

# Introduction to Computer Science via Scheme

Luke Hunsberger

© 2019 Luke Hunsberger  
All rights reserved

Spring 2019



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>I</b>	<b>Non-Destructive Programming in Scheme</b>	<b>5</b>
<b>2</b>	<b>Scheme Expressions vs. Scheme Data</b>	<b>7</b>
2.1	Primitive Data Expressions . . . . .	7
2.1.1	Numbers . . . . .	7
2.1.2	Booleans . . . . .	8
2.1.3	The Empty List (or Null) . . . . .	9
2.1.4	The <i>Void</i> Datum . . . . .	9
2.1.5	Symbols . . . . .	9
2.2	String Expressions . . . . .	10
2.3	Summary . . . . .	10
<b>3</b>	<b>Evaluating Scheme Data</b>	<b>11</b>
3.1	Evaluating Numbers, Booleans, the Empty List, the <i>void</i> Datum, and Strings . . . . .	11
3.2	Evaluating Symbols . . . . .	12
3.3	Summary . . . . .	14
<b>4</b>	<b>Introduction to DrScheme</b>	<b>15</b>
4.1	Entering Expressions into the Interactions Window . . . . .	16
4.2	DrScheme’s <i>Run</i> Button . . . . .	17
4.3	Summary . . . . .	17
<b>5</b>	<b>Built-In Functions</b>	<b>19</b>
5.1	Built-in Functions for Arithmetic . . . . .	19
5.2	Contracts . . . . .	20
5.3	Built-in Functions for Integer Arithmetic . . . . .	21
5.4	The Built-in <code>eval</code> Function . . . . .	21
5.5	The Built-in Functions <code>printf</code> and <code>void</code> . . . . .	22
5.5.1	The <code>printf</code> Function . . . . .	22
5.5.2	The <code>void</code> Function . . . . .	22
5.6	Summary . . . . .	23
<b>6</b>	<b>Non-Empty Lists</b>	<b>25</b>
6.1	The Syntax and Semantics for Non-Empty Lists . . . . .	25
6.2	Evaluating Non-Empty Lists: the Default Rule . . . . .	26
6.2.1	A More Formal Description of the Default Rule . . . . .	29
6.3	Summary . . . . .	31

<b>7</b>	<b>Special Forms</b>	<b>33</b>
7.1	The <code>define</code> Special Form . . . . .	34
7.1.1	The Syntax of the <code>define</code> Special Form . . . . .	34
7.1.2	The Semantics of the <code>define</code> Special Form . . . . .	34
7.2	The <code>quote</code> Special Form . . . . .	37
7.2.1	The Syntax of the <code>quote</code> Special Form . . . . .	37
7.2.2	The Semantics of the <code>quote</code> Special Form . . . . .	37
7.2.3	Alternate Syntax for <code>quote</code> Special Forms . . . . .	39
7.3	Summary . . . . .	39
<b>8</b>	<b>Predicates</b>	<b>41</b>
8.1	Type-Checker Predicates . . . . .	41
8.2	Comparison Predicates . . . . .	43
8.3	Summary . . . . .	44
<b>9</b>	<b>Defining Functions</b>	<b>47</b>
9.1	Defining Functions vs. Applying Them to Inputs . . . . .	47
9.2	The <code>lambda</code> Special Form . . . . .	48
9.3	The Syntax and Semantics of Lambda Expressions . . . . .	50
9.3.1	The Syntax of a Lambda Expression . . . . .	50
9.3.2	The Semantics of a Lambda Expression . . . . .	51
9.3.3	Applying a <code>lambda</code> Function to Input Values . . . . .	52
9.4	Summary . . . . .	57
<b>10</b>	<b>Some practicalities</b>	<b>59</b>
10.1	More Fun with the Built-in <code>printf</code> Function . . . . .	59
10.2	The Built-in <code>load</code> Function . . . . .	62
10.3	Comments . . . . .	63
10.4	Summary . . . . .	65
<b>11</b>	<b>Conditional Expressions I</b>	<b>67</b>
11.1	The <code>if</code> Special Form . . . . .	67
11.2	Simplifying Conditional Expressions . . . . .	70
11.3	The <code>when</code> Special Form . . . . .	73
11.4	Summary . . . . .	75

# List of Figures

1.1	A function with no side effects (left) vs. a function with side effects (right)	2
2.1	The universe of Scheme data, partitioned according to data type	8
3.1	The evaluation function in Scheme	12
3.2	Sample evaluations in Scheme	12
3.3	A sample environment	13
4.1	A programmer's interaction with DrScheme	16
4.2	Using the <i>Run</i> button to evaluate (the data denoted by) multiple expressions	18
6.1	The evaluation of the list denoted by <code>(+ 2 3)</code>	28
6.2	DrScheme's use of different colors to distinguish side-effect printing from output values	31
7.1	The side effect of <code>define</code> : inserting a new entry into the Global Environment	35
9.1	Applying the function corresponding to <code>(lambda (x) (* x x))</code> to the value 4	53
9.2	Applying the <code>sphere-volume</code> function to the value 2	54



# Chapter 1

## Introduction

Most kinds of communication are based on some kind of language, whether written, spoken, drawn or signed. To be used successfully, the *syntax* and *semantics* of a language must be (explicitly or implicitly) understood.

- The *syntax rules* of a language specify the legal words, expressions, statements or sentences of that language.
- The *semantic rules* of a language specify what the legal words, expressions, statements or sentences *mean*.

For example, the syntax rules of the English language tell us that *person*, *tall*, *told*, *the*, *a*, *me* and *joke* are legal words, and that “*The tall person told me a joke*” is a legal sentence, whereas *pkrs*, *shrel* and *fdadfa* are not legal words, and “*Person tall told a me the*” is not a legal sentence. The semantic rules tell us what each of the words mean (e.g., what objects the nouns denote and what processes the verbs convey), as well as what the entire sentence means (e.g., that a particular tall person told me a joke). For another example, the syntax rules of French tell us that *je*, *vais*, *au*, *tableau* and *noir* are legal words, and that “*Je vais au tableau noir*” is a legal sentence; and the semantic rules tell us that this sentence means that I am going to the blackboard.

Just as people use so-called *natural languages* (e.g., English or French) to communicate with one another, people use *programming languages* to communicate with computers. Over the years, many programming languages have been introduced, having names such as *Java*, *Scheme*, *Python*, *C*, *C++*, *Fortran*, *Lisp*, *Haskell*, *Basic*, *Algol*, *Javascript*, and many others. Like any natural language, each programming language has an associated set of syntax rules that specify the legal expressions (or statements or sentences or programs) that can be used in that language, and a set of semantic rules that specify what the legal expressions *mean*.

- ★ The *meaning* of a computer program includes the data denoted by expressions, the computations to be performed on that data, and any auxiliary *actions* to be done (e.g., printing information on the computer screen or changing the value of a variable stored in the computer’s memory).

