1.1

A relational structure is an ordered quadruple.

1. The first object is the (non-empty) domain, A.
2. The second object is a set of n-ary relations on A, where n depends on the element number of the relation in the set.
3. The third object is a set of n-ary functions on A, where n depends on the element number of the function in the set.
4. The fourth object is a set of elements of A. These are called "distinguished elements".

It may be remarked that the third set is a subset of the second set, and the fourth a subset of the first. If 3 and 4 are empty, it is said to be a purely relational structure.

Suppose now that 2 is a single-element set with only the binary relation of equality. Bridge comments: "To include the equality relation in the presentation of a structure is, in a sense, superfluous. It does not imply any 'structure' in the underlying domain. However, in section 1.2 we introduce a language associated with a given structure which depends intrinsically (on 1, 2, 3 and 4)…and two different languages arise depending on whether or not equality is included." **(1)**

An example is given wherein a relational structure with non-empty 1/2/3/4 is reduced to one with non-empty 1/2/4 and ultimately a purely relational structure, i.e. non-empty 1/2 only. From a mathematical viewpoint, remarks Bridge, these relational structures are identical; but as relational structures they are distinct. **(1)**

In construing as a relational structure a field with many different types of objects, replacing the, say, binary functions with ternary relations gives us a way to distinguish between them by 'imposing additional structure'. (We have the power to say that there exists no operation involving the addition of a vector and a scalar.) **(2)**

The ordered triplet consisting of

1. the set K producing the distinguished elements
2. The function coding the arity of the relations (λ)
3. The function coding the arity of the functions (μ)

defines the *type* of relational structure. If they are identical, the two structures are said to be of the *same type*.

If there are bijections taking us from this triplet for one structure to the triplet for another, these two structures are called of the *similar type.* "…can be readily re-indexed so as to become structures of the same type." **(3)**

A relational structure S is a *substructure* of the relational structure R if:

1. The domain of S (call it B) is a subset of the domain of R (call it A).
2. The intersection of the nth relation in set #2 of R with the nth power of B gives a set which is effectively a manifestation of that very relation such that it is restricted to work within the bounds of the substructure. The *restriction* of R_n on B. These restricted versions of those relations should constitute set #2 of S.
3. It follows that the set #3 of S is a subset of the set #3 of R.
4. Each distinguished element of R is an element of B.

If S is a substructure of R, R is an *extension* of S.

A mapping from the domain of one relational structure to the domain of another relational structure is a *homomorphism* if the same mapping also maps from the sets #2, #3 and #4 of the first relational structure to those of the other. In Bridge's words, "*h* is a mapping which preserves structure."

Question: Suppose *h*: C->D is a homomorphism; C, D are relational structures. Is the image of C under *h,* C(h), a substructure of D?

Answer: Yes. Let C be (A, R, f, K) and D be (B, S, g, J). Now, we are given a homomorphism h: C->D. Note that the range is always a subset of the co-domain. Take A(h) to be the range; this is then a subset of B, the co-domain. Arguing similarly for the other sets, we conclude that C(h) is a substructure of D.

A homomorphism is an *isomorphism* if its inverse exists. (Bridge has not used that particular term ('inverse'), however, so perhaps it is not mathematically rigorous to call it that.)

If the range is equal to the co-domain for an isomorphic mapping, it is called an *embedding.*

To prove: That h: C->D is an isomorphism and h^{-1}: D->C iff h is a bijection and the conditions for being a homomorphism are satisfied.

Proof: Suppose h is a bijection. Then h^{-1} exists. h(A)=B=>h^{-1}o h(A)= h^{-1} (B)=> h^{-1} (B)=A (by definition of inverse). Therefore, h^{-1} is a mapping from B to A. We argue similarly for set #2, #3 and #4 and our proof is then complete.

Question: Give an example of non-isomorphic structures C, D of the same type such that there exists a homomorphism *h* which is a one-one map from the domain A of C to the domain B of D.

Answer: Let A=N, B=Q'. Since the cardinality of N is less than that of Q', no mapping from one to the other can be surjective. Thus, the any structures with these domains are non-isomorphic. We can now pick any arbitrary one-one homomorphism *h* between A and B and provide the desired structures C and D.

<u>1.2</u>

First-order language consists of

1. Variables (enumerably many) **(4)**
2. Constants (enumerably many)
3. Predicates (Relations)
4. Functions
5. Logical connectives: Negation+Implication/Negation+Conjunction/…
6. The universal quantifier
7. Brackets

The language depends only on the structure (or rather, the *type* of structure) ("a one-one correspondence"). A structure is a *realization* of a language.

The smallest set X which includes all variables and constants, alongside all those functions whose domain and range is a subset of the set of all constants and variables (equivalently, of the set of all terms), forms the set of all terms, X.

Subformulae are defined. Its action is to break up a formula into its atomic constituents in a reversible manner. (ϕ can be reconstructed if subform(ϕ) is given.) "(A formula) is atomic iff its subform consists of a single element."

Remark: Well-formed formulae could have been defined equally well with the negation and any other logical connective in place of conjunction, and with the existential instead of the universal quantifier. **Regimentation is merely an act of convenience. (5)**

Scope (exhibited occurrence), bound and free variables are defined.

An interesting analogy: “The effect of the (definite) integral sign is similar to that of the quantifier…it binds.”

A term cannot be substituted into a variable if another variable in the term ends up being ‘accidentally’ bound upon substitution.

A formula with no free variables is a sentence. The *closure* of a formula consists of binding its free variables with universal quantifiers to turn it into a sentence.

<u>1.3</u>

“A first order language L, even though associated with a given (type of) structure, is technically a syntactic object with no semantic significance.”

Interpretation: An object in the language L is taken to a correspondent in the structure S; a formula becomes an assertion about the structure.

Given an *assignment* (a sequence of element(s) from the domain A), Bridge defines how to interpret a formula (what each term in ϕ denotes in S). Variables are replaced by corresponding elements in the assignment, constants and functions turn into that which they are in one-one correspondence with (and the latter changes also in tandem with how the terms in its argument change).

Note: An assignment is the same as the ‘domain of interpretation’ that was spoken of by the professor.

An assignment is said to *satisfy* an atomic formula in the backdrop of a structure iff the assigned formula is an element of the set of all relations in S (set #2). Various lemmas are subsequently stated.

Exercises 1.12, answers:

a) If $q_0<q_1$
b) If $q_1<q_2$

Theorem: An assignment satisfies a formula iff imposing those terms in it denoting its free variables onto any other assignment leads always to another assignment which satisfies the formula. (Only the free variables affect the truth-value of a formula under an assignment.)

Proven by induction on the length of the formula. Base step: This is verified to be true for an atomic formula. Induction step: Assume the theorem to be true for all subformulae of a formula. Now, since we have committed to three items in our language (negation, conjunction and universal quantifier), we show one by one that the theorem holds for formulae of those forms on the strength of our induction hypotheses, and then we are done. **(5)**

Either a sentence or its negation satisfies all assignments. (Either a sentence is true, or its negation is.) Since a sentence has no free variables, this follows immediately from the previous theorem.

Lemma: An assignment satisfies a formula iff replacing a bound variable with another variable not in the formula yields a new formula which is satisfied by the assignment.

