```
Proof: By induction on t.
Case: t = true
... show P(true) ...
Case: t = false
... show P(false) ...
Case: t = ift_1 then t_2 else t_3
... show P(ift_1 then t_2 else t_3), using P(t_1), P(t_2), and P(t_3) ...
(And similarly for the other syntactic forms.)
```

For many inductive arguments (including the proof of 3.3.3), it is not really worth writing even this much detail: in the base cases (for terms t with no subterms) P(t) is immediate, while in the inductive cases P(t) is obtained by applying the induction hypothesis to the subterms of t and combining the results in some completely obvious way. It is actually *easier* for the reader simply to regenerate the proof on the fly (by examining the grammar while keeping the induction hypothesis in mind) than to check a written-out argument. In such cases, simply writing "by induction on t" constitutes a perfectly acceptable proof.

3.4 Semantic Styles

Having formulated the syntax of our language rigorously, we next need a similarly precise definition of how terms are evaluated—i.e., the *semantics* of the language. There are three basic approaches to formalizing semantics:

1. *Operational semantics* specifies the behavior of a programming language by defining a simple *abstract machine* for it. This machine is "abstract" in the sense that it uses the terms of the language as its machine code, rather than some low-level microprocessor instruction set. For simple languages, a *state* of the machine is just a term, and the machine's behavior is defined by a *transition function* that, for each state, either gives the next state by performing a step of simplification on the term or declares that the machine has halted. The *meaning* of a term t can be taken to be the final state that the machine reaches when started with t as its initial state.²

^{2.} Strictly speaking, what we are describing here is the so-called *small-step* style of operational semantics, sometimes called *structural operational semantics* (Plotkin, 1981). Exercise 3.5.17 introduces an alternate *big-step* style, sometimes called *natural semantics* (Kahn, 1987), in which a single transition of the abstract machine evaluates a term to its final result.

It is sometimes useful to give two or more different operational semantics for a single language—some more abstract, with machine states that look similar to the terms that the programmer writes, others closer to the structures manipulated by an actual interpreter or compiler for the language. Proving that the behaviors of these different machines correspond in some suitable sense when executing the same program amounts to proving the correctness of an implementation of the language.

2. *Denotational semantics* takes a more abstract view of meaning: instead of just a sequence of machine states, the meaning of a term is taken to be some mathematical object, such as a number or a function. Giving denotational semantics for a language consists of finding a collection of *semantic domains* and then defining an *interpretation function* mapping terms into elements of these domains. The search for appropriate semantic domains for modeling various language features has given rise to a rich and elegant research area known as *domain theory*.

One major advantage of denotational semantics is that it abstracts from the gritty details of evaluation and highlights the essential concepts of the language. Also, the properties of the chosen collection of semantic domains can be used to derive powerful laws for reasoning about program behaviors—laws for proving that two programs have exactly the same behavior, for example, or that a program's behavior satisfies some specification. Finally, from the properties of the chosen collection of semantic domains, it is often immediately evident that various (desirable or undesirable) things are impossible in a language.

3. *Axiomatic semantics* takes a more direct approach to these laws: instead of first defining the behaviors of programs (by giving some operational or denotational semantics) and then deriving laws from this definition, axiomatic methods take the laws *themselves* as the definition of the language. The meaning of a term is just what can be proved about it.

The beauty of axiomatic methods is that they focus attention on the process of reasoning about programs. It is this line of thought that has given computer science such powerful ideas as *invariants*.

During the '60s and '70s, operational semantics was generally regarded as inferior to the other two styles—useful for quick and dirty definitions of language features, but inelegant and mathematically weak. But in the '80s, the more abstract methods began to encounter increasingly thorny technical problems,³ and the simplicity and flexibility of operational methods came

^{3.} The *bête noire* of denotational semantics turned out to be the treatment of nondeterminism and concurrency; for axiomatic semantics, it was procedures.



