n(x) = f_m(x) \wedge f_n(y) = f_m(y) \rightarrow$

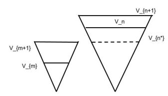

Figure 11.1: One possible relation relationship between V_m, V_{m+1}, V_n and the V_*, V_{n+1} we will construct.

$(f_n(x) \in_n f_n(y) \leftrightarrow f_m(x) \in_m f_m(y))$ by $V_n \geq_{set} V_m$. It is also obvious that applying the truth-functional connectives $\wedge$, $\vee$ and $\neg$ to formulae which have this property always yields new formulas which has this property.

Assume $f_n(v_1) = f_m(v_1) \wedge ...$ for all variables $v_1 \dots v_k$ free in $(\forall x)\phi(x)$. Assuming that the relevant theorem holds for ϕ, we will show that it holds for $(\forall x)\phi(x)$ as well.

Suppose, for contradiction, that we had $t_n((\forall x)\phi(x))$, but not $t_m((\forall x)\phi(x))$. Then we have $\Box_{Vn}[V_{n+1} \geq_{\mathbf{x}} V_n \rightarrow t_{n+1}(\phi)]$. But we also have $\neg t_m((\forall x)\phi(x)) = \Diamond_{Vm}[V_{m+1} \geq_{\mathbf{x}} V_m \wedge \neg t_{m+1}(\phi)]$.

Our strategy will be to build (holding fixed V_n, V_m) a $V_{n+1} \geq_{\mathbf{x}} V_n$ which mimics the assignment of $f_{m+1}(v)$ which makes $t_{m+1}(\phi)$ come out false, and then show that this scenario must be one in which $\neg t_{n+1}(\phi)$ (contrary to the $\Box_{Vn}$ claim above, which can be imported into this $\Diamond_{V_n,V_m}$ context).

By the Hierarchy Extending Lemma (proved in appendix C.5) we can have V_{n+1} extending V_n, containing an isomorphic copy, V_{n*}, of V_{m+1}.

Lemma 11.1.1 (Hierarchy Extending Lemma). *If $\mathscr{V}(V_a) \wedge \mathscr{V}(V_b)$ then it's possible (holding fixed V_a, V_b) that there is a V_{a+} extending V_a and Z' functioning like an isomorphism from V_b to an intial segment of V_{a+}*

This isomorphism between sets (and the unique isomorphism between copies of the numbers) naturally induces an isomorphism for f, so now we have V_n*, f_n* isomorphic to V_{m+1}, f_{m+1}. Then by the Isomorphism Lemma (proved in appendix B.2), we can infer from the falsehood of $t_{m+1}(\phi)$ to the falsehood of $t_{n*}(\phi)$ (where $t_{n**}(\phi)$ is the version of $t_n(\phi)$ which talks about V_n*, f_n* rather than V_n, f_n[1]).

Ismorphism Lemma *If ϕ is a formula employing only relation symbols $R_1 \ldots R_n$ and quantifiers restricted to objects related by $R_1 \ldots R_n$* [2] *outside of all $\Box$s and $\Diamond$s, ϕ' is the result of replacing each R_i in ϕ with a corresponding R'_i (which does not occur anywhere in ϕ), and $x_1, ...x_n$ are all the free variables in ϕ then:*

$\vdash \langle R_1, \ldots R_m \rangle \cong_Z \langle R'_1, \ldots R'_m \rangle \to (\forall x_1)...(\forall x_n)...(\forall x'_1)...(\forall x'_n)[Z(x_1, x'_1) \wedge ...Z(x_n, x'_n) \to (\phi(x_1...x_n) \leftrightarrow \phi'(x'_1...x'_n)]$

By Simple Comprehension (and the Multiple Definitions Lemma proved on 8.5), we can have f_{n+1} such that $\vec{V_{n+1}} \geq_x \vec{V_n}$ and $f_{n+1}(x) = f_{n*}(x)$. Now we want to show that f_{n+1} and f_{n*} will agree on all variables free in ϕ, and then deploy (a version of) the inductive hypothesis to infer from $\neg t_*(\phi)$ to $\neg t_{n+1}(\phi)$.

[Directly, the inductive hypothesis only gives us $\vdash V_{n+1} \geq V_n \wedge f_{n+1} = f_n(v_1) \wedge ...f_{n+1}(v_k) = f_n(v_k) \to (t_{n+1}(\psi) \leftrightarrow t_n(\psi))$. But applying $\Box$I, lets us deduce that this claim is logically necessary. Then applying $\Box$ Relabling (proved on **??**) lets us derive that the corresponding claim with all instances

[1] That is, t_{n*} is just like t_n, but replaces all appeals to set_{n+1} with $set_{n+1}*$ and so forth.