Proven again by induction.

Question: Construct a formula which is satisfied by an assignment with free occurrences of a variable such that replacement of a bound variable in it with that one yields a new formula which is not satisfied by the assignment.

Answer: In the domain Q, there exists an x such that x is strictly less than y.

Lemma: If a term is free for a variable, then an assignment satisfies the formula yielded by substituting the variable with the term iff the original formula is satisfied by an assignment wherein the interpretation of that variable is set to be the same as the interpretation of the substituted term.

Proof: Suppose that *k* variables occur in the term t. Since *t* is free for v_i , there no free occurrences of it within the scope of a quantifier. However, there may be bound occurrences of the other variables in *t* in the formula. Using the previous lemma, we substitute these variables with some others not occurring in it; and then we may proceed normally to prove the given lemma by induction.

Question: Let A, B be structures of the same type and h: A->B be an isomorphism. Show that if a formula is satisfied under A by an assignment (a_1, a_2…) then the formula is satisfied under B by the assignment ($h(a_1)$, $h(a_2)$…).

Answer: If the formula is satisfied under A by the given assignment, it means that the denoted formula is an element of the set of all relations in A. By definition, an isomorphism maps an element from a given structure to a corresponding element in the other structure and so under the action of *h,* our formula becomes an element of the set of all relations in B.

Exercise 1.16, answers:

a) If x is bound, our result follows from lemma 1.25. ("The meaning of ϕ in any given interpretation is unchanged.")
   If x is free: The if-statement already tells us that the formula is satisfied for all assignments.
b) Since our relational structure does not have equality but only ordering, it is intuitively evident that, between two assignments, two elements are equivalent if they follow the same order. Theorem 1.23's requirement of a_i=b_i may then be weakened to an ordering requirement rather than an equality requirement and we may proceed with the proof.

Some definitions:

1. If a formula is satisfied in a structure/realization of a language for all assignments, it is said to be *valid.* The structure/realization is then a *model* for the formula. **(6)**
2. If a formula is valid in all realizations of a language, it is *universally valid.*
3. A formula is *satisfiable* if there is for some realization of a language some sequence such that it is satisfied.
4. A formula is *refutable* if its negation is satisfiable.

Question: Show that a formula is universally valid iff its negation is not satisfiable.

Answer: This means that there exists no realization in the language which satisfies the formula's negation for an assignment. Converting the existential quantifier to a universal, we obtain that for all realizations, the formula is valid, and we are done.

Question: Show that a formula is valid in a realization iff its universally quantified version is; and that it is satisfiable in a realization iff its existentially quantified version is.

Answer: Second one: Choose your assignment such that each element in the assignment is the element asserted to exist in the existentially quantified version of the formula. We have thus found an assignment which satisfies the formula in the given realization and the proof is complete. First one: Follows from definition 1.21, iv).

"There are many similarities between tautologies and universally valid formulae, but there is one very important difference. There is a simple procedure to determine whether ϕ is a tautology. It is sufficient to examine the truth table of ϕ. In general, however there is no finite procedure to determine of a given first order formula whether or not it is universally valid."

In general, validity, universal validity, satisfiability and refutability are undecidable. [Church (1936)]

Some cases in which they are decidable:

Lemma 1.28: Given a formula ϕ there are only finitely many non-isomorphic realizations of the language of ϕ with a domain of (finite) cardinality N.

Proof: The maximum number of interpretations is found, and found to be finite, by computing permutations and combinations. And if there exist only finitely many interpretations, it follows that there exist only finitely many realizations.

In certain cases, then, a formula's validity is decidable: A formula is valid in a realization with finite domain iff it is satisfied by all the finitely many interpretations.

Lemma 1.29: Suppose $\phi \in$ Form(L) contains only k unary predicate letters and no functions or constants. Then ϕ is universally valid iff ϕ is valid in all domains with $\leq 2^k$ elements.

Proof: Suppose the negation of a formula is satisfied by an assignment under a given realization, that is, ϕ is not valid under a given realization. Then there exists another realization with domain $\leq 2^k$ elements in which ϕ is not valid. First, we define an equivalence relation on the domain of the given realization, which is that two elements are related if they *both* belong to a given predicate. Define a set A* which is the set of all equivalence classes. Evidently, there are at most 2^k elements in this set, since an element is either in a predicate or not in it, and so there are at most that many equivalence classes. This is our domain. We then redefine our predicates such that R^*_i is the set of equivalence classes of all the elements which lie in R_i. Induction now tells us that ~ ϕ is satisfied by this new structure. This tells us that if ϕ is not valid in some structure then it is not valid in some structure with $\leq 2^k$ elements in its domain. Moreover, obviously, if ϕ is not valid in some structure with $\leq 2^k$ elements in its domain, it is not valid in some structure. Inverting the implication, we obtain the desired lemma.

Note: The proof fails if the formula contains functions or constants because ~ ϕ may then be satisfied in some structure and yet not be satisfied in any structure with $\leq 2^k$ elements in its domain. Take ~ϕ as "A proper subset X of A is in one-one correspondence with A." over the domain N=A. (This can never be satisfied over a finite domain!)

Lemma 1.30: A formula without any quantifiers is valid in a realization iff it is valid in its substructure.

Proof: The assignment is a sequence of A if it is a sequence of A*, by hypothesis. By induction, the denotation of this assignment in A* is equal to its denotation in A. We then use induction again on the length of the formula to prove the lemma.

Note: If the formula contains quantifiers, the proof fails at the second induction. The atomic case deviates from the preliminary result on terms and is false in general, since there may exist a b in A (which would then not be in A*, since the formula is valid in the substructure) which, when denoted into the bound term, does not satisfy the assignment. (The universal quantifier requires that all b satisfy.)

Theorem: Suppose a sentence is of the following form: Existential quantifiers on n variables, followed by the universal quantifier on y vector, followed by a ψ such that ψ contains no quantifiers, functions or constants. If this sentence is satisfiable, then it is satisfiable in a domain with at most n elements.

Proof: Take a realization with *k* predicates and domain A in which the sentence is valid. Then there exists some sequence in the domain which, when substituted into the *n* variables with the existential quantifier, yields a valid sentence. Furthermore, any sequence in the domain may be substituted in place of y vector with the validity of the sentence retained. The sentence is now quantifier-free, and application of lemma 1.30 tells us that this final sentence is also valid in all substructures of this realization.

Define a substructure with domain containing the *n* elements substituted into the variables. Since the sentence is also valid in this realization and all its substructures, we conclude that all structures with ≤n elements in their domain are models for the sentence.

Corollary: If the sentence is of the form described above, it is universally valid iff it is valid in all domains with at most n elements.

Proof: Apply the previous theorem to ~sentence.

Definition: If the union of a set of formulae with ϕ is a subset of the set of all formulae in a realization of a language and all formulae in said set are satisfied for a given sequence under a given realization, the set of formulae *logically implies* ϕ and ϕ is satisfied for this sequence under the given realization. A logical consequence of the empty set is a universally valid formula.

Material implication v/s logical implication: The former is intra-model, the latter is inter-model. (Syntactic v/s semantic.)