Figure 3-1: Booleans (B)

to seem more and more attractive by comparison—especially in the light of new developments in the area by a number of researchers, beginning with Plotkin's Structural Operational Semantics (1981), Kahn's Natural Semantics (1987), and Milner's work on CCS (1980; 1989; 1999), which introduced more elegant formalisms and showed how many of the powerful mathematical techniques developed in the context of denotational semantics could be transferred to an operational setting. Operational semantics has become an energetic research area in its own right and is often the method of choice for defining programming languages and studying their properties. It is used exclusively in this book.

3.5 Evaluation

Leaving numbers aside for the moment, let us begin with the operational semantics of just boolean expressions. Figure 3-1 summarizes the definition. We now examine its parts in detail.

The left-hand column of Figure 3-1 is a grammar defining two sets of expressions. The first is just a repetition (for convenience) of the syntax of terms. The second defines a subset of terms, called *values*, that are possible final results of evaluation. Here, the values are just the constants true and false. The metavariable v is used throughout the book to stand for values.

The right-hand column defines an evaluation relation⁴ on terms, written

^{4.} Some experts prefer to use the term *reduction* for this relation, reserving *evaluation* for the "big-step" variant described in Exercise 3.5.17, which maps terms directly to their final values.

 $t \rightarrow t'$ and pronounced "t evaluates to t' in one step." The intuition is that, if t is the state of the abstract machine at a given moment, then the machine can make a step of computation and change its state to t'. This relation is defined by three inference rules (or, if you prefer, two axioms and a rule, since the first two have no premises).

The first rule, E-IFTRUE, says that, if the term being evaluated is a conditional whose guard is literally the constant true, then the machine can throw away the conditional expression and leave the then part, t_2 , as the new state of the machine (i.e., the next term to be evaluated). Similarly, E-IFFALSE says that a conditional whose guard is literally false evaluates in one step to its else branch, t_3 . The E- in the names of these rules is a reminder that they are part of the evaluation relation; rules for other relations will have different prefixes.

The third evaluation rule, E-IF, is more interesting. It says that, if the guard t_1 evaluates to t'_1 , then the whole conditional if t_1 then t_2 else t_3 evaluates to if t'_1 then t_2 else t_3 . In terms of abstract machines, a machine in state if t_1 then t_2 else t_3 can take a step to state if t'_1 then t_2 else t_3 if *another* machine whose state is just t_1 can take a step to state t'_1 .

What these rules do *not* say is just as important as what they do say. The constants true and false do not evaluate to anything, since they do not appear as the left-hand sides of any of the rules. Moreover, there is no rule allowing the evaluation of a then- or else-subexpression of an if before evaluating the if itself: for example, the term

```
if true then (if false then false else false) else true
```

does not evaluate to if true then false else true. Our only choice is to evaluate the outer conditional first, using E-IF. This interplay between the rules determines a particular *evaluation strategy* for conditionals, corresponding to the familiar order of evaluation in common programming languages: to evaluate a conditional, we must first evaluate its guard; if the guard is itself a conditional, we must first evaluate *its* guard; and so on. The E-IFTRUE and E-IFFALSE rules tell us what to do when we reach the end of this process and find ourselves with a conditional whose guard is already fully evaluated. In a sense, E-IFTRUE and E-IFFALSE do the real work of evaluation, while E-IF helps determine where the work is to be done. The different character of the rules is sometimes emphasized by referring to E-IFTRUE and E-IFFALSE as *computation rules* and E-IF as a *congruence rule*.

To be a bit more precise about these intuitions, we can define the evaluation relation formally as follows.

3.5.1DEFINITION: An *instance* of an inference rule is obtained by consistently replacing each metavariable by the same term in the rule's conclusion and all its premises (if any). П

For example,

if true then true else (if false then false else false) \rightarrow true

is an instance of E-IFTRUE, where both occurrences of t₂ have been replaced by true and t_3 has been replaced by if false then false else false.

- 3.5.2 DEFINITION: A rule is *satisfied* by a relation if, for each instance of the rule, either the conclusion is in the relation or one of the premises is not.
- 3.5.3 DEFINITION: The *one-step evaluation* relation \rightarrow is the smallest binary relation on terms satisfying the three rules in Figure 3-1. When the pair (t, t') is in the evaluation relation, we say that "the evaluation statement (or judgment) $t \rightarrow t'$ is derivable."