[2] i.e., only quantifiers in clauses of the form $(\forall x)(x \in Ext(\hat{R}) \to ...)$ and $(\exists x)(x \in Ext(\hat{R}) \wedge ...)$, where $\hat{R}$ is a collection of one or more relations from within $R_1 \ldots R_n$, an Ext(R) is as defined as per 5.2.1.

of f_n replaced by f_{n*} (note that t_{n+1} makes no mention of V_n or f_n) is logically necessary, hence actually true. Thus we have $\vdash V_{n+1} \geq V_{n*} \wedge f_{n+1} = f_{n*}(v_1) \wedge ... f_{n+1}(v_k) = f_{n*}(v_k) \rightarrow (t_{n+1}(\psi) \leftrightarrow t_{n*}(\psi))$, as desired.]

Now it remains to prove the antecedent of the claim above. Clearly f_{n+1} and f_{n*} agree on x. But we can also show that they agree on all the other variables free in ϕ, since for all variables v free in $(\forall x)\phi(x)$, $f_{n+1}(v) = f_n(v) = f_m(v) = f_{m+1}(v) = f_{n*}$, with the last equivalence holding because V_{n*}, f_{n*} was constructed to be isomorphic to f_{m+1}.[3] We also know that $V_{n+1} \geq V_{n*}$, by our construction of V_{n+1}.

Thus, by (our suitably massaged version of) the inductive hypothesis we know that ϕ is true on V_{n*}, f_{n*} iff it is true V_{n+1}, f_{n+1}. Since ϕ is false on V_{n*}, it must be false on V_{n+1} as well. Thus, contrary to the $\Box_{V_n}$ claim above, we have a logically possible scenario (holding fixed V_n) in which $V_{n+1} \geq_{\mathbf{x}} V_n \wedge \neg t_{n+1}(\phi)$. So $\perp$.

The same argument with m and n swapped shows that we can't have $t_m((\forall x)\phi(x))$ be true while $t_n((\forall x)\phi(x))$ is false. Thus we have $t_n((\forall x)\phi(x)) \leftrightarrow t_m((\forall x)\phi(x))$ as desired. $\Box$

[3] To spell this argument out rigorously, consider some possible isomorphisms relating $V_n, V_{n+1}, V_m, V_{m+1}$ and V_{n*}. By the fact that $V_n \geq_{set} V_m \vee V_m \geq_{set} V_n$, we know that the identity map $i_{n,m}$ behaves like [alt: it would be possible to have an I which behaves like an identity map, and this would constitute an...] an isomorphism from V_n (or an initial segment of it) to V_m. Consider what happens when we compose this map with the isomorphism z from V_{m+1}, f_{m+1} to V_{n*}, f_{n*}. What we get is an isomorphism from initial segment of V_{n+1} to an initial segment of V_{n+1}. By the Isomorphism Agreement Lemma, we know that this map cannot disagree with the identity isomorphism from V_n to V_n.

Now consider an aribitrary variable v which is not x and occurs in ϕ. It suffices to show that the composite function $z(i_{n,m}(f_n(v))) = f_{n*}(v)$, because then the fact that this composite function behaves like the identity gives us $f_{n*}(v) = f_n(v) = f_{n+1}(v)$.

By hypothesis we have $f_n(v) = f_m(v)$. By the fact that $V_{m+1} \geq_x V_m$, we have $f_m(v) = f_{m+1}(v)$, and hence $i_{n,m}(f_n(v)) = f_m(v) = f_{m+1}(v)$. By the fact that z as an isomorphism between V_{m+1}, f_{m+1} and V_{n*}, f_{n*}, we have $z(f_{m+1}(v)) = f_{n*}(v)$. Putting these two facts together gives $z(i_{n,m}(f_n(v))) = f_{n*}(v)$ as desired.

11.2 Vindictation of FOL

Now let us return to our ultimate task: showing that every first-order logical argument in the language of set theory can be transformed into an argument in the deduction system described above, which takes us from the translation of the premises for this argument to the translation of its conclusion.