Question: Show that, in theorem 1.31, if the sentence is permitted to have m distinct constants, it is satisfiable in a domain with at most m+n elements. Show also that the theorem fails if the sentence has functions.

Answer: Obviously, alongside the *n* elements denoted by the existentially quantified variables, we need *m* more for there to be something to denote for the m constants, and so the result extends to a domain of at most m+n elements. Moreover, the theorem fails in the presence of functions because we can have a function with an arity greater than n.

Question: Show that ϕ I-- ψ iff I-- ϕ->ψ.

Answer: From the premise, whenever/if ϕ is satisfied, so is ψ. Then the conclusion also follows from the definition.

Exercise 1.20, answers:

a) Yes
b) No

Exercise 1.21, answer: Use the procedure described at the end of lemma 1.28 (which will work because the domain is finite, as prescribed in theorem 1.31) for all m+1 formulas and lemma 1.22, b).

Question: Show that, if a universal formula is valid under a structure, it is valid under all of its substructures; and that if an existential formula is valid under a structure, it is valid under all of its superstructures.

Answer: Merely emulate the proof of lemma 1.30. In the universal case: The issue that comes up with quantifiers pointed out in the note does not hold, since if the formula is satisfied for all b in A, it is also satisfied for all A* C A. In the existential case: If there is a sequence in A* which satisfies the formula, there will also be a sequence in A which satisfies, since A* C A.

<u>Points of interest/queries:</u>

1) It is observed that two relational structures and languages can be distinct and yet identical from a mathematical viewpoint. Is this significant? Does this allow us to distinguish between objects mathematics is as such blind to, such as between the modulus function and the root-of-square function?
2) It was said that purely relational structures are technically simpler but with reduced complexity than the original structure. How is it that we see the converse occurring here?
3) What is the relationship between being of *similar type* and being *isomorphic*?
4) What was the difficulty which not imposing this condition entailed?
5) The anthropocentric guider of convenience/simplicity crops up again.
6) How do we determine the mapping which takes an object in the language to its correspondent in the realization?
7) Let it be noted that notions such as ‘proven’ and ‘proof by induction’ rely on there being an “unambiguous segment of language” (in the professor’s words) which I am in suspicion of.
8) The backdrop in which the statement made may be said to be true.

<u>**2.1**</u>

Rules of inference (here, modus ponens and generalization) and a certain set of axioms are taken.

Suppose the union of a set of formulae with ϕ is a well-formed formula. Then, a *derivation* of ϕ from this set is a finite sequence of formulae which ends with ϕ such that the other formulae in the sequence are either axioms, or a part of the set, or is an immediate consequence of the previous ones by modus ponens or generalization (note that the quantified variable is not to occur free in the set).

ϕ is a *theorem* of the *predicate* calculus if it is derivable from the empty set. Else, it is a theorem of the set in question (and all its supersets). This should also inspire the idea of a minimal, finite set from which the formula is derivable. This set will consist of formulae in the derivation which occur neither as axioms nor as immediate consequences of our rules of inference, but only as the second condition. **(1)**

Determining whether a given formula is an axiom is decidable. Determination of whether a given sequence is a derivation is decidable.

We now call a certain set of our axioms *tautologies.* Furthermore, a formula is an *instance of a tautology* if it can be produced by replacing each statement letter of a tautology with some formula.

Exercise 2.1, answers:

a) Yes, from axiom 1.
b) No.
c) Yes, from axiom 4.

Lemma: Each instance of a tautology is universally valid.

Proof: We observe that, for negation and conjunction, satisfaction mimics truth value. Now, define an assignment of truth values to the statement letters of the tautology from which the instance of a tautology was formed such that it yields T iff the formula substituting the statement letter is satisfied. It follows that all subformulae of the instance built up only from the substituted formulae are satisfied iff the corresponding subformulae of the tautology yields T under an assignment of truth values. Since the tautology itself yields T for all assignments of truth values, it follows that the instance of a tautology being considered is satisfied for all realizations and assignments.

It is desirable to take all instances of tautologies as axioms instead of just tautologies, since our implicit aim with formalization is to generate as theorems all universally valid formulae, and we wish then to be able to prove all instances of tautologies. Note that determining whether a given formula is a tautology is decidable (truth tables).

Lemma: Each instance of a tautology is a theorem. Moreover, the derivation of this formula consists of either axioms (1-5) or of a formula which is an immediate consequence of the ones preceding it by modus ponens.

Proof: Suppose the instance of a tautology being considered is yielded by substituting a set of formulae in place of the statement letters of a certain tautology. The completeness theorem (proven later) tells us that this tautology is derivable by a sequence of the form described above. Replace the statement letters in the tautological formulae in this derivation by whatever subformulae in order to obtain the corresponding formulae in the derivation for the instance, and the statement letters in the others by some arbitrary formula in order to obtain their corresponding formulae in the derivation for the instance. Then the derivation sequence for the instance becomes of the required form.

Question: Show that the converse of the second assertion of the above lemma is true, i.e. all derivations of that form yield an instance of a tautology.

Answer: Going back from the supplied derivation of the instance to the original derivation of the tautology will prove the converse. This can be done by reversing the substitution operation performed to prove the lemma.

Any schemata which can generate all instances of tautologies can be called *complete*.

A universally valid formula is not automatically an instance of a tautology.

Lemma: Each instance of A6 and A7 is universally valid.

Proof for A6: Follows from lemma 1.22 b), lemma 1.26 and definition of universal validity.

Proof for A7: Utilize definition 1.21 iv) followed by lemma 1.22 b) on the condition. Then apply lemma 1.25 on the first formula in the condition followed by lemma 1.22 b) and the definition of universal validity and we are done.

"Thus in order to construct a formal system in which all universally valid formulae may be derivable it is necessary to have as axioms formulae which are not instances of tautologies or rules other than modus ponens (or both)."

2.2

The Soundness Theorem: Suppose the union of a set of formulae with ϕ is a well-formed formula. Then, if ϕ is a theorem of (derivable from) the set, the set entails ϕ.

Proof: Consider the minimal set which ϕ is a theorem of. We use induction on the length of the derivation for ϕ. If a formula in the derivation happens to be an axiom, it is universally valid and thus entailed by the set. If the formula happens to be a part of the set, it is trivially entailed by the set. What is left to be shown is that the remaining formulae in the derivation present by virtue of R1 or R2 are also entailed by the set, and then we are done.

R1: Take a realization in which the set is satisfied for some assignment. By the induction hypothesis, the formulae which gives us the one in question are also satisfied and so by definition, the formula given by R1 is satisfied. Definition 1.33 enables us to wrap it up.

R2: Suppose a formula in the derivation is yielded by R2 from a previous formula in the derivation by replacing the k^{th} variable quantified. Since it does not occur free in the set, if the set is satisfied by any assignment in a realization, it is satisfied for all assignments yielded by replacing the k^{th} element of this one by anything else (lemma 1.25); and so the antecedent formula is satisfied for all assignments yielded by replacing the k^{th} element of this by anything else. This completes the proof.

Definition: A set is consistent if it never implies both a formula and its negation.

Corollary: The empty set is consistent.