For most computer programming languages, the constituents of the language, whether they are called expressions, statements or entire programs, usually comprise sequences of typewritten characters. For example, the following character sequences are legal building blocks of a Java program:

- `int x = 5;`
- `for (int i=0; i < 5; i++) System.out.println(i);`
- `public class Sample { }`

And the following character sequences are legal building blocks of a Scheme program:

- `(define x 5)`
- `(+ 2 3)`
- `(printf "Hi there...")`

The semantic rules of Java stipulate that the legal statement, `int x = 5;`, represents an instruction to the computer to create space for a variable named `x` whose value will, at least initially, be the integer *five*. Similarly, the semantic rules of Scheme stipulate that the legal expression, `(define x 5)`, when *evaluated*, should cause the computer to create a new variable named `x` whose value will be the integer *five*.

Although people can effectively communicate with one another using a natural language based on an informal, imprecise, intuitive understanding of its syntax and semantics, trying to program a computer based on an informal, imprecise, intuitive understanding of the syntax and semantics of a given programming language typically leads to trouble. Therefore, it is important to be explicit about the syntax and semantics of the programming language being used.

- ★ Indeed, while programming, it is extremely important to have an accurate *mental model* of the computations you are effectively asking the computer to perform.

To enable us to enter the world of programming as quickly and painlessly as possible, it is helpful to use a programming language whose rules of syntax and semantics are relatively simple. Scheme is just such a language.

- ★ Although Scheme has a relatively simple *computational model* (i.e., syntax and semantics), it is as computationally powerful as any programming language.

In contrast, the Java programming language has a much more complicated set of syntax rules, and a correspondingly complicated computational model—without any theoretical increase in computational power. Therefore, in this class, we begin with Scheme.

- ★ The concepts you learn in this class will be helpful to you when learning any other programming language in the future.

In summary, to be effective, programmers need to have an accurate mental model of the operation of whatever computer they are programming. The complexity of their mental model depends in large part on the kind of programming language they are using. One of the significant advantages of the Scheme programming language is that it is based on a fairly simple *computational model*. Scheme’s computational model is based on the *Lambda Calculus* invented by the mathematician Alonzo Church in the 1930s, well before the advent of modern computers. Internalizing Scheme’s model of computation will make you an effective Scheme programmer in no time!

## Functions

Scheme is an example of a *functional* programming language. The main thing that you, as a Scheme programmer, will do is design *functions* for solving problems. For our purposes, a function is something that takes *zero or more* inputs, and generates a *single* output, as illustrated on the lefthand side of Fig. 1.1. For example, you might define a Scheme function whose input is a scoresheet for some game, and whose output is the sum of the scores on that scoresheet.

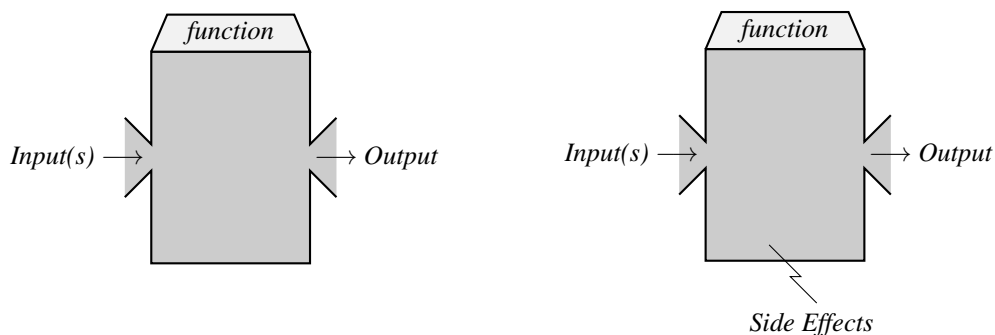


Figure 1.1: A function with no side effects (left) vs. a function with side effects (right)



In certain cases, we may also consider functions that generate *side effects*, as illustrated on the righthand side of Fig. 1.1. An example of a harmless, but very useful side effect is that of causing information to be displayed onscreen. For example, the above-mentioned function might not only compute the sum of the scores on a given scoresheet, but also have the side effect of displaying the contents of that scoresheet on a computer screen.

★ Functions that have either no side effects or only harmless side effects are called *non-destructive*.

As you will discover in Part I of this book (*Non-Destructive Programming in Scheme*) a wide variety of extremely useful computations can be performed by non-destructive functions. Furthermore, non-destructive functions tend to be very easy to write and *debug* (i.e., to find errors and fix them).

Nonetheless, as Part II (*Destructive Programming in Scheme*) reveals, there are also many areas where *destructive* functions (i.e., functions having destructive side effects) can be extremely useful. The most basic example of a destructive side effect is one that *modifies* the value assigned to a variable or to a slot within a data structure. For example, the above-mentioned function might not only compute the sum of the scores on the scoresheet, but also destructively modify the scoresheet by entering a new score into one of its slots. Although this kind of side effect may sound harmless, it can greatly complicate the task of writing and debugging functions. (For example, does the computed sum include the newly entered score?) Therefore, when we encounter destructive functions, starting with Chapter ??, we shall do so very carefully.

## A Note about the Approach

This textbook takes a very careful, bottom-up approach to the computational model of Scheme. Each of the first several chapters introduces a small portion of the computational model, highlighting the syntax and semantics of each construct that is presented. Although this approach can seem slow at first, it leads to faster results in the long run because it helps to avoid many common pitfalls that can frustrate programmers who are relying on a casual understanding of the computational model being used.



## **Part I**

# **Non-Destructive Programming in Scheme**



## Chapter 2

# Scheme Expressions vs. Scheme Data

In our daily lives, we frequently use character sequences to denote both concrete and abstract data. For example, the character sequence *dog* can be used to denote a dog; and the character sequence *34* can be used to denote the number *thirty-four*. Of course, this book itself largely consists of a bunch of character sequences that denote all sorts of things. Well, in its physical form, it is a bunch of pages that are covered with ink marks. The ink marks represent characters that, in turn, form sequences of characters that denote other things. The point is: we are so used to using character sequences to denote (or represent) things that we tend to take it for granted. When programming computers, it is important to have a solid understanding of the legal character sequences, what they mean, and what computations they may lead to.

Any program in Scheme is a sequence of (usually typewritten) characters. The syntax rules of Scheme tell us which character sequences constitute legal Scheme programs.

- ★ The building blocks of a Scheme program are character sequences called *expressions*.

In other words, each Scheme program consists of one or more Scheme expressions. For example, as we'll soon discover, `3`, `#t` and `()` are legal expressions in Scheme.

- ★ In Scheme, each legal expression denotes a datum (i.e., a piece of data).
- ★ The semantics of Scheme tells us which datum each legal expression denotes.

For example, in Scheme, the legal expressions `3`, `#t` and `()` respectively denote the following pieces of data: *the number three*, *the truth value true*, and *the empty list*.

As illustrated in Fig. 2.1, the universe of Scheme data is partitioned into *data types* having names such as: *numbers*, *booleans* (i.e., truth values), *symbols* and *functions*, among many others.

- ★ Each datum belongs to one and only one data type.

For example, a Scheme datum might be a number or a symbol, but cannot be both. The rest of this chapter addresses expressions that denote some of the most commonly used types of Scheme data, beginning with *primitive data expressions*.