Proposition 11.2.1. *If $\gamma_1, \ldots, \gamma_m, \theta$ are sentences in the language of set theory and $\gamma_1 \ldots \gamma_m \vdash_{FOL} \theta$ then $t(\gamma_1) \ldots t(\gamma_m) \vdash (\theta)$ (where $\vdash_{FOL}$ represents provability in first order logic and $\vdash$ represents provability in the formal system just introduced).*

Proof. First note that it suffices to show that whenever $\gamma_1, \ldots, \gamma_m, \theta$ are as in the proposition above we also have $\mathscr{V}(V_0), t_0(\gamma_1) \ldots t_0(\gamma_m) \vdash t_0(\theta)$. For, by the lemma below, any such proof can be transformed into a proof one witnessing $t(\gamma_1) \ldots t(\gamma_m) \vdash t(\theta)$.

Lemma 11.2.2. *It $\mathscr{V}(V_0), t_0(\gamma_1) \ldots t_0(\gamma_m) \vdash t_0(\theta)$ then $t(\gamma_1) \ldots t(\gamma_m) \vdash t(\theta)$.*

Proof. Recall that when ψ is a sentence of set theory, $t(\psi) = \Box(\mathscr{V}(V_0) \rightarrow t_0(\psi))$. If we have a proof witnessing $\mathscr{V}(V_0), t_0(\gamma_1) \ldots t_0(\gamma_m) \vdash t_0(\theta)$, then we also have one witnessing $\vdash \mathscr{V}(V_0) \wedge t_0(\gamma_1), \ldots, t_0(\gamma_m) \rightarrow t_0(\theta)$. The latter proof can be transformed into one witnessing $t(\gamma_1) \ldots t(\gamma_m) \vdash t(\theta)$ as follows.

Start with the assumptions that $t(\gamma_1) \ldots t(\gamma_m)$, i.e., $\Box(\mathscr{V}(V_0) \rightarrow t_0(\gamma_i))$ for each i. To get the desired conclusion that $t(\theta) = \Box(\mathscr{V}(V_0) \rightarrow t_0(\theta))$, we will make a $\Box$I argument. Consider an arbitrary logically possible scenario (holding fixed nothing). Suppose, for conditional introduction, that $\mathscr{V}(V_0)$.

Our assumptions $t(\gamma_i)$ are content-restricted to the empty set (because their outer connective is a $\Box$ with no subscripts), so they can be assumed to remain true within the arbitrary logically possible scenario under consideration. By applying $\Box$E to each $t(\gamma_i)$, we get the corresponding $\mathscr{V}(V_0) \to t_0(\gamma_i)$ claim. Thus we can derive that $t_0(\gamma_i)$.

By the assumptions in the first paragraph we can also prove from empty premises that $\mathscr{V}(V_0) \wedge t_0(\gamma_1), \ldots, t_0(\gamma_m) \to t_0(\theta)$. Thus we can infer that $t_0(\theta)$. Discharging our assumption that $\mathscr{V}(V_0)$ gives us the conditional $\mathscr{V}(V_0) \to t_0(\theta)$ within the logically possible scenario being considered.

Since all of our assumptions are content-restricted to the empty set, the rules for $\Box$I allow us to conclude that $t(\theta) = \Box[\mathscr{V}(V_0) \to t_0(\theta)]$, as desired. $\square$

Thus it suffices to show that if $\mathscr{V}(V_0), t_0(\gamma_1) \ldots t_0(\gamma_m) \vdash t_0(\theta)$ then $\mathscr{V}(V_0), t_0(\gamma_1) \ldots t_0(\gamma_m) \vdash t_0(\theta)$.

We will prove this claim by induction on the length of FOL proofs. But first I must explain what I mean by FOL proofs.

Given a rule such as ($\wedge$I) If $\Gamma_1 \vdash_{FOL} \theta$ and $\Gamma_2 \vdash_{FOL} \psi$, then $\Gamma_1, \Gamma_2 \vdash_{FOL} (\theta \wedge \psi)$ we refer to $\Gamma_1 \vdash_{FOL} \theta$ and $\Gamma_2 \vdash_{FOL} \psi$ as the first and second inputs to the rule, $\Gamma_1, \Gamma_2 \vdash_{FOL} (\theta \wedge \psi)$ as the output of the rule.