Proof: Application of theorem 2.6 to the contrary statement contradicts definition 1.21.

Definition: A *theory* T in L is a set of sentences of L which is *deductively closed.* (For each sentence S, T entails S/S is derivable from/a theorem of T iff S is an element of T.)

Lemma: A set is inconsistent iff it implies all well-formed formulae.

Corollary: A theory is consistent iff it is unequal to the set of all sentences in L.

The Deduction Theorem: If Σ U {ϕ, ψ} is a well-formed formula and Σ U {ϕ} I— ψ then Σ I— ϕ--> ψ.

Proof: If the condition is true, then either ψ is a theorem of Σ or whenever ψ is a theorem of some subset of Σ U {ϕ}, ϕ is a formula in this subset. In the first case, the result follows by modus ponens on A1 as a theorem of Σ. In the second case, we use induction on the derivation of ψ from the subset. The first two formula-types are easy enough.

R1: Write out the induction hypothesis for the two formulae which need to be true (and so here, theorems of the set of the first two types) for applying modus ponens. The required formula then turns out to be an instance of A2. Apply R1 twice on this and receive the desired result.

R2: Write out the induction hypothesis for the formula which needs to be true to apply generalization. R2 followed by A7 yields the result.

Exercise 2.3, answers:

a) We know that this statement is true for the notion of satisfiability. Now we only need to show that if a formula is satisfiable, it is consistent. Satisfiability means that there is a model for the formula, and the completeness theorem tells us that a set of formulae are consistent iff they have a model: In other words, iff all of them are satisfiable.
b) In light of the above discovery, it suffices to prove this statement for satisfiability; that is, all the formulae in a set of well-formed formulae are satisfiable iff all the formulae in all its finite subsets are satisfiable. The proof for this is trivial.

Question: If a formula is built up from a set of formulae using only propositional connectives, show that if each formula being in bi-implication with another formula is a theorem, the formula built up by substituting those other formulae in the original formula being in bi-implication with the original formula is a theorem.

Answer: Note that in the proof for lemma 2.3, a variant of this was used: If the antecedent is true, the corresponding subformulae being in bi-implication with each other is a theorem. Taking it forward from here is easy enough (the soundness theorem allows us to take the truth-value/satisfaction equivalence over to derivability).

Exercise 2.5, answers:

a) Follows from R2 (generalization) that if ϕ, then ϕ*.
b) A constant will always be free for substitution for the free occurrence of a variable. Therefore, it follows from A6 that if ϕ, then ϕ'.
c) First, we take A7 to be true. Then we know that there exists at least one consistent model wherein both ϕ and ϕ* are valid. Now, take A7 as false, i.e., take as our axiom ~A7. We still obtain the result of there being at least one consistent model wherein both the formulaes are valid. This tells us that A7 is independent of the rest of the axioms.
d) Taking A6 to be false gives us a model where ϕ is valid but ϕ' isn't; taking it to be true gives us a model wherein both are true. Essentially, we obtain that there is at least one consistent model in both cases. We are therefore done.

2.4

Take a mapping f between the wff of the language and a co-domain which is a subset S of the wff such that ϕ is a theorem iff f(ϕ) is. Take another mapping h with the same domain and co-domain S' such that ϕ<->h(ϕ) is a theorem.

Semantically, the difference between the two properties is as follows: The first is a bi-implication between two theorems. The second is a theorem about a bi-implication.

Lemma: The set of all the sentences in a language L is a possible subset S (i.e. there exists a map between it and Form(L) satisfying the above condition). A possible map consists of universally quantifying all the free variables in the formula. The bi-implication follows from R2/A6+R1.

Question: Show that there exists a formula in the language for which the bi-implication between a formula and its universally quantified version is not a theorem.

Answer: We need to show that the second property does not hold for the given map on sent(L). One way around, the implication will always hold by R2. To make it fail the other way round, keep a variable in the formula quantified existentially and so not free for any term. A6 can no longer be applied.

The *prenex normal form* is defined.

Theorem: The second (stronger) property holds for PNF(L).

Proof: A set of six fairly easily derivable theorems are listed out which enable us to 'pull out' the quantifiers in a formula by interconversions and allow us to end up at the prenex normal form. The theorem is then proven.

Exercise 2.6, answers:

a) Convert the if-then in terms of negation and conjunction. Turn the universal to an existential using i) and reconvert to if-then.
b) Convert the if-then in terms of negation and conjunction. Turn the existential to a universal using ii) and reconvert to if-then.

Etc.

Exercise 2.8: We need to show here that all quantifierless formulas are the disjunction of the conjunctions of a set of formulas which are either atomic or the negation of an atomic. Easy proof by induction on the length of the formula. The only case remaining: Negation outside. Pull it in.

Exercise 2.9: Convert the if-then in terms of negation and conjunction. Change the existential to a universal. We obtain the negation of a contradiction.

Exercise 2.10, answers:

a) Replace c in each step of derivation of psi with exists (x).
b) Refer to the bit in lemma 3.5.

Exercise 2.11: One with all means all for one. All for one does not mean one with all.

The Principle of Duality: We are given ϕ to be satisfied by some assignment in a realization M. ϕ^D obtained by interchanging conjunction with disjunction and universal quantifier with existential is satisfied by some assignment in a realization M^D obtained by changing the relations of M such that each relation is replaced by its negation-space with respect to the domain. Furthermore, ϕ is a theorem iff ~ ϕ^D is a theorem.

Proof: For #1, we use induction on the length of the formula. An atomic formula is merely a predicate between elements; we have no quantifiers or logical connectives. Say this is satisfied by some assignment in M. The negation-space predicate of this predicate is therefore not satisfied by this assignment. The negation of the negation-space predicate of this predicate is therefore satisfied. The proof is done for the atomic case.

If ϕ is ~ψ: The statement holds for ψ by the induction hypothesis. However, since our statement is a biimplication between ψ and ~ψ, it can be inverted and stated for ~ψ without much ado.

If ϕ is ϕ_1 & ϕ_2: Our transformed formula becomes 'ϕ'_1 or ϕ'_2'; alternatively, '~(~ ϕ'_1 & ~ϕ'_2)'. ϕ_1 and ϕ_2 as well as their negations adhere to this statement, by the induction hypothesis. Definition 1.2, iii) finishes the job for us.

If ϕ is a universally quantified formula, $Ax\psi$: Our transformed formula becomes $Ex\psi'$. By hypothesis, the statement holds for ψ; furthermore, it does so for any assignment on it. It is thus shown for this kind of formula as well.

The second part can be similarly done.

2.5

L_E are languages with the special, lone predicate of equality. Two extra axioms are added in its honor.

These two extra axioms really only enforce the special predicate as an equivalence relation (the next lemma proves its symmetry and transitivity, and its reflexivity was an axiom; we obtain as a corollary that it is an equivalence relation). It can be conceivably interpreted as the special equivalence relation of equality; however, it *can* also be interpreted otherwise.

Theorem: A9 holds in L_E for all formulae, not just atomic ones.

Proof: Induction on the length of the formula.

"A consequence of theorem 2.18 is that any model for a theory in L_E can be contracted to one in which = is interpreted by equality and the sets of sentences valid in the two models are precisely the same." Let us see how this is done to gain clarity on this statement.