The force of the word "smallest" here is that a statement $t \rightarrow t'$ is derivable iff it is justified by the rules: either it is an instance of one of the axioms E-IFTRUE and E-IFFALSE, or else it is the conclusion of an instance of rule E-IF whose premise is derivable. The derivability of a given statement can be justified by exhibiting a *derivation tree* whose leaves are labeled with instances of E-IFTRUE or E-IFFALSE and whose internal nodes are labeled with instances of E-IF. For example, if we abbreviate

 $s \stackrel{\text{def}}{=} if true then false else false$

 $t \stackrel{\text{def}}{=} \text{if s then true else true}$

 $u \stackrel{\text{def}}{=}$ if false then true else true

to avoid running off the edge of the page, then the derivability of the statement

if t then false else false \rightarrow if u then false else false

is witnessed by the following derivation tree:

$$\frac{\overline{s \rightarrow false}}{t \rightarrow u}^{E-IFTRUE} E-IF}_{E-IF}$$
if t then false else false \rightarrow if u then false else false

Calling this structure a tree may seem a bit strange, since it doesn't contain any branches. Indeed, the derivation trees witnessing evaluation statements

will always have this slender form: since no evaluation rule has more than one premise, there is no way to construct a branching derivation tree. The terminology will make more sense when we consider derivations for other inductively defined relations, such as typing, where some of the rules do have multiple premises.

The fact that an evaluation statement $t \rightarrow t'$ is derivable iff there is a derivation tree with $t \rightarrow t'$ as the label at its root is often useful when reasoning about properties of the evaluation relation. In particular, it leads directly to a proof technique called *induction on derivations*. The proof of the following theorem illustrates this technique.

3.5.4 THEOREM [DETERMINACY OF ONE-STEP EVALUATION]: If $t \rightarrow t'$ and $t \rightarrow t''$, then t' = t''.

Proof: By induction on a derivation of $t \rightarrow t'$. At each step of the induction, we assume the desired result for all smaller derivations, and proceed by a case analysis of the evaluation rule used at the root of the derivation. (Notice that the induction here is not on the length of an evaluation sequence: we are looking just at a single step of evaluation. We could just as well say that we are performing induction on the structure of t, since the structure of an "evaluation derivation" directly follows the structure of the term being reduced. Alternatively, we could just as well perform the induction on the derivation of $t \rightarrow t''$ instead.)

If the last rule used in the derivation of $t \rightarrow t'$ is E-IFTRUE, then we know that t has the form if t_1 then t_2 else t_3 , where $t_1 = t$ rue. But now it is obvious that the last rule in the derivation of $t \rightarrow t''$ cannot be E-IFFALSE, since we cannot have both $t_1 = t$ rue and $t_1 = f$ alse. Moreover, the last rule in the second derivation cannot be E-IF either, since the premise of this rule demands that $t_1 \rightarrow t'_1$ for some t'_1 , but we have already observed that true does not evaluate to anything. So the last rule in the second derivation can only be E-IFTRUE, and it immediately follows that t' = t''.

Similarly, if the last rule used in the derivation of $t \rightarrow t'$ is E-IFFALSE, then the last rule in the derivation of $t \rightarrow t''$ must be the same and the result is immediate.

Finally, if the last rule used in the derivation of $t \rightarrow t'$ is E-IF, then the form of this rule tells us that t has the form if t_1 then t_2 else t_3 , where $t_1 \rightarrow t'_1$ for some t'_1 . By the same reasoning as above, the last rule in the derivation of $t \rightarrow t''$ can only be E-IF, which tells us that t has the form if t_1 then t_2 else t_3 (which we already know) and that $t_1 \rightarrow t''_1$ for some t''_1 . But now the induction hypothesis applies (since the derivations of $t_1 \rightarrow t'_1$ and $t_1 \rightarrow t''_1$ are subderivations of the original derivations of $t \rightarrow t'$ and

 $t \rightarrow t''$), yielding $t'_1 = t''_1$. This tells us that $t' = ift'_1$ then t_2 else $t_3 = ift''_1$ then t_2 else $t_3 = t''$, as required.