Definition I will take FOL proofs to be finite sequences of numbered lines, where each line corresponds to an application of one of our FOL closure conditions. Specifically, a proof is a sequence of lines where the n-th line is given by $\langle \theta_n, R_n, \vec{l_n}, \Gamma_n \rangle$, such that θ_n is a formula of FOL, Γ_n is a set of first order formulas and $\vec{l}$ is a tuple of line numbers $c_i \leq n$ where i indexes the inputs required by rule R (where we regard the assumption rule as having itself as an input). So, for example, if R was $\wedge I$ the tuple $\vec{l}$ would be a pair of line numbers while if R was As it would be empty. Finally, we require that R yield $\Gamma_n \vdash \theta_n$ when acting on the inputs $\Gamma_{c_i} \vdash \theta_{c_i}$ for c_i in $\vec{l}$. Note that the validity of a proof is a simple syntactic matter to check. I will say that a proof **witnesses** the fact that $\Gamma \vdash_{FOL} \theta$ if its last line takes the form $\langle \theta, R(...), \vec{l_n}, \Gamma \rangle$.

Note that it follows from the definition above that truncating an n-line proof of θ at some line $m < n$ $\langle \theta_{n'}, R(c_1, c_2), \Gamma_{n'} \rangle$ yields a strictly shorter m-line proof that $\Gamma_{n'} \vdash_{FOL} \theta_{n'}$.

I will proceed by induction on the length of the proof to be translated.

We first establish the claim holds for proofs of height 1. Only two kinds of proofs of length 1 are possible.

First, there are proofs generated by a single application of the assumption introduction rule, i.e., proofs witnessing $\phi \vdash_{FOL} \phi$. Clearly in all such cases we have $\mathscr{V}(V_0), t_0(\phi) \vdash t_0(\phi)$.

Second, there single line proofs of $\vdash_{FOL} v = v$ produced by an application of (=I). In all such cases we have $\mathscr{V}(V_0), \vdash t_0(v = v)$, since $t_0(v = v)$ is $f_0(v) = f_0(v)$.

Now assume, by way of induction, that whenever θ is any *formula* in the language of set theory, Γ is a set of such formulas and there is a proof of θ from Γ with height at most n, then $\mathscr{V}(V_0), \Gamma \vdash t_0(\theta)$ and establish the claim for proofs of height $n + 1$

We consider each possibility for the final rule used in the proof of θ

We first note that since our definition of t_n above (specifically, given the fact that we have defined t_0 to commute with $\wedge, \vee, \neg, \rightarrow$), it is trivial to demonstrate the inductive step in every case where the proof ends with one of the rules below.

($\wedge$I) If $\Gamma_1 \vdash_{FOL} \theta$ and $\Gamma_2 \vdash_{FOL} \psi$, then $\Gamma_1, \Gamma_2 \vdash_{FOL} (\theta \wedge \psi)$.

($\wedge$E) If $\Gamma \vdash_{FOL} (\theta \wedge \psi)$ then $\Gamma \vdash_{FOL} \theta$; and if $\Gamma \vdash_{FOL} (\theta \wedge \psi)$ then $\Gamma \vdash_{FOL} \psi$.

($\vee$I) If $\Gamma \vdash_{FOL} \theta$ then $\Gamma_1 \vdash_{FOL} h\theta \vee \psi$; if $\Gamma \vdash_{FOL} \psi$ then $\Gamma \vdash_{FOL} \theta \vee \psi$.

($\vee$E) If $\Gamma_1 \vdash_{FOL} (\theta \vee \psi)$, $\Gamma_2, \theta \vdash_{FOL} \phi$ and $\Gamma_3, \psi \vdash_{FOL} \phi$, then $\Gamma_1, \Gamma_2, \Gamma_3 \vdash_{FOL} \phi$.

($\rightarrow$ I) If $\Gamma, \theta \vdash_{FOL} \psi$, then $\Gamma \vdash_{FOL} (\theta \rightarrow \psi)$.

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

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

(DNE) If $\Gamma \vdash_{FOL} \neg\neg\theta$ then $\Gamma \vdash_{FOL} \theta$.

($\perp$ I) If $\Gamma \vdash_{FOL} \psi \wedge \neg\psi$ then $\Gamma \vdash \perp$.

($\perp$ E) If $\Gamma, \theta_{FOL} \vdash \perp$ then $\Gamma \vdash_{FOL} \neg\theta$.

For example, to check the clause for $\rightarrow$I, suppose that we have a length $n+1$ proof of $\theta \rightarrow \psi$ whose final rule is $\rightarrow$I. This ensures that we have a height-n proof of ψ from the hypotheses Γ, θ. By inductive hypothesis, we have a proof of $t_0(\psi)$ from $\mathscr{V}(V_0), t_0[\Gamma], t_0(\theta) \vdash t_0(\psi)$. Adding an application of $\rightarrow$I to the end of this proof produces a proof of $\mathscr{V}(V_0), t_0[\Gamma] \vdash t_0(\theta) \rightarrow t_0(\psi)$.