## 2.1 Primitive Data Expressions

A *primitive* datum is one that is atomic, in the sense that it is not composed of smaller parts that a Scheme program can access. Examples of primitive data in Scheme include *numbers*, *booleans* and the *empty list*. A *primitive data expression* is an expression that denotes a primitive datum.

### 2.1.1 Numbers

According to the syntax rules of Scheme, character sequences such as `3`, `-44`, `34.9` and `85/6` are legal Scheme expressions. According to the semantics of Scheme, these expressions respectively denote the numbers *three*,

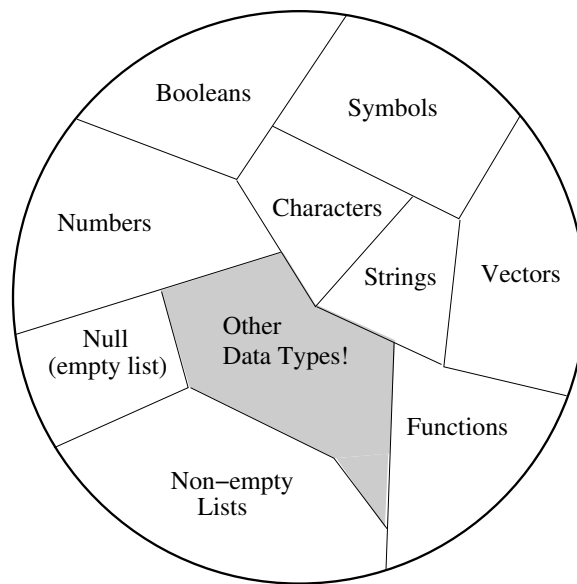


Figure 2.1: The universe of Scheme data, partitioned according to data type

*negative forty-four, thirty-four point nine and eighty-five sixths*. Each of these numbers is an example of a primitive Scheme datum.

For the purposes of this course, it is not necessary to explicitly write down the full set of syntax rules for numerical expressions in Scheme. We will only need the most basic sorts of numerical expressions in Scheme, most of whose rules are undoubtedly already familiar to you through whatever math classes you may have taken in years gone by.

**Character sequences vs. the data they denote.** It is extremely important to distinguish character sequences (e.g., `3`) from the data they denote (e.g., the *number three*). To highlight this distinction, we use the following notation:

Character Sequence  $\longrightarrow$  Datum

For example, we can use this notation to describe the data denoted by the previously seen character sequences:

`3`  $\longrightarrow$  the number *three*  
`-44`  $\longrightarrow$  the number *negative forty-four*  
`85/6`  $\longrightarrow$  the number *eighty-five sixths*

In some cases, multiple Scheme expressions denote the same datum. For example, each of the following character sequences denotes the number *zero* in Scheme: `0`, `000` and `000000`, as indicated below.

`0`  $\longrightarrow$  the number *zero*  
`000`  $\longrightarrow$  the number *zero*  
`000000`  $\longrightarrow$  the number *zero*

As programmers, we only get to type the numerical expressions (i.e., character sequences); however, behind the scenes, the computer is working with the numbers (i.e., Scheme data) denoted by those expressions.

### 2.1.2 Booleans

According to the syntax rules of Scheme, the character sequences, `#t` and `#f`, are legal Scheme expressions. According to the semantics of Scheme, these expressions respectively denote the truth values *true* and *false*, as illustrated below:

#t → the *true* truth value

#f → the *false* truth value

Again, keep in mind the difference between the character sequences and the truth values they denote. The *boolean* data type consists solely of these two truth values (i.e., pieces of data). As programmers, we type the character sequences #t and #f; behind the scenes, the computer is working with the corresponding truth values.

### 2.1.3 The Empty List (or Null)

According to the syntax rules of Scheme, the character sequence, `()`, is a legal Scheme expression. According to the semantics of Scheme, it denotes the *null* datum, which is also called *the empty list*.

`()` → the *empty list*

(We'll encounter non-empty lists later on.) The *null* data type includes only this one datum.

### 2.1.4 The Void Datum

Scheme includes a data type called *void* that contains only one datum, called the *void* datum. As will be seen later on (e.g., in Section 5.5), the *void* datum is used to represent “no value”. For example, a printing function, whose job is to display a bunch of textual information as a harmless *side effect*, will typically return the *void* datum as its *output value*. (Recall that there is a sharp distinction in Scheme between the output value of a function and any side effects it might have.)

- ★ Although the *void* datum is a primitive datum, there is no corresponding primitive data expression that denotes the *void* datum.

In other words, there is no Scheme expression that can be put into the box below to denote the *void* datum.

→ the *void* datum

How, then, can we get our hands on the *void* datum should we ever want to? That's an open question for now.

### 2.1.5 Symbols

Another kind of primitive data in Scheme is a *symbol*. Symbols are typically used as names for things in a Scheme program. For example, each of the *built-in functions* in Scheme has a corresponding symbol that serves as its name. More generally, symbols can be used as names for *any* kind of Scheme data. That is, symbols can be used as *variables* in a Scheme program. For example, the symbol *income* might be used as a variable whose value is some amount of money.

To provide programmers with a great degree of flexibility when dealing with symbols, the syntax rules for symbol expressions in Scheme are very liberal. For example, `miles-per-gallon`, `_LEGAL_SYMBOL_`, `*Legal-Symbol*` and `!even@me?` are all legal expressions in Scheme that denote symbols. Because they are so flexible, it would be a bit tedious to explicitly write down all of the syntax rules for symbol expressions in Scheme. Fortunately, it is not necessary. For our purposes, the following general guidelines will suffice:

- Any sequence of letters, whether lower-case, upper-case, or a mixture of the two, is a legal symbol expression in Scheme. Examples include: `hello`, `goodBye` and `gasMileage`.
- Any character sequence consisting of letters and punctuation characters such as hyphens, asterisks, question marks and exclamation points is a legal symbol expression in Scheme. Examples include: `new-world`, `gas-mileage`, `*CONSTANT*`, `_WIDTH_`, `roll-dice!` and `symbol?`.
- Commonly used one-character expressions, such as `*`, `+`, `-` and `/`, also constitute legal symbol expressions in Scheme.

The semantics of Scheme specifies the datum denoted by each legal symbol expression. For example, the legal expression, `hello`, denotes the symbol *hello*; and the legal expression, `*`, denotes the asterisk symbol.

hello → the symbol *hello*  
 \* → the asterisk symbol

Again, it is important to keep in mind the difference between the typewritten character sequences (e.g., `hello` and `bye-bye`) and the *symbols* (i.e., the Scheme data) that they denote (e.g., the symbol *hello*, and the symbol *bye-bye*). This distinction is hard to write down because we use symbols to denote character sequences, and we also use symbols to denote the symbols denoted by character sequences.) In addition, it is important to remember that symbols are primitive data; they do not have any parts that can be accessed by a Scheme program. For example, the symbol denoted by the expression `hello` does not have any parts; it is indivisible. It may help to think of it as a billiard ball with *hello* written on it.