3.5.5 EXERCISE [★]: Spell out the induction principle used in the preceding proof, in the style of Theorem 3.3.4. □

Our one-step evaluation relation shows how an abstract machine moves from one state to the next while evaluating a given term. But as programmers we are just as interested in the final results of evaluation—i.e., in states from which the machine *cannot* take a step.

3.5.6 DEFINITION: A term t is in *normal form* if no evaluation rule applies to it i.e., if there is no t' such that $t \rightarrow t'$. (We sometimes say "t is a normal form" as shorthand for "t is a term in normal form.")

We have already observed that true and false are normal forms in the present system (since all the evaluation rules have left-hand sides whose outermost constructor is an if, there is obviously no way to instantiate any of the rules so that its left-hand side becomes true or false). We can rephrase this observation in more general terms as a fact about values:

3.5.7 THEOREM: Every value is in normal form.

When we enrich the system with arithmetic expressions (and, in later chapters, other constructs), we will always arrange that Theorem 3.5.7 remains valid: being in normal form is part of what it *is* to be a value (i.e., a fully evaluated result), and any language definition in which this is not the case is simply broken.

In the present system, the converse of Theorem 3.5.7 is also true: every normal form is a value. This will not be the case in general; in fact, normal forms that are not values play a critical role in our analysis of *run-time errors,* as we shall see when we get to arithmetic expressions later in this section.

3.5.8 THEOREM: If t is in normal form, then t is a value.

Proof: Suppose that t is not a value. It is easy to show, by structural induction on t, that it is not a normal form.

Since t is not a value, it must have the form ift_1 then t_2 else t_3 for some t_1 , t_2 , and t_3 . Consider the possible forms of t_1 .

If $t_1 = true$, then clearly t is not a normal form, since it matches the left-hand side of E-IFTRUE. Similarly if $t_1 = false$.

If t_1 is neither true nor false, then it is not a value. The induction hypothesis then applies, telling us that t_1 is not a normal form—that is, that there is some t'_1 such that $t_1 \rightarrow t'_1$. But this means we can use E-IF to derive $t \rightarrow if t'_1$ then t_2 else t_3 , so t is not a normal form either.

It is sometimes convenient to be able to view many steps of evaluation as one big state transition. We do this by defining a multi-step evaluation relation that relates a term to all of the terms that can be derived from it by zero or more single steps of evaluation.

- 3.5.9 DEFINITION: The *multi-step evaluation* relation \rightarrow^* is the reflexive, transitive closure of one-step evaluation. That is, it is the smallest relation such that (1) if $t \rightarrow t'$ then $t \rightarrow^* t'$, (2) $t \rightarrow^* t$ for all t, and (3) if $t \rightarrow^* t'$ and $t' \rightarrow^* t''$, then $t \rightarrow^* t''$.
- 3.5.10 EXERCISE $[\star]$: Rephrase Definition 3.5.9 as a set of inference rules.

Having an explicit notation for multi-step evaluation makes it easy to state facts like the following:

- 3.5.11 THEOREM [UNIQUENESS OF NORMAL FORMS]: If $t \rightarrow^* u$ and $t \rightarrow^* u'$, where u and u' are both normal forms, then u = u'.
 - *Proof:* Corollary of the determinacy of single-step evaluation (3.5.4).

The last property of evaluation that we consider before turning our attention to arithmetic expressions is the fact that *every* term can be evaluated to a value. Clearly, this is another property that need not hold in richer languages with features like recursive function definitions. Even in situations where it does hold, its proof is generally much more subtle than the one we are about to see. In Chapter 12 we will return to this point, showing how a type system can be used as the backbone of a termination proof for certain languages.

Most termination proofs in computer science have the same basic form:⁵ First, we choose some well-founded set *S* and give a function *f* mapping "machine states" (here, terms) into *S*. Next, we show that, whenever a machine state t can take a step to another state t', we have f(t') < f(t). We now observe that an infinite sequence of evaluation steps beginning from t can be mapped, via *f*, into an infinite decreasing chain of elements of *S*. Since *S* is well founded, there can be no such infinite decreasing chain, and hence no infinite evaluation sequence. The function *f* is often called a *termination measure* for the evaluation relation.