Definition: The *normal contraction* of a realization in L_E is a realization of the same *type* as the original one constructed by reidentifying each element of the realization (the elements in the domain, the elements in the predicates, the input-output of the functions, the distinguished elements) with the equivalence class induced by the equivalence relation =. We add that a realization is called *normal* iff = is interpreted by equality.

Theorem: The normal contraction of any realization in L_E is normal. A formula is satisfied by an assignment in a realization if the formula produced by replacing the elements in its assignment with their equivalence classes is satisfied in the normal contraction.

Proof: The first claim is straightforward; the equivalence classes of two related elements will always be equal.

For the second claim: We show first by induction on the length of a term (proving it for elements, distinguished elements and functions) that the equivalence class of an assigned term in a realization equals the term given to the equivalence class of the same assignment in the normal realization. Then we use induction on the length of the formula to prove the claim.

Corollary: A set of formulae has a model iff the set has a normal model.

Definition: We define a formula which says 'There is exactly one x such that P(x)' in the language $L_{E.}$

Pulling the quantifiers out and applying A4 with theorem 2.18 on this definition solves exercise 2.14.

Question: A definition is given. We must show that it translates to saying 'There are exactly two x such that P(x)'.

Answer: The definition tells us that there is one y which satisfies the formula and is unequal to x, and also that there exists an x which satisfies it. Observe that if something else were to satisfy it, it would just be equal to y, since all that is not x is given to be y. This completes the answer.

Question: Consider a formula which says 'There are exactly *n* x such that P(x)' in the language. Construct now a sentence which is valid in a realization iff the domain of the realization has exactly *n* elements.

Answer: Make a polynomial with *n* distinct roots and assert that there exist *n* x which satisfy the equation. Assert also that there does not exist anything which does not satisfy the equation. The sentence is made.

"In a language with equality associated with a normal structure it is possible to express properties of the structure concerning the cardinality of some subsets of its domain."

Lemma: Take a realization M with cardinality n. For all cardinals greater than n, it is possible to construct a realization with that cardinality such that the realization is also a superstructure of M. Furthermore, a formula is satisfied by an assignment in M iff it is satisfied by the same assignment in this superstructure.

The domain of the structure to the domain of the superstructure such that all the elements in the superstructure which are not in the structure are taken to some fixed element in the structure; the mapping is the identity for the rest. Ensure then that each element in the domain of the superstructure has the same properties as its image in the substructure. For a special superstructure like this, we may proceed to complete the proof by induction.

Corollary: If a formula is satisfied by an assignment in some realization for a language, it is also satisfied by an assignment for a superstructure of this realization such that the superstructure has a cardinality greater than *n* for any *n*.

Proof: Take the domain of the superstructure to be the union of the domain of the structure and a disjoin set X with a cardinality greater than *n*. Define a map and the other objects in the superstructure as before. Since no element in X is in A, for all b in X, b* is a_0, and we are done.

"Since there is no upper bound on the size of a model of a satisfiable formula there is no syntactic way of limiting the size of a class of structures of a certain kind."

Question: Show that any model of the formula (given) necessarily has infinite cardinality.

Answer: The existential quantifier can only be extracted after the universal, and so will be deeper in than it. Therefore, the PNF will not be of the given form.

The first and the last segments in this formula tell us that for all x, there exists some y which is necessarily unequal to x such that P(x,y). Putting z=x in the second segment, we obtain also that if P(x,y), then ~P(y,x), since otherwise we get P(x,x). It is now clear that the third segment gives us an infinitely cascading Russel-esque situation requiring an infinite domain.

Points of interest/queries:

1) Is the question of finding said minimal finite set a decidable one? Or is it undecidable?

The Completeness Theorem: A set of sentences is consistent iff it has a model.

Proof: Recall that theorem 2.6 has already told us the converse of this statement. Now for the other side.

A subset of sentences of L is *complete* if for all sentences in L, either the sentence or its negation is a theorem of the set.

“The completeness…is, as stated here, entirely a syntactic property. However, the completeness theorem for the predicate calculus implies that the property may equally well be defined in terms of semantic concepts.”

L’ is an *alphabetic extension* of L if L’ is obtained from L by adding new constant symbols only.

Take a subset of sentences in L and a subset of sentences in L’ such that the subset in L is a subset of the subset in L’. Take a formula in L with precisely one free variable the existentially quantified version of which is a theorem of the subset in L. If L’ has a constant (sometimes called a Henkin constant) such that replacing the variable with that constant makes it a theorem of the subset of L’, the subset of L’ is a *full extension* of the subset of L. If a subset is its own full extension in L, it is called *full*.

With these definitions at hand, here are the steps by which we perform the proof.

1) Any consistent subset of sentences in L can be embedded in a consistent full-extension.
2) Any consistent subset of sentences in L can be embedded in a complete consistent subset of sentences in L.
3) Any consistent subset of sentences in L can be embedded in a subset of sentences in an alphabetic extension of L, L’ such that this subset is complete, consistent and full.
4) A complete consistent full subset of sentences in L’ has a model.
5) Using 3) and 4): If a subset of sentences of L is consistent, it has a model.

Lemma 3.5: There exists a consistent full-extension for a consistent subset of sentences in L.

Proof: Define the set of all formulas with precisely one free variable whose existential quantified versions are theorems of the original subset. Now add to L fresh constants in one-one correspondence with the elements of this set. The union of the original subset and the elements of the defined set instantiated with the fresh constants is a full-extension of the original subset. Now to prove that this is consistent:

Suppose not. This means that both some formula and its negation is a theorem with finitely long derivations respectively. For each, then, there is a finite subset of the second set in the union forming the full-extension, the union of which with the original subset yields the formulae and is thus inconsistent. Therefore, the negation of this very finite subset is also a theorem of their union. Applying the deduction theorem and a tautology, we obtain that the negation of this finite subset is a theorem of the original subset.

But also, the existentially quantified version of each formula in the finite subset is also a theorem of the original subset. It is straightforward enough to show that this is contradicted by the negation of the finite subset being a theorem. We obtain a reductio ad absurdum the presumed consistency of the original subset, and we are done.

Lemma 3.6: If the negation of a sentence is not a theorem of the union of a subset with that sentence, said union is consistent.

Proof: By lemma 2.10, the subset is consistent from the if-condition. Suppose the union is inconsistent. Then the negation of the sentence is a theorem. Emulate the previous lemma to show that the negation of a sentence is a theorem also of the subset. This is a contradiction.

1) is done.

Lemma 3.7: There exists a complete and consistent extension to a consistent subset of sentences in a countable language L.

Proof: An example is given, essentially.

2) is done.

Theorem: There exists a complete, full and consistent extension to a subset of sentences in a countable language L.

Proof: Lemmas 3.5 and 3.7 are used alternatingly and countably often to construct countable languages and sentence sets with countable cardinality. Take the union of the languages and the union of the sets constructed by lemma 3.7 upon the sets constructed by lemma 3.5. This subset-language pair have the desired properties.

3) is done.