## 2.2 String Expressions

This section introduces the *string* data type in Scheme. Unlike all of the data types discussed above, strings are non-primitive in Scheme: each string has parts, called characters, that can be accessed by a Scheme program. However, we will not be focusing on the non-primitive nature of strings in this book. In other words, although strings are non-primitive in Scheme, we will, in what follows, treat them as though they were primitive. Why, then, do we introduce them here? Because, as will be seen, Scheme’s very useful printing functions use strings!

- ★ The *string* data type will not be a focus of this book (i.e., it will not play an important role in our understanding of Scheme’s computational model); instead, we will only use strings when we want our Scheme programs to print out useful information.

Syntactically, a string expression is a sequence of characters delimited by double-quotes. For example, `"hi"` and `"Howdy!"` are legal string expressions in Scheme. The semantics of Scheme stipulates that each string expression denotes a string datum (i.e., a non-primitive sequence of characters), as illustrated below.

`"hi"` → the string *"hi"*  
`"Howdy!"` → the string *"Howdy!"*

## 2.3 Summary

This chapter introduced the syntax and semantics for a variety of data types in Scheme, including: numbers, booleans, the empty list, the *void* datum, symbols, and strings. Examples of legal expressions that denote these kinds of data are given below.

Numbers: 342, -81, 34/9, 21.832, etc.  
 Booleans: #t, #f  
 The empty list: ()  
 Symbols: x, miles-per-gallon, dollarsPerGallon, \*, +, /, etc.  
 Strings: "hi", "Howdy!"

For each legal expression (i.e., piece of syntax), the semantics specifies the datum denoted by that expression. This book uses a single arrow (→) to represent *denotation*. For example, the fact that the character sequence `34` denotes the number *thirty-four* is represented by: `34` → the number *thirty-four*.

Although there are expressions that denote data belonging to most of the data types listed above, there is no legal expression in Scheme that denotes the *void* datum.

Of all the data types addressed above, only strings are non-primitive in Scheme; however, investigating the non-primitive aspects of strings shall not be a focus of this book. But have no fear: Chapter 6 will introduce *non-empty lists*, a non-primitive type of data that plays a central role in Scheme’s computational model.



## Chapter 3

# Evaluating Scheme Data

We have seen that a variety of character sequences (e.g., `34`, `xyz`, `()` and `#t`) constitute legal expressions according to the syntax rules of Scheme. In addition, we've seen that each legal expression denotes a piece of data of a particular kind. For example, `34` denotes the number *thirty-four*, and `xyz` denotes the symbol *xyz*. The character sequences are expressions; the data they denote belong to the universe of Scheme data. As programmers, we type character sequences; the computer deals with the corresponding Scheme data.

This chapter addresses the one thing that a Scheme computer does—namely, it *evaluates* Scheme data. The following observations are important to keep in mind:

- ★ Evaluation is done by the computer, not the programmer.
- ★ Evaluation involves Scheme data, not expressions/character sequences.

Because evaluation is the one-and-only thing that a Scheme computer does, it is important to carefully describe it. The good news is that the process of evaluation can be described fairly succinctly for many kinds of Scheme data.

We begin by noting that evaluation is a *function*—in the mathematical sense (i.e., something that takes zero or more inputs, and generates a single output). In particular, the evaluation function takes one Scheme datum as its input, and generates another Scheme datum as its output, as illustrated in Fig. 3.1.

The result of applying the *evaluation* function depends on the type of data that it is applied to. Thus, in what follows, we describe what the evaluation function does for each kind of data we have seen so far.

- ★ In most cases, the application of the evaluation function to a Scheme datum does not directly generate any side effects. However, there are some important exceptions that shall be highlighted as they are encountered—in Chapters 7, ?? and ??.

### 3.1 Evaluating Numbers, Booleans, the Empty List, the *void* Datum, and Strings

The *evaluation* function acts like the *identity function* when applied to numbers, booleans, the *empty list*, the *void* datum, or strings, as illustrated in Fig. 3.2. Since drawing the kinds of black boxes shown in Figs. 3.1 and 3.2 takes up so much space, from now on we'll use a simpler, text-based notation to represent the application of the *evaluation* function to some datum, as illustrated below.

Input Datum     $\implies$     Output Datum

The double arrow ( $\implies$ ) is reserved solely for representing the application of the *evaluation* function to some Scheme datum (called the input) to generate some, possibly quite different Scheme datum (called the output).

- ★ Instead of saying that the *evaluation* function generates the output datum when applied to a certain input datum, we may say that the output datum is the result of *evaluating* the input datum (or that the input datum *evaluates* to the output datum). Keep in mind that when we say such things, we are talking about the application of the one-and-only *evaluation* function.




Figure 3.1: The evaluation function in Scheme




Figure 3.2: Sample evaluations in Scheme

Here are some more examples illustrating the trivial behavior of the *evaluation* function when applied to numbers, booleans, the *empty list*, the *void datum*, or strings:

the number <i>zero</i>	⇒	the number <i>zero</i>
the boolean <i>true</i>	⇒	the boolean <i>true</i>
the <i>empty list</i>	⇒	the <i>empty list</i>
the <i>void datum</i>	⇒	the <i>void datum</i>
the string “ <i>hi there</i> ”	⇒	the string “ <i>hi there</i> ”

Of course, if the *evaluation* function acted like the identity function for *every* kind of input, then it would not be very interesting. (It would just be the identity function.) The following section addresses one of the important cases where the evaluation function does something a little more interesting.

## 3.2 Evaluating Symbols

In Scheme, symbols are frequently used as *variables*. In math, variables frequently have values associated with them. For example, the variable  $x$  may have the value 3. So it is with Scheme. For this reason, the evaluation of symbols is different from the evaluation of numbers, booleans, the *empty list*, the *void datum*, or strings. In particular, symbols typically do not evaluate to themselves; instead, they evaluate to the value associated with them. (Keep reading!)

**Environments in Scheme.** In Scheme, symbols are evaluated with respect to an *environment*. For example, the symbol  $x$  might evaluate to the number *three* in one environment, but to the boolean *false* in another environment. Although the word *environment* may sound mysterious, an environment in Scheme is really nothing more than a table of entries, where each entry pairs a symbol  $s$  with its *value*  $v$ . For example, the sample environment illustrated in Fig. 3.3 pairs the symbol  $num$  with the value *three*, and the symbol  $xyz$  with the value *two*.

Symbol	Value
<i>num</i>	number <i>three</i>
<i>xyz</i>	number <i>two</i>
<i>boolie</i>	boolean <i>true</i>
<i>stringy</i>	string “ <i>hello</i> ”

Figure 3.3: A sample environment

- ★ The *value* of a symbol *s* with respect to some environment is simply whatever datum appears in the entry for the symbol *s* in that environment. (If there is no such entry, then the value for *s* with respect to that environment is undefined.)
- ★ The value associated with a symbol in an environment can be a Scheme datum of *any* type.

For example, the symbol *num* evaluates to the number *three* with respect to the environment shown in Fig. 3.3. Similarly, the symbol *xyz* evaluates to the number *two* in that environment; and the symbol *boolie* evaluates to the boolean *true*.