3.5.12 THEOREM [TERMINATION OF EVALUATION]: For every term t there is some normal form t' such that $t \rightarrow^* t'$.

Proof: Just observe that each evaluation step reduces the size of the term and that size is a termination measure because the usual order on the natural numbers is well founded.

^{5.} In Chapter 12 we will see a termination proof with a somewhat more complex structure.

3.5.13 EXERCISE [RECOMMENDED, $\star\star$]:

1. Suppose we add a new rule

if true then
$$t_2$$
 else $t_3 \rightarrow t_3$ (E-FUNNY1)

to the ones in Figure 3-1. Which of the above theorems (3.5.4, 3.5.7, 3.5.8, 3.5.11, and 3.5.12) remain valid?

2. Suppose instead that we add this rule:

$$\frac{t_2 \rightarrow t'_2}{\text{if } t_1 \text{ then } t_2 \text{ else } t_3 \rightarrow \text{if } t_1 \text{ then } t'_2 \text{ else } t_3} \qquad (\text{E-FUNNY2})$$

Now which of the above theorems remain valid? Do any of the proofs need to change? $\hfill \Box$

Our next job is to extend the definition of evaluation to arithmetic expressions. Figure 3-2 summarizes the new parts of the definition. (The notation in the upper-right corner of 3-2 reminds us to regard this figure as an extension of 3-1, not a free-standing language in its own right.)

Again, the definition of terms is just a repetition of the syntax we saw in §3.1. The definition of values is a little more interesting, since it requires introducing a new syntactic category of *numeric values*. The intuition is that the final result of evaluating an arithmetic expression can be a number, where a number is either 0 or the successor of a number (but not the successor of an arbitrary value: we will want to say that succ(true) is an error, not a value).

The evaluation rules in the right-hand column of Figure 3-2 follow the same pattern as we saw in Figure 3-1. There are four computation rules (E-PREDZERO, E-PREDSUCC, E-ISZEROZERO, and E-ISZEROSUCC) showing how the operators pred and iszero behave when applied to numbers, and three congruence rules (E-SUCC, E-PRED, and E-ISZERO) that direct evaluation into the "first" subterm of a compound term.

Strictly speaking, we should now repeat Definition 3.5.3 ("the one-step evaluation relation on arithmetic expressions is the smallest relation satisfying all instances of the rules in Figures 3-1 *and* 3-2..."). To avoid wasting space on this kind of boilerplate, it is common practice to take the inference rules as constituting the definition of the relation all by themselves, leaving "the smallest relation containing all instances..." as understood.

The syntactic category of numeric values (nv) plays an important role in these rules. In E-PREDSUCC, for example, the fact that the left-hand side is pred (succ nv_1) (rather than pred (succ t_1), for example) means that this rule cannot be used to evaluate pred (succ (pred 0)) to pred 0, since this



Figure 3-2: Arithmetic expressions (NB)

would require instantiating the metavariable nv_1 with pred 0, which is not a numeric value. Instead, the unique next step in the evaluation of the term pred (succ (pred 0)) has the following derivation tree:

$$\frac{1}{pred \ 0 \rightarrow 0} \xrightarrow{\text{E-PREDZERO}} \text{E-SUCC}$$
succ (pred 0) \rightarrow succ 0
pred (succ (pred 0)) \rightarrow pred (succ 0)

3.5.14 EXERCISE [******]: Show that Theorem 3.5.4 is also valid for the evaluation relation on arithmetic expressions: if $t \rightarrow t'$ and $t \rightarrow t''$, then t' = t''.

Formalizing the operational semantics of a language forces us to specify the behavior of *all* terms, including, in the case at hand, terms like pred 0 and succ false. Under the rules in Figure 3-2, the predecessor of 0 is defined to be 0. The successor of false, on the other hand, is not defined to evaluate to anything (i.e., it is a normal form). We call such terms *stuck*.