Now, consider the possibility that our proof of length $n+1$ concludes with an application of $\forall$I.

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

Given an $n+1$-line proof of $\forall v\theta$, we can extract an n-line proof witnessing $\gamma_1 ... \gamma_n \vdash_{FOL} \theta$ where the variable v does not occur free in any member of Γ. So, by inductive hypothesis, we have $\mathscr{V}(V_0), t_0[\Gamma] \vdash t_0(\theta)$ and we must establish $\mathscr{V}(V_0), t_0[\Gamma] \vdash t_0(\forall v\theta) = \Box_{V_0}[V_1 \geq_{\mathbf{v}} V_0 \rightarrow t_1(\theta)]$.

So assume $\mathscr{V}(V_0), t_0[\Gamma]$. Now consider an arbitrary scenario in which $V_1 \geq_{\mathbf{v}} V_0$ (while the facts about V_0 are held fixed). If we could infer $t_1(\theta)$ the result would follow immediately.

But the inductive hypothesis (and the fact that $\mathscr{V}(V_0)$ is content restricted to V_0, and by observation 2 in chapter 10.3, so is $t_0[\Gamma]$) we can infer that $t_0(\theta)$. Since v is free in θ in different choices for $V_1 \geq_{\mathbf{v}} V_0$ could (in principle) change the truth value of $t_1(\theta)$.

However, one important further fact is available: our inductive hypothesis says that $t_0(\theta)$ is derivable from $\mathscr{V}(V_0), t_0[\Gamma]$ and these formula make no mention of V_0's assignment of v. Because the elements of Γ don't mention v as a free variable (and this is the only variable on which f_1 is allowed to differ from f_2 $t_0(\gamma)$ cannot change when we go from considering V_0 to V_1: by the translation lemma we have $V_1 \geq V_0 \rightarrow (t_1[\gamma_i] \leftrightarrow t_0[\gamma_i])$, hence $t_1[\Gamma]$. Observation 3 in chapter 10.3, says that derivability facts don't change when we swap subscripts, so we have $\mathscr{V}(V_1), t_1[\Gamma] \vdash t_1(\theta)$ because we have $\mathscr{V}(V_0), t_0[\Gamma] \vdash t_0(\theta)$. Thus we can deduce that $t_1(\theta)$, as desired.

Thus, $\mathscr{V}(V_0), t_0[\Gamma] \vdash V_1 \geq_{\mathbf{v}} V_0 \rightarrow t_1(\theta)$. However, as $\mathscr{V}(V_0), t_0[\Gamma]$ are content restricted to V_0 by $\Box$I we derive $\mathscr{V}(V_0), t_0[\Gamma] \vdash \Box_{V_0}[V_1 \geq_{\mathbf{v}} V_0 \rightarrow$

$t_1(\theta)]$.

The last two clauses we need to check can be handled very quickly, once we prove the following lemma.

Lemma 11.2.3. *For any formula ϕ, $\vdash \mathscr{V}(V_i) \wedge f_i(v) = f_i(v') \to (t_i(\phi) \leftrightarrow t_i(\phi'))$, where ϕ' is obtained from ϕ by replacing zero or more occurrences of v (not necessarily all such uses) with v', provided that no bound variables are replaced, and all substituted occurrences of v' are free.*

Proof. I will argue by induction on the complexity of the formula ϕ.

This fact is immediate in the base case where ϕ has the form $v = v'$ or $v \in v'$, (for some variables v and v', not necessarily distinct).

It is also immediate that if this claim holds for ρ, ψ then it holds for $\phi = \neg\rho$, $\phi = \rho \wedge \psi$ etc.

Suppose $\phi = (\exists x)(\psi)$ and the result above is true for ψ. There are two possibilities to consider.

First consider the case where x is neither v_1 or v_2. In this case we have $t_i(\phi') = \Diamond_{V_i}[V_{i+1} \geq_x V_i \wedge t_{i+1}(\psi')]$. To establish the bi-conditional it is enough to show we can derive $t_i(\phi')$ from $t_i(\phi)$ and vice versa. We do this by entering the diamond context, importing $f_i(v_1) = f_i(v_2)$ and using the inductive hypothesis to infer $t_{i+1}(\psi)$ from $t_{i+1}(\psi')$ (and vice versa).