For another example, if an environment  $\mathcal{E}_0$  contains an entry that pairs the symbol *xyz* with the number *two*, while another environment  $\mathcal{E}_1$  contains an entry that pairs the symbol *xyz* with the boolean *false*, then evaluating *xyz* with respect to the environment  $\mathcal{E}_0$  will yield the number *two*, while evaluating *xyz* with respect to  $\mathcal{E}_1$  will yield the boolean *false*.

Okay, that’s true enough. However, while there can be lots of different environments in Scheme, the focus of attention for the next several chapters shall be on the most important environment in Scheme: the *Global Environment*. The Global Environment is the environment that is used by default.

- ★ When we are talking about evaluating some Scheme datum, unless we explicitly say something to the contrary, we shall assume that we are talking about evaluating that datum with respect to the Global Environment.

It may help to think of an environment as a room that has a table of symbol/value pairs tacked to one of its walls. When a symbol needs to be evaluated in that room/environment, the evaluation function simply fetches the symbol’s value from the relevant entry in that table.

**Evaluating symbols in the Global Environment.** If the Global Environment contains an entry that pairs the symbol *xyz* with the number *two*, then the symbol *xyz* will evaluate to the number *two*:

the symbol *xyz*     $\implies$     the number *two*

Since the value that is paired with a symbol in the Global Environment can be a Scheme datum of any type, it might be that the boolean *true* is the value for the symbol *pq*. Similarly, the empty list might be the value associated with the symbol *my-empty-list*, as illustrated below.

the symbol *pq*                     $\implies$     the boolean *true*  
the symbol *my-empty-list*     $\implies$     the *empty list*

Symbols can even evaluate to other symbols. For example, if the Global Environment contains an entry associating the symbol *bar* with the symbol *foo* (where *bar* corresponds to the value), then the following would hold:

the symbol *foo*     $\implies$     the symbol *bar*

On the other hand, if a symbol does not have a corresponding entry in the Global Environment, then evaluating that symbol with respect to the Global Environment is undefined. A little later on, in Chapter 7, we’ll see how to insert entries into the Global Environment, thereby enabling us to create and use variables of our own.

- ★ An environment is a context within which Scheme data get evaluated. However, an environment is *not* a Scheme datum. Thus, environments in Scheme are not available for direct inspection.

### 3.3 Summary

At the core of the Scheme computational model is the process of *evaluation*. Evaluation is a function that takes a Scheme datum as its input and generates a (possibly different) Scheme datum as its output. For each type of data, the semantics of Scheme specifies how instances of that data type are evaluated (i.e., what output is produced). Numbers, booleans, the empty list, the *void* datum, and strings all evaluate to themselves (i.e., the evaluation function works like the identity function for instances of those data types). However, a symbol is evaluated differently: by looking for a corresponding entry in the relevant environment. The default environment is called the Global Environment.

This book uses the double arrow ( $\implies$ ) to represent the process of evaluation. For example, if the Global Environment contains an entry associating the symbol  $x$  with the number *eighty-six*, this fact can be represented as follows:

$$\text{the symbol } x \implies \text{the number } \textit{eighty-six}$$

It is important to remember that:

- (1) each expression—which is a character sequence—denotes a Scheme datum; and
- (2) each Scheme datum evaluates to a (possibly different) Scheme datum.

For example:

$$x \longrightarrow \text{the symbol } x \implies \text{the number } \textit{eighty-six}$$

## Chapter 4

# Introduction to DrScheme

This chapter introduces the piece of software known as *DrScheme*.<sup>1</sup> This software simulates the operation of a computer that understands the Scheme programming language. It also enables us to interact with that simulated computer. In effect, we use DrScheme as an intermediary between us and that simulated computer. We interact with the simulated computer as follows:


- We type some character sequence into DrScheme’s *Interactions Window* (i.e., the lower window-pane in DrScheme’s window).
- DrScheme takes the datum denoted by that character sequence and feeds it into the evaluation function (i.e., DrScheme evaluates that datum), generating some output datum.
- DrScheme displays some typewritten text in the *Interactions Window* describing the output datum to us.

This process is illustrated in Fig. 4.1, where everything in the shaded box is carried out behind the scenes by DrScheme. Notice that our interaction with DrScheme is through the character sequences (i.e., expressions) we type into the Interactions Window; and the character sequences that DrScheme displays to us in response. We never get to “touch” the Scheme data denoted by our character sequences. (What would it mean to touch a number anyway?) For this reason, it is extremely important that we maintain an accurate mental model of what’s going on in that simulated world. In other words, we need to have an accurate understanding of Scheme’s computational model.

More formally, when we type a sequence of characters,  $C_{in}$ , into the *Interactions Window*, and then hit the *Return* (or *Enter*) key, DrScheme does the following:

- (1) It figures out which Scheme datum,  $D_{in}$ , is denoted by the character sequence  $C_{in}$ ;
- (2) It feeds that Scheme datum as input to the evaluation function, which generates an output datum,  $D_{out}$  (i.e.,  $D_{in}$  evaluates to  $D_{out}$ ).
- (3) Finally, it displays some typewritten text,  $C_{out}$ , in the *Interactions Window* that describes the output datum,  $D_{out}$ .

This process is illustrated below.



<sup>1</sup>The DrScheme software is freely available from [drscheme.org](http://drscheme.org). For the purposes of this book, DrScheme and DrRacket, which is freely available from [drracket.org](http://drracket.org), may be considered to be equivalent. Thus, DrRacket may be used in place of DrScheme, if desired.

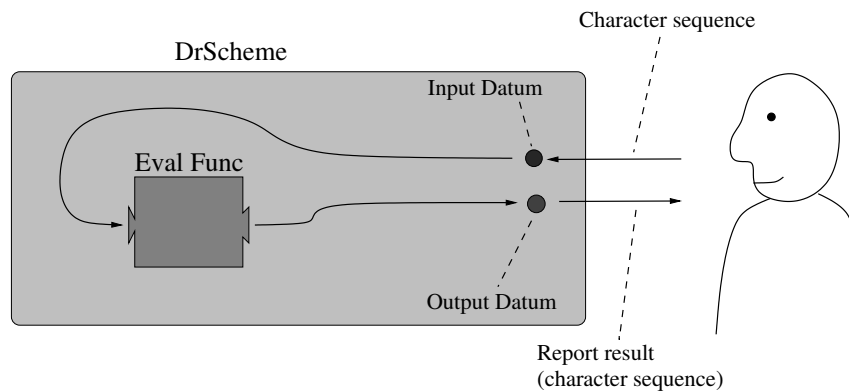


Figure 4.1: A programmer's interaction with DrScheme

Keep in mind that we only see the character sequences,  $C_{in}$  and  $C_{out}$ ; we do not see the Scheme data,  $D_{in}$  and  $D_{out}$ . (What does a Scheme datum look like anyway?) We can more succinctly describe this process as follows:

$$C_{in} \longrightarrow [ D_{in} \Longrightarrow D_{out} ] \longrightarrow C_{out}$$

where the single arrow ( $\longrightarrow$ ) represents the translation from character sequences to the denoted Scheme data (in either direction), the double arrow ( $\Longrightarrow$ ) represents the application of the evaluation function, and the square brackets indicate that we don't get to see the Scheme data,  $D_{in}$  and  $D_{out}$ .