3.5.15 DEFINITION: A closed term is *stuck* if it is in normal form but not a value. \Box

"Stuckness" gives us a simple notion of *run-time error* for our simple machine. Intuitively, it characterizes the situations where the operational semantics does not know what to do because the program has reached a "meaningless state." In a more concrete implementation of the language, these states might correspond to machine failures of various kinds: segmentation faults, execution of illegal instructions, etc. Here, we collapse all these kinds of bad behavior into the single concept of "stuck state."

3.5.16 EXERCISE [RECOMMENDED, *******]: A different way of formalizing meaningless states of the abstract machine is to introduce a new term called wrong and augment the operational semantics with rules that explicitly generate wrong in all the situations where the present semantics gets stuck. To do this in detail, we introduce two new syntactic categories

badnat ::=	non-numeric normal forms:
wrong	run-time error
true	constant true
false	constant false
badbool ::=	non-boolean normal forms:
wrong	run-time error
nv	numeric value

and we augment the evaluation relation with the following rules:

if badbool then t_1 else $t_2 \rightarrow wrong$	(E-IF-WRONG)
succ badnat \rightarrow wrong	(E-Succ-Wrong)
pred badnat \rightarrow wrong	(E-Pred-Wrong)
iszero badnat → wrong	(E-IsZero-Wrong)

Show that these two treatments of run-time errors agree by (1) finding a precise way of stating the intuition that "the two treatments agree," and (2) proving it. As is often the case when proving things about programming languages, the tricky part here is formulating a precise statement to be proved—the proof itself should be straightforward.

3.5.17 EXERCISE [RECOMMENDED, *******]: Two styles of operational semantics are in common use. The one used in this book is called the *small-step* style, because the definition of the evaluation relation shows how individual steps of computation are used to rewrite a term, bit by bit, until it eventually becomes a value. On top of this, we define a multi-step evaluation relation that allows us to talk about terms evaluating (in many steps) to values. An alternative style,

called *big-step* semantics (or sometimes *natural semantics*), directly formulates the notion of "this term evaluates to that final value," written $t \Downarrow v$. The big-step evaluation rules for our language of boolean and arithmetic expressions look like this:

$\vee \Downarrow \vee$	(B-VALUE)
$\frac{\texttt{t}_1 \Downarrow \texttt{true} \qquad \texttt{t}_2 \Downarrow \texttt{v}_2}{\texttt{ift}_1 \texttt{then} \texttt{t}_2 \texttt{else} \texttt{t}_3 \Downarrow \texttt{v}_2}$	(B-IFTRUE)
$\frac{t_1 \Downarrow false t_3 \Downarrow v_3}{\text{if } t_1 \text{ then } t_2 \text{ else } t_3 \Downarrow v_3}$	(B-IFFALSE)
$\frac{\texttt{t}_1 \Downarrow \texttt{nv}_1}{\texttt{succ } \texttt{t}_1 \Downarrow \texttt{succ } \texttt{nv}_1}$	(B-SUCC)
$\frac{\mathtt{t}_1 \Downarrow \mathtt{0}}{\mathtt{pred}\mathtt{t}_1 \Downarrow \mathtt{0}}$	(B-PredZero)
$\frac{\mathtt{t}_1 \Downarrow \mathtt{succ} \mathtt{nv}_1}{\mathtt{pred} \mathtt{t}_1 \Downarrow \mathtt{nv}_1}$	(B-PredSucc)
$\frac{\texttt{t}_1 \Downarrow \texttt{0}}{\texttt{iszero}\texttt{t}_1 \Downarrow \texttt{true}}$	(B-IszeroZero)
$\frac{t_1 \Downarrow succ nv_1}{szero t_1 \Downarrow false}$	(B-IszeroSucc)

Show that the small-step and big-step semantics for this language coincide, i.e. $t \rightarrow^* v$ iff $t \Downarrow v$.

3.5.18 EXERCISE [★★ +→]: Suppose we want to change the evaluation strategy of our language so that the then and else branches of an if expression are evaluated (in that order) *before* the guard is evaluated. Show how the evaluation rules need to change to achieve this effect. □

3.6 Notes

The ideas of abstract and concrete syntax, parsing, etc., are explained in dozens of textbooks on compilers. Inductive definitions, systems of inference rules, and proofs by induction are covered in more detail by Winskel (1993) and Hennessy (1990).