If $x = v_1$ then $\phi[v_1/v_2] = (\exists v_1 \psi)[v_1/v_2] = \phi$.[4] If $x = v_2$, by hypothesis v_1 may not appear free in ϕ least v_2 be bound in $\phi[v_1/v_2]$, so again $\phi[v_1/v_2] = \phi$.

Exactly analogous reasoning works when $\phi = (\forall x)(\psi)$.

$\square$

[4] Remember that $[v_1/v2]$ only substitutes v_1 for v_2 in cases where v_1 occurs free.

Next, suppose that our proof tree of height $n+1$ concludes with an application of ∀E.

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

Then there is proof with height at most n of $\Gamma \vdash \forall v\theta$. From this we must establish $\mathscr{V}(V_0), t_0(\Gamma) \vdash t_0(\theta(v|v'))$.

So assume $\mathscr{V}(V_0), t_0(\Gamma)$. By the inductive hypothesis we can infer $\Box_{V_0}[V_1 \geq_{\mathbf{v}} V_0 \to t_1(\theta)]$. We note that, holding V_0 fixed it is logically possible that $V_1 \geq_{\mathbf{v}} V_0$ and $f_1(v)$ to $f_0(v')$. We can import $\Box_{V_0}[V_1 \geq_{\mathbf{v}} V_0 \to t_1(\theta)]$ into this logical possibility context and thereby infer $\Diamond_{V_0} t_1(\theta)$.

Because $f_1(v) = f_0(v') = f_1(v')$, we can derive that $t_1(\theta(v|v'))$ must also be true in this scenario by the lemma above. The fact that v is free for v' in θ means that all the substituted instances of v' are free (so the requirements for this Lemma 11.2.3 are fulfilled). Now we can use the Translation Lemma to derive that $t_0(\theta(v|v'))$ is also true in this scenario, as follows. Recall that Translation Lemma tells us that when $v_1 ... v_k$ are all the free variables in θ we have $\vdash V_n \geq_{set} V_m \wedge f_n = f_m(v_1) \wedge ... f_n(v_k) = f_m(v_k) \to (t_n(\phi) \leftrightarrow t_m(\phi))$.

By definition, $\theta(v|v')$ substitutes v' for every free instance of v in θ. So v (the only variable on which f_1 and f_0 can disagree) never occurs free in $\theta(v|v')$. So we can derive any sentence of this form in any context, including inside our current $\Diamond$ context. Thus we can derive that $t_0(\theta(v|v')) \leftrightarrow t_1(\theta(v|v'))$, and hence infer that $t_0(\theta(v|v'))$.

Thus we have $\Diamond_{V_0} t_0(\theta(v|v'))$. Finally we can use the fact that $t_0(\theta(v|v'))$ is content-restricted to V_0, f_0 to infer the truth of $t_0(\theta(v|v'))$ simpliciter, as desired.

Finally, consider the possibility that our proof of height $n+1$ concludes with an application of =E.

(=E) If $\Gamma_1 \vdash v_1 = v_2$ and $\Gamma_2 \vdash \theta$, then $\Gamma_1, \Gamma_2 \vdash \theta'$, where θ' is obtained from θ by replacing zero or more occurrences of v with v' [(so unlike in $\theta(v|v')$ not all free instances of v must be replaced)], provided that no bound variables are replaced, and all substituted occurrences of v' are free.

In this case, we have a proof with height 1 for $\Gamma_1 \vdash_{FOL} v_1 = v_2$ and one of height n for $\Gamma_1, \vdash_{FOL} \theta$. So by inductive hypothesis, we have $\mathscr{V}(V_0), t_0(\Gamma_1) \vdash t_0(v = v')$, i.e., $f_0(v) = f_0(v')$ and $\mathscr{V}(V_0), t_0(\Gamma_1)) \vdash t_0(\theta)$.

So assume $\mathscr{V}(V_0), t_0(\Gamma_1), t_0(\Gamma_2)$. Then we can derive $t_0(\theta)$ and $f_0(v_1) = f_0(v_2)$, by the fact just noted. We need to show that $t_0(\theta')$ can be derived. As required by the Lemma above, we know that θ' is obtained from θ by replacing zero or more occurrences of v_1 with v_2, where all substituted occurrences of v_2 are free. Thus we can prove $\mathscr{V}(V_0) \rightarrow (t_0(\theta) \leftrightarrow t_0(\theta'))$ from empty premises. So we can can conclude that $t_0(\theta')$.

Thus we have $\mathscr{V}(V_0), t_0(\Gamma_1)), t_0(\Gamma_2)) \vdash t_0(\theta')$, as desired.

$\square$