- ★ When entering expressions into the Interactions Window, the datum  $D_{in}$  is evaluated with respect to the Global Environment.

## 4.1 Entering Expressions into the Interactions Window

We can use DrScheme to confirm some of the things discussed in previous chapters. In particular, we can enter character sequences (i.e., expressions) into the Interactions Window and then examine the results reported by DrScheme. In each case, we only get to see the character sequences we type in, and those reported back by DrScheme; we do not get to see the Scheme data that is manipulated by the Scheme computer.

### Example 4.1.1: DrScheme's Interactions Window

*The following interactions demonstrate that numbers, booleans, the empty list and strings all evaluate to themselves:*

```
> 3
3
> #t
#t
> ()
()
> "Howdy!"
"Howdy!"
```

*In the Interactions Window, DrScheme uses the > character to prompt the user for input. Everything following the > character is typed by the programmer. The text on the next line is that generated by DrScheme in response. Thus, the above example shows four separate interactions.*

In these simple examples, the character sequence displayed by DrScheme happens to be the same as that typed by the programmer. However, recall that, behind the scenes, DrScheme is doing quite a bit more than these examples suggest. In particular:

$$\begin{aligned} 3 &\longrightarrow [ \text{ the number } \textit{three} \implies \text{ the number } \textit{three} ] \longrightarrow 3 \\ \#\text{t} &\longrightarrow [ \text{ the boolean } \textit{true} \implies \text{ the boolean } \textit{true} ] \longrightarrow \#\text{t} \\ () &\longrightarrow [ \text{ the empty list } \implies \text{ the empty list } ] \longrightarrow () \\ \text{"Howdy!"} &\longrightarrow [ \text{ the string } \textit{"Howdy!"} \implies \text{ the string } \textit{"Howdy!"} ] \longrightarrow \text{"Howdy!"} \end{aligned}$$

### Example 4.1.2

The following interactions demonstrate that several different character sequences can be used to denote the number zero:

```
> 0
0
> 000
0
> 000000
0
```

As this example illustrates, DrScheme need not use the same character sequence as the one we entered when reporting back that the result of evaluating the number zero is the number zero. Instead, DrScheme chooses the most compact character sequence.

- ★ For convenience, we may say that DrScheme is evaluating the expressions we type into the Interactions Window, when of course we mean that DrScheme is evaluating *the data denoted by* the expressions we type into the Interactions Window.

## 4.2 DrScheme's *Run* Button

Although manually typing individual expressions into the Interactions Window and viewing DrScheme's responses can be quite useful, it is often desirable to ask DrScheme to evaluate a large number of Scheme expressions. (Re-read the above note about "evaluating expressions".) To avoid endless typing and re-typing (e.g., when fixing errors), the upper window-pane of DrScheme, called the *Definitions Window*, can be used to edit—and, if desired, save—any number of Scheme expressions. Afterward, clicking the *Run* button in DrScheme's toolbar causes DrScheme to evaluate each of the expressions currently residing in the Definitions Window, one after the other, as if we had manually typed them into the Interactions Window, as illustrated in Fig. 4.2.

- ★ When using the *Run* button, DrScheme only reports the *results* of evaluating the expressions from the Definitions Window.

More generally, the Definitions Window can be used to hold the contents of an entire Scheme program. In such cases, clicking the *Run* button would cause all of the expressions in that program to be evaluated, one implication being that any functions defined in that program could then be used.

## 4.3 Summary

The DrScheme software simulates a Scheme computer that we, as programmers, can interact with. We type expressions (i.e., character sequences) into the Interactions Window, and DrScheme responds by displaying some (possibly different) character sequence. However, something very important happens in between:

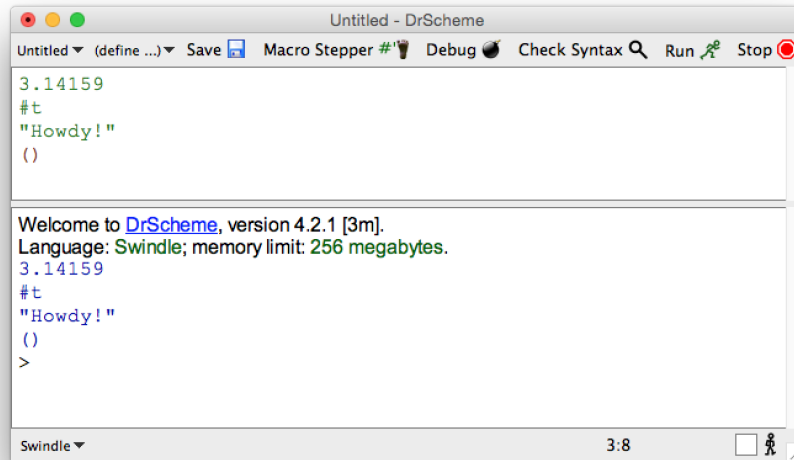


Figure 4.2: Using the *Run* button to evaluate (the data denoted by) multiple expressions

- (1) the input character sequence  $C_{\text{in}}$  denotes some Scheme datum  $D_{\text{in}}$ ;
- (2) DrScheme evaluates  $D_{\text{in}}$ , yielding some datum  $D_{\text{out}}$ ; and
- (3) DrScheme displays a character sequence  $C_{\text{out}}$  that describes  $D_{\text{out}}$  to us.

This process is concisely summarized by:

$$C_{\text{in}} \longrightarrow [ D_{\text{in}} \implies D_{\text{out}} ] \longrightarrow C_{\text{out}}$$

where the stuff between the square brackets is invisible to us. Since such important computations are happening behind the scenes, it is important that we, as programmers, have an accurate mental model of what Scheme is doing.

DrScheme's Definitions Window can be used to hold multiple expressions. Clicking the *Run* button causes each of those expressions to be evaluated in turn, with the results reported in the Interactions Window.



## Chapter 5

# Built-In Functions

For convenience, Scheme includes a variety of *built-in* functions. Examples include the addition function, the multiplication function, and a printing function, among many others.

- ★ Each built-in function is a Scheme datum that is primitive, like numbers and booleans, in the sense that they don't have any parts that a Scheme program can access. Thus, a built-in function is a black box to us.

If you are wondering what character sequences in Scheme denote built-in functions, the answer may surprise you:

- ★ There are no Scheme expressions that denote built-in Scheme functions.<sup>1</sup>

This surprising fact leads to another question: How can a Scheme programmer make use of built-in functions if none of them are denoted by any Scheme expressions? The answer is indicated by the following observation.

- ★ Although the *Input Datum* shown in Fig. 4.1 can never be a function, the *Output Datum* can be.

In particular, for each built-in function, there is an entry in the Global Environment that associates a particular symbol with that function. Therefore, *evaluating* that symbol generates the corresponding function as an output value. In other words, if the *Input Datum* from Fig. 4.1 is a symbol that serves as the name of a built-in function, then the corresponding *Output Datum* will be that function. That is: we gain access to a built-in function by evaluating the symbol that serves as its *name*.