Definition: The *canonical structure* determined by a subset of sentences in a language with at least one constant is as follows.

a) Domain: Closed terms of L
b) Relations: n-tuples of closed terms which are a theorem of the subset under some predicate
c) Functions: n-ary functions which are in L
d) Distinguished elements: Constants in L

Theorem: A sentence is valid in the canonical structure for a consistent, complete and full subset of sentences iff the sentence is a theorem of said subset.

Proof: By induction on the sentence length. Note that we must thus modify the induction hypothesis and can assume truth of the theorem only for sub-sentences.

Start from assuming the if. The atomic case follows from the definitions. For negation: Use definition 1.21, the induction hypothesis and the consistency + completeness of the subset. For conjunction: Definition 1.21, the induction hypothesis and A4.

For the final case: First apply definition 1.21 and the induction hypothesis, followed by A6. (The sentence not being a theorem of the subset implies its negation is, which contradicts our assumption post application of theorem 2.15)

4) is done.

Theorem: Any consistent subset of sentences in a countable language has a model.

Proof: Theorem 3.8 tells us that there is an alphabetical extension of this language which has a complete consistent full subset of sentences which extends the original subset. Theorem 3.10 tells us that this extension has a model. The reduct of this model to the original language is the required model.

The Compactness Theorem: A subset of sentences in a countable language has a model iff every finite subset of it has a model. (Proven by contradiction.)

3.3

For languages with equality, if a set has a model, then it also has a normal model.

Another canonical structure is defined, but with one modification: The term is replaced by its equivalence class under the equality relation.

Theorem: In a countable language with equality, a subset is consistent iff it has a (normal) model.

The language made to give our subset a model was made using a countable sequence of countable languages by adjoining at each stage a countable set of new constants. Thus, the final language must also be countable.

The previous theorem can now be strengthened in the following way: In a countable language with equality, a subset of sentences is consistent iff it has a countable model.

Another restating of the completeness theorem:

“A countable consistent set of sentences is satisfiable in a subset of the natural numbers (with suitable relations, functions and individual constants).”

3.4

In order to extend lemma 3.7 to languages with a cardinality greater than that of N, the axiom of choice must be assumed: “Any set can be well ordered.”

Lemma 3.17: “Suppose k is an infinite cardinal. Then k = k*(aleph-naught) = k*k.”

Refer elsewhere for a proof of the above two.

Corollary: If a first-order language L has cardinality k, then the set of sentences of L has cardinality k.

Lemma 3.19: There exists a complete and consistent extension to a consistent subset of sentences in any first-order language L.

Proof: The set of sentences of L can be well-ordered. Transfinite induction is used to complete the proof in analogy to lemma 3.7’s proof.

Question: Prove the above by applying Zorn’s lemma to the partially ordered set of consistent extensions of the subset in L.

Answer: The solution lies in the fact that each sentence not in the subset will be the maximal element for some partially ordered set. We pick them out in this way and add them if its negation is not already in the subset.

The remainder of the completeness theorem’s proof is unchanged, since that was the only place where L’s countability was directly exploited.

The Generalized Completeness Theorem: If a subset of sentences in a language (or a language with equality) with cardinality k is consistent, then the subset has a model (or a normal model) with cardinality no greater than k.

Theorem: A countable consistent set of quantifier-free formulae of L is satisfiable in a countable structure.

Proof: Ultimately similar to the proof of theorem 3.10 with slight differences: We take in the domain of our structure not only closed terms but *all* terms, and the last step of induction can be eliminated. (Also note that completeness is not formally defined for formulae.)

Question: Let there be a quantifier-free formula with h terms and sub-terms. Show that if it is satisfiable, then it is satisfiable in a structure with not more than h elements.

Answer: Refer to theorem 1.31.

3.5

Theorem: If a subset of sentences in a first-order language with equality has arbitrarily large finite models, it has an infinite model.

Proof: Define an alphabetical extension of the language by adjoining a countably infinite set of distinct new constants. Consider a subset of sentences of this new language which is the union of our original subset with the set of all sentences of the given form (wherein each sentence involves two of the new constants adjoined). Consider now a finite subset of *this* subset.

This finite subset is the union of some subset of our original subset and a finite set of sentences of the form given above—say, p of those sentences.

By our hypothesis, the original subset has a model with more than 2p elements. Expand this model to a realization for the alphabetical extension by adjoining distinct new constants corresponding to the (not more than) 2p new constants occurring in the finite subset. This realization would then satisfy that finite subset.

Now applying the compactness theorem: Since an arbitrary finite subset has a model, so must the subset of the alphabetical extension; furthermore, the model of the infinite subset must have an infinite domain.

The reduct of this is the required model.

Corollary: There is no set of sentences such that it is satisfied by a model iff that model is finite.

A notable contrast to the result of exercise 2.16! "It is possible to characterize (in a language with equality) all (normal) structures with a given fixed cardinality, or indeed those structures with domains of less than some fixed finite cardinality. Finiteness itself cannot be characterized in the same way even when infinite sets of sentences are considered."

Theorem: Let a first-order language with equality have a unary function, two binary functions and the constant '0'. Consider a realization of it with the domain N, the unary function interpreted as successor, the binary ones as multiplication and division and '0' being zero.

Consider the subset of sentences in the language which are valid in this realization. There exists a structure which is not isomorphic to the original realization such that this subset is valid in it.

Proof: Consider the alphabetical extension of the language obtained by adding a single new constant 'c'. Define a new set of sentences which is the union of the original set and the set of sentences which are of the form ~c=n and n is the n^{th} successor of 0.

Take an arbitrary finite subset of this new set. There then exists a k such that for all m not less than it the sentence ~c=m is not an element of this subset. The subset is then valid in a realization with domain up to k. By the compactness theorem, the new set of sentences defined has a model.

Suppose now that the reduct of this model to the original language is isomorphic to the original realization. This implies the existence of an isomorphism going from the latter to the former. So, for some integer n, f(n) = c.

But in the reduct, ~Ex=n[c] is valid and in the original model, Ex=n[n] is valid. This gives us the contradiction we need.

(The proof is based on the fact that c ends up being a ‘transfinite’ number.)

Definition: The structure <M, R> where R is a binary relation on M is a *graph* if R is irreflexive and symmetric on M. If a, b in M such that aRb then a is *connected* to b. The graph is said to be *k-colorable* if there is a partition of M into not more than k subsets such that no two connected elements are in the same subset.

Theorem: A graph is k-colorable iff every finite substructure of it is k-colorable.

Proof: It is obvious that if the graph itself is k-colorable, so is any finite substructure of it. Now for the converse.

Let us consider a first-order language of which this graph is a realization. Let the language have the required number of constants to match the graph’s domain along with k unary predicates. Consider now the following set of sentences in the language (as defined). This set has the same cardinality as M only if M is infinite; the set may exclude finitely many elements in order to satisfy its conditions.

Suppose every finite subgraph possesses the property in question. Consider a finite subset of the set of sentences defined above and the finite subset of the realization with a reduced domain corresponding to the sentence subset. By our assumption, this realization (with the original set of relations restricted to those in its domain) is k-colorable.

Consider the k disjoint subsets given to us by our coloring, and take some constant *a* in the reduced realization which acts as the interpretation for all the constants in the language which have nothing to correspond to in the realization.