The rest of this chapter introduces some of the most commonly used built-in functions.

### 5.1 Built-in Functions for Arithmetic

DrScheme provides a variety of built-in functions for doing basic arithmetic computations. For example, when DrScheme is first started up, the Global Environment is automatically populated with entries that ensure that each of the following evaluations holds:

the symbol `+`  $\implies$  the *addition* function  
the symbol `-`  $\implies$  the *subtraction* function  
the symbol `*`  $\implies$  the *multiplication* function  
the symbol `/`  $\implies$  the *division* function

Thus, a Scheme programmer can refer to these built-in functions indirectly, by asking DrScheme to evaluate the corresponding symbols.

---

<sup>1</sup>Indeed, there are no Scheme expressions that denote *any* kind of function, whether built-in or not!

**Example 5.1.1: Accessing the built-in arithmetic functions**

That the abovementioned entries do indeed exist in the Global Environment can be confirmed by DrScheme, as illustrated below:<sup>a</sup>

```
> +
#<procedure:+>
> -
#<procedure:->
> *
#<procedure:*>
> /
#<procedure:/>
```

The behind-the-scenes work involved in these interactions can be summarized as follows:

$$\begin{aligned}
 + &\longrightarrow [ \text{the } + \text{ symbol} \implies \text{the addition function} ] \longrightarrow \#<\text{procedure}:+> \\
 - &\longrightarrow [ \text{the } - \text{ symbol} \implies \text{the subtraction function} ] \longrightarrow \#<\text{procedure}:-> \\
 * &\longrightarrow [ \text{the } * \text{ symbol} \implies \text{the multiplication function} ] \longrightarrow \#<\text{procedure}:*> \\
 / &\longrightarrow [ \text{the } / \text{ symbol} \implies \text{the division function} ] \longrightarrow \#<\text{procedure}:/>
 \end{aligned}$$

<sup>a</sup>This text uses the terms, *function* and *procedure*, interchangeably; however, the term *function* seems better suited given that Scheme is typically referred to as a *functional* programming language.

Notice that the character sequences reported by DrScheme need not be legal pieces of Scheme syntax. (Recall that there is no legal piece of Scheme syntax that denotes a primitive function.) Instead, a character sequence such as `#<procedure:+>` is DrScheme's best attempt to describe to us the fact that the output datum is a function—namely, the function associated with the `+` symbol.

- \* Although we are required to type legal Scheme expressions into the Interactions Window, DrScheme is allowed to write whatever it wants when it seeks to describe the results of an evaluation.

## 5.2 Contracts

To be able to make proper use of a built-in function, it is important to know its name, the kinds of inputs it can be applied to, the order in which it expects its inputs, some sort of description of the output it is supposed to generate and, if applicable, any side effects it might have. This kind of information is typically gathered together into a *contract*, as illustrated by the following examples.

**Example 5.2.1: Contracts for some built-in functions**

Here is a contract for the built-in addition function:

<i>Name:</i>	<code>+</code>
<i>Inputs:</i>	$x_1, x_2, \dots, x_n$ ; zero or more numerical inputs
<i>Output:</i>	The sum, $x_1 + x_2 + \dots + x_n$
<i>Side Effects:</i>	None

Notice that the contract describes what the output value should be, but it does not go into the underlying details about how that output value is actually computed. Similar remarks apply to the following contract for the built-in subtraction function:

<i>Name:</i>	–
<i>Inputs:</i>	$x_1, x_2, \dots, x_n$ ; one or more numerical inputs
<i>Output:</i>	If $n = 1$ (i.e., if there is only one input), then the output is $-x_1$ otherwise, the output is the value, $x_1 - x_2 - x_3 \dots - x_n$
<i>Side Effects:</i>	None

For example, applying the subtraction function to the two inputs, 10 and 3, yields the output 7; but applying it to the single input 5 yields the output  $-5$ .

- ★ Since most functions encountered in this course will not have any side effects, we shall follow the convention that if a contract does not mention side effects, then the function can be assumed to not have any.

The rest of this chapter presents several other commonly used built-in functions. The next chapter will show how to apply functions to inputs (i.e., make them do something). Later chapters will show how to create functions of our own design and give them names by inserting appropriate entries into the Global Environment.

### 5.3 Built-in Functions for Integer Arithmetic

You may recall the process of doing integer division in grade school. For example, you may have been shown that 17 divided by 3 yields an answer of 5 with remainder 2. (The answer is often called the *quotient*—but I always had trouble remembering that.) DrScheme provides two built-in functions, called `quotient` and `remainder`, that together can be used to carry out integer division: `quotient` provides the answer; `remainder` provides the remainder. The contracts for these functions are given below:

<i>Name:</i>	<code>quotient</code>
<i>Inputs:</i>	<i>numer</i> , <i>denom</i> , two integers
<i>Output:</i>	The (integer) answer that results from dividing <i>numer</i> by <i>denom</i> , ignoring any remainder.

<i>Name:</i>	<code>remainder</code>
<i>Inputs:</i>	<i>numer</i> , <i>denom</i> , two integers
<i>Output:</i>	The (integer) remainder left over from dividing <i>numer</i> by <i>denom</i> .

Scheme also includes a built-in function, called `integer?`, for checking whether a given datum is an integer.

<i>Name:</i>	<code>integer?</code>
<i>Input:</i>	<i>datum</i> , anything
<i>Output:</i>	<code>#t</code> if <i>datum</i> is an integer; otherwise, <code>#f</code>

### 5.4 The Built-in `eval` Function

The evaluation function that is so important to the computational model of Scheme is itself provided as a built-in function. In particular, the Global Environment contains an entry that associates the `eval` symbol with the built-in evaluation function, as demonstrated by the following interaction:

```
> eval
#<procedure:eval>
```

Since it is a primitive, built-in function, we don't get to see *how* the evaluation function operates; however, we have started to discover *what* the evaluation function does—at least for some kinds of Scheme data. Subsequent chapters will address what the evaluation function does for other kinds of Scheme data. Once we understand what the evaluation function does for each kind of Scheme data, we could think about writing down a contract for it.

- ★ Like numbers, booleans, the empty list, the *void* datum, and strings, functions evaluate to themselves.

In other words, if you feed a Scheme function as input to the evaluation function, the output will be that same function. For example, the addition function evaluates to the addition function; the multiplication function evaluates to the multiplication function; and the evaluation function applied to itself yields itself (!), as summarized below.

$$\begin{aligned} \text{the addition function} &\implies \text{the addition function} \\ \text{the multiplication function} &\implies \text{the multiplication function} \\ \text{the evaluation function} &\implies \text{the evaluation function} \end{aligned}$$

A demonstration of functions evaluating to themselves will be given in the next chapter.

## 5.5 The Built-in Functions `printf` and `void`

Recall from Section 2.1.4 that Scheme includes a data type called *void* whose only datum is also called *void*. The purpose of the *void* datum is to represent “no value”. For example, a function whose main job is to do a bunch of side-effect printing might return the *void* datum as its output value, representing “no output value”. In such cases, DrScheme would display all of the side-effect printing, but would not display anything for the *void* output value. (Since *void* represents “no value”, DrScheme does not feel compelled to display anything for *void*.)