The subset of sentences is modelled by the structure (as defined). Since the subset itself was arbitrary, by the compactness theorem, there exists a model for the original set of sentences.

From exercise 1.22 a), it follows that any substructure of this model is also a model for the set. Now, consider the purely relational version of the substructure whose domain is equal to the distinguished elements of the original. This is isomorphic to the original graph.

Define k disjoint subsets of the graph by an isomorphism: An element belongs to a subset (*i*) iff the corresponding constant in the language is an element of the intersection of the substructure’s domain and the corresponding predicate (*i*).

From the third condition for the set of sentences, each element belongs to at least one subset. From the fourth, this subset is unique for the element. From the fifth, we obtain that two connected elements must be in two distinct subsets.

Hence, proved.

Question: Show that no finite set of sentences exists in a first-order language with equality such that it is satisfied in a realization iff the domain of the realization is infinite.

Answer: We need to show that every set of sentences either has a finite model, or that there exists an infinite model which fails to satisfy it.

Suppose it has no finite model. If all infinite models satisfy it, take its negation, which would then be satisfied iff the model is finite. We have a contradiction to corollary 3.23.

Hence, proved.

Question: Let A, B be two theories in a first-order language such for any realization of the language, one theory is valid iff the other is not. Show that both theories are finitely axiomatizable.

Answer: The union of the two theories has no model. Therefore, by the compactness theorem, no finite subtheory of the union has a model. Let X be a finite subtheory of A and Y of B. The union of X and Y has no model. Therefore, any structure is a model of X iff it is not a model of Y. If it is not a model of Y, it is not a model of B. Therefore, any model of X is a model of A. Since X was finite, it follows that A is finitely axiomatizable (and similarly, B).

Herbrand's theorem: A closed existential formula A is universally valid iff there is a universally valid formula which is a disjunction of instantiations of A.

Significance: An existential formula is universally valid in first-order logic if the formula obtained by the disjunction of all its instantiations is universally valid in zero-order logic.

Question: Show using the compactness theorem that any partially ordered set can be totally ordered by somehow extending the former.

Answer: Induction followed by compactness.

3.6

A 'truth assignment' is defined, and on its heels, so is 'tautology'.

An axiomatization is given for propositional calculus.

Propositional consistency and propositional completeness are defined. A set of propositional formulas has a model if there exists a truth assignment which gives all its constituent formulas a T.

Theorem: A set of formulae in a propositional language is consistent iff it has a model.

Proof: The first way round—the soundness theorem's equivalent—is straightforward enough. If a formula is a theorem of a set, we can show by induction on the length of the derivation (using the first two properties of the axiomatization) that a model for the set is also a model for the formula. It follows that a set of sentences with a model must be consistent. Now for the other way round.

Once again, we make a complete consistent set from the given consistent set. Give this a canonical truth assignment: A proposition is T iff it is a theorem. The assignment is extended in the usual way to the rest of the language's formulae. We can then show that a formula is T iff it is a theorem of the set.

Corollary 3.29: A is complete/The theorems of A are precisely the tautologies; any set of formulae in the language has a model iff all its finite subsets have a model.

4.1

The Downward Löwenheim-Skolem Theorem: If a subset of sentences with infinite cardinality k in a first-order language with equality has a model, then the subset has a model of cardinality not greater than k.

Proof: If the subset has a model, it is consistent. Only if the cardinality of the language is also k can the subset have infinite cardinality k. The final result follows from theorem 3.20.

A special case of the above theorem which may be proven without assuming the axiom of choice: If the subset of sentences is countable and satisfiable, then it has a countable model.

The Skolem paradox: It is possible to construct a consistent theory T in a countable first-order language L in which the theorem "The set of real numbers is uncountable" is derivable. At the same time, the above theorem implies that T has a countable model.

Resolution: This simply means that there is no function *in the model* which is a bijection between N and R. "The elements in the domain representing the reals may indeed be countable; no contradiction occurs because the one-one function from it onto N will be outside the model."

The Upward Löwenheim-Skolem Theorem: If in a first-order language with equality of cardinality k there exists a subset of sentences with an infinite model, then the subset has a model of cardinality p for each infinite p not less than k.

Proof: Define an alphabetical extension (with cardinality p) of the language by adjoining the required set of distinct new constants. Consider a subset of sentences of this new language which is the union of our original subset with the set of all sentences of the given form (wherein each sentence involves two of the new constants adjoined). Consider now a finite subset of *this* subset.the

By hypothesis, our original subset has some infinite model. This model can in turn be expanded to a realization for the language's alphabetical extension. There then exists a model for our finite subset. By the compactness theorem, the set of sentences formed by the union also has a model.

Since the set of constants adjoined to the extension has cardinality p, the realization has cardinality not less than p. But also, by the downward Löwenheim-Skolem theorem, the set of sentences formed by the union (with cardinality p) has a model of cardinality not greater than p. This implies that the model has cardinality exactly p. Since the original sentence is also valid in this model, the proof is complete.

"Just as we may infer from theorem 3.22 that no set of sentences of a first-order language (with equality) is valid in precisely the finite realizations of the language, so we may infer from [this theorem] that there is no set of sentences characterizing the realizations of a given cardinality."

The upshot of the DLST is relativization. The upshot of the ULST is the inability to characterize a given infinity.

4.2

A model-theoretic characterization of completeness is provided.

One realization is *elementarily equivalent* to another iff for all sentences in the given first-order language, the satisfaction of the sentence in the first realization implies its satisfaction in the second.

This is an equivalence relation (for any realization in the language, a given sentence or its negation will be satisfied). Note that this property can hold between two non-isomorphic structures. (That it must hold for isomorphic ones is given by ex 1.15.)

Question: If there exists a realization of L with cardinality K not less than w, show that for each infinite cardinality there exists a realization of cardinality not less than it which is elementarily equivalent to the original realization.

Answer: For each infinite cardinality, there is a consistent model. By the completeness theorem, this means that this is a model for a complete set of sentences. By ULST, this set of sentences has a model of each infinite cardinality greater than this one; furthermore, by lemma 4.5, all these models will be elementarily equivalent.

Lemma: A subset of sentences in a first-order language is complete iff all its models are elementarily equivalent.

Proof: Either an arbitrary sentence or its negation is a theorem of the subset. By the completeness theorem, one of the two is a logical consequence of the subset. So, for any model of the subset, either the sentence is satisfied (in the first case) or its negation is (in the second case). The converse is straightforward.

A consistent theory in a first-order language without equality is *categorical* if all its models are isomorphic. The theory is *K-categorical* if it has a model of cardinal K, and if any two models of cardinal K are isomorphic.

Clearly, a categorical theory can only have finite models (under threat of violating the ULST). Since isomorphic structures are always elementarily equivalent, it follows trivially that categorical theories are always complete.

A partial converse lemma: In a first-order language without equality, a complete theory with a finite model is categorical.

Proof: Let A be the given finite model (with n elements in its domain) of the complete theory. The sentence asserting that there are n elements in its domain is satisfied in it. Since the theory is complete, this sentence is an element of it. It follows that any other model (say, B) of the theory must also have n elements in its domain. By lemma 4.5, A and B are elementarily equivalent. What remains to be shown is that A and B are isomorphic.