- ★ If a function’s output is *void*, then we may say that the function does not generate any output value.

The built-in functions, `printf` and `void`, introduced below, are examples of functions whose output value is invariably the *void* datum.

### 5.5.1 The `printf` Function

Scheme provides a built-in `printf` function that can be used to display information in the Interactions Window.

- ★ The display of textual information by the `printf` function is an example of a harmless *side effect*.
- ★ The output value generated by the `printf` function is always the *void* datum.

Although the `printf` function has additional capabilities that won’t be explored until Chapter 10, the contract for the simplest use of the `printf` function, given below, will suffice for now.

Name:	<code>printf</code>
Input:	<i>str</i> , a string
Output:	the <i>void</i> datum (i.e., “no value”)
Side Effect:	displays the contents of the string <i>str</i> in the Interactions Window (without the double-quotes)

### 5.5.2 The `void` Function

Recall from Section 5.5 that there is no legal Scheme expression that we can type into the Interactions Window that denotes the *void* datum. However, should you ever need to get your hands on the *void* datum, there is a built-in function, called `void`, that does nothing but generate the *void* datum as its output. Here is its contract:

Name:	<code>void</code>
Inputs:	Any number of inputs
Output:	the <i>void</i> datum
Side Effects:	none

## 5.6 Summary

There are no Scheme expressions that *denote* functions! However, that is not a problem because there *are* Scheme expressions that denote Scheme data that *evaluate* to functions. (Denotation vs. evaluation.) In particular, the Global Environment comes pre-populated with entries that associate certain symbols with various *built-in* functions. For example, the `+` symbol is associated with the built-in *addition* function; and the `*` symbol is associated with the built-in *multiplication* function. As a result, we can effectively refer to the built-in functions by name, as illustrated below:

```
+ → [ the + symbol ⇒ the built-in addition function ] → #<procedure:+>
```

Note that DrScheme is not required to follow the rules of Scheme syntax when displaying information in the Interactions Window.

So that we may use the built-in functions properly, each function has an associated *contract* that specifies its name (a symbol), its inputs (how many and their types), its output, and any side effects it might have. The information found in the contracts for the built-in functions is available online, for example, using the *HelpDesk* feature of the DrScheme program. Later on, when we learn how to specify functions of our own design (cf. Chapter 9), we will include a contract for each new function.

The evaluation function itself is provided as a built-in function—it is the value associated with the `eval` symbol.

### Built-In Functions Introduced in this Chapter

Basic Arithmetic:	<code>+</code> , <code>-</code> , <code>*</code> , <code>/</code>
Integer Arithmetic:	<code>quotient</code> , <code>remainder</code> , <code>integer?</code>
Evaluation Function:	<code>eval</code>
Basic Printing:	<code>printf</code>
Generating the <i>void</i> datum:	<code>void</code>



## Chapter 6

# Non-Empty Lists

Previously, we have seen many examples of primitive data: numbers, booleans, the empty list, the *void* datum, symbols, and primitive functions. Recall that each primitive datum is atomic in the sense that it has no parts that we, as Scheme programmers, can access. In contrast, strings are non-primitive data: they have parts, called characters, that are accessible to Scheme programmers. However, instead of exploring the non-primitive nature of strings, this chapter explores another kind of non-primitive data: *non-empty lists*.

★ As will soon be revealed, non-empty lists play a very important role in Scheme's computational model.

A non-empty list is an ordered sequence of Scheme data. For example, a list might contain items such as the symbol `+`, the number *three*, and the boolean *true*. Other examples of non-empty lists are given below:

- a list containing the number *three* and the number *four*
- a list containing the `+` symbol, the number *three*, and the number *four*
- a list containing: (1) the symbol *eval*, and (2) a subsidiary list containing the `+` symbol, the number *three*, and the number *four*

The last example illustrates that a list can contain elements that are themselves lists.

★ A non-empty list is, by itself, a Scheme datum. It is a Scheme datum that happens to contain other Scheme data as its *elements*.

### 6.1 The Syntax and Semantics for Non-Empty Lists

Since a non-empty list is a Scheme datum, a natural question arises: what kinds of character sequences can the programmer use to denote non-empty lists (i.e., what are the syntax rules for non-empty lists)? We begin with sample character sequences that the programmer can use to denote the Scheme lists described above:

`(3 4)` → a list containing the number *three* and the number *four*

`(+ 3 4)` → a list containing the `+` symbol, the number *three*, and the number *four*

`(eval (+ 3 4))` → a list containing:

(1) the symbol *eval*, and

(2) a subsidiary list containing the `+` symbol, the number *three*, and the number *four*

As these examples illustrate, if  $E_1, E_2, \dots, E_n$  are legal Scheme expressions (i.e., character sequences), then the character sequence

$$(E_1 E_2 \dots E_n)$$

is a legal character sequence. (That's the syntax!) Furthermore, that character sequence denotes a list containing the  $n$  items denoted by  $E_1, E_2, \dots, E_n$ . (That's the semantics!) Thus, if

$$\begin{aligned}
 E_1 &\longrightarrow D_1 \\
 E_2 &\longrightarrow D_2 \\
 &\dots \\
 E_n &\longrightarrow D_n
 \end{aligned}$$

(i.e., each  $E_i$  is a Scheme expression that denotes a Scheme datum,  $D_i$ ), then the character sequence

$$(E_1 E_2 \dots E_n)$$

is a legal character sequence that denotes a list  $\mathcal{D}$  containing the  $n$  items  $D_1, D_2, \dots, D_n$ .

For example, the character sequences  $+$ ,  $3$  and  $4$  are legal Scheme expressions that respectively denote the  $+$  symbol, the number *three*, and the number *four*. Thus, the character sequence,  $(+ 3 4)$ , is a legal Scheme expression that denotes a list containing the  $+$  symbol, the number *three*, and the number *four*. In this example, the expressions  $E_1, E_2$  and  $E_3$  are  $+$ ,  $3$  and  $4$ , respectively; and the Scheme data  $D_1, D_2$  and  $D_3$  are the  $+$  symbol, the number *three*, and the number *four*.

Since  $(+ 3 4)$  denotes a list, if we type this character sequence into the Interactions Window, the *Input Datum* will be that list. (It may help to refer back to Fig. 4.1.) However, DrScheme will then *evaluate* that list—because DrScheme always evaluates the Input Datum to generate the Output Datum. Therefore, we need to talk about how non-empty lists are evaluated.

## 6.2 Evaluating Non-Empty Lists: the Default Rule

As already seen, the empty list evaluates to itself; however, the evaluation of a non-empty list is altogether different. This section presents the *Default Rule for evaluating non-empty lists*. Exceptions to the Default Rule—the so-called *special forms*—will be covered later on.

### Example 6.2.1

*We begin with some examples that confirm that something new is happening when DrScheme evaluates non-empty lists.*

```

> (+ 2 