Construct a finite sequence of alphabetical extensions of L. We show by induction that the corresponding sequence of models (which are obtained by declaring suitably many distinguished elements from the domain) are elementarily equivalent. Finally, we construct an isomorphism between the two.

The base step of the induction is already given to us to be true ($A \equiv B$). We assume truth for some r and take the $r+1^{th}$ alphabetical extension. Suppose the models for this next extension are *not* elementarily equivalent. Fixing the fresh distinguished element in A: So, for each element in the domain of B (excluding those already selected as distinguished elements), when it is included in the next iteration of B's distinguished elements, there exists a sentence which is satisfied in A and not in B.

Using this, a sentence is constructed in the previous language of the sequence which is satisfied in (the previous iteration of) A but not in B. Thus, we prove that A and B are elementarily equivalent throughout the sequence. The remaining part of the proof, consisting of explicitly defining an isomorphism, is simple enough.

Note that this proof works only when A and B are finite (thus failing for the elementarily equivalent but non-isomorphic models in theorem 3.25.) This is because the formula defined to bring about a contradiction must be finite (and consequently, the domain of B must be finite).

Question: Show that if a set of sentences in a language with finitely many non-logical symbols is categorical, then there exists a finite subset of it such that all the sentences of the original set are theorems of the subset.

Answer: The set of sentences is in a language with finitely many non-logical symbols, but with no restriction on the number of variables, may yet be infinitely many.

But since it is categorical, follows that it has a finite model. Now we have a model of both finite domain and finitely many theorems, since the number of non-logical symbols remain only finitely many. A corresponding finite set of sentences can now be constructed.

Cantor's dense order theorem: DLO is aleph naught categorical.

Proof: We show that an order-preserving isomorphism exists between any two arbitrary countable models.

Proven by induction on a sequence of finite subsets of A and B with an isomorphism between them (A_n, B_n, h_n). Let the first triplet be the null set each and assume the n^{th} one has been defined.

If the n^{th} element of A is an element of A_n, the $n+1^{th}$ terms in the sequence remains the same. If not, choose A_{n+1} as the union of A_n with the n^{th} element of A and choose some fresh element of B (to set as its image) to be added to get B_{n+1}.

Similarly, check if the n^{th} element of B is an element of B_n and make the construction.

By induction, the union of all the isomorphisms gives us the required function.

Lemma: The DLO of reals is non-isomorphic with the DLO substructure obtained by deleting 0 from it.

Proof: Let us assume there exists an order-preserving isomorphism h going from the former to the latter.

Consider the sequence p_n in A which is the pre-image to the sequence 1/n in B. Since order is preserved, *both* sequences are decreasing with increasing n, and since R is complete, $\{p_n\}$ converges to some p, and $h(p) < 0$.

There exists an r such that $h(p) < r < 0 < h(p_n) = 1/n$ for every n.

But since p is the greatest lower bound of the sequence, $h^{-1}(r) < p$. This contradicts h being order-preserving.

Thus, no such h can exist.

Corollary: DLO is not categorical for the cardinality of reals.

Theorem: If a theory is categorical for some uncountable cardinality, it is categorical for each uncountable cardinality.

The proof has been omitted here.

Theorem: A theory in a language of cardinality k with no finite models is complete if it is categorical for some infinite cardinality not less than k.

Proof: Suppose not, i.e. T is incomplete but categorical for some infinite cardinality p not less than k. Then there exists a sentence in L such that neither it nor its negation is a theorem of T. By lemma 3.6, the union of T with either the sentence or its negation is consistent. Models exist for each of these extensions of T, neither of which are finite (by hypothesis); thus, models exist of cardinality p (by ULST). But since T is p-categorical, both the models are isomorphic.

But isomorphic structures satisfy precisely the same sentences. We have the required contradiction.

Corollary: (Q, >) and (R, >) are elementary equivalent. (Both are models for DLO, which is a complete theory.)

An explicit example is given of a theory which is not k-categorical for any infinite cardinality k.

4.3

Definition: A substructure is an *elementary substructure* of its superstructure if the satisfaction of a given formula under a certain assignment in the substructure implies its satisfaction in the superstructure. The superstructure is its *elementary extension*.

Definition: An embedding between the substructure and the superstructure is an *elementary embedding* if the range is an elementary substructure of the co-domain. If such an embedding exists, the substructure is said to be *elementarily embeddable* in the superstructure.

Lemma: Let f: A→B be a map between the two structures.

a) f is an embedding iff for every atomic or negated atomic formula in L, its satisfaction by an assignment a in A implies its satisfaction by the assignment f(a) in B.
b) f is an elementary embedding iff for every formula in L, its satisfaction by an assignment a in A implies its satisfaction by the assignment f(a) in B.

Proof: The proof of a) is left as an exercise. Now for b): Suppose f is an elementary embedding. Since the domain and the range are isomorphic, ex 1.15's result combined with the definition of elementary substructure allows us to conclude the result we want. Now for the other way round.

Suppose the satisfaction by an assignment a in A implies its satisfaction by the assignment f(a) in B. Using this assumption, it is established that f is an isomorphism; subsequently, it follows from definition that its range is an elementary substructure of the co-domain, and the proof is complete.

Even when a substructure turns out to be elementarily equivalent to its superstructure, three possibilities remain: The former is an elementary substructure of the latter; the former is *not* and elementary substructure of the latter but an elementary embedding exists between the two; neither is the former an elementary substructure of the latter nor does an elementary embedding exist between the two. An example is provided of each case.

Note regarding example 2: The given structures are elementarily equivalent because we have not introduced any constants into our language (and so one cannot obtain a 'sentence' by replacing x with 1.)

Definition: Define an alphabetical extension of L obtained by adjoining to L a new individual constant for each element in the domain of a realization for it. The new structure corresponding to the expanded language (called here the 'expanded structure') has an intended interpretation for each constant.

Lemma: Given f: A→B, the mapping is an elementary embedding between two structures iff the expanded structure of the domain is elementarily equivalent to the expanded structure of the co-domain.

Proof: The first way round: Use lemma 4.21, b) on the assumption and apply lemma 1.26 on the formulae (after appropriately associating each missing constant in the smaller language with a variable), and the proof is done. The converse is similarly proven.

Corollary: A substructure is an elementary substructure of its superstructure iff the expanded structures of the two (in the same domain) are elementarily equivalent.

Definition: A characterization of the number of alternations in quantifiers in the PNF of a formula is defined.

A way to procure all the wff of a language on the basis of this definition is given.

Definition: A substructure is an n-elementary substructure of its superstructure if the elementary substructure condition holds specifically for all the formulae of universal-PNF number n (as defined previously).

Definitions: The *open diagram* of a realization of a language is the set of all formulae in its extended structure such that the formula is satisfied in it, and is either an atomic sentence or its negation in the extended language.

The *complete diagram* of the realization is the set of all sentences in its extended structure such that the sentence (which must be an element of the set of all sentences in the extended language) is satisfied in it.

Evidently, the open diagram of a model is a subset of its complete diagram, and the latter is complete (being the set of all true sentences in a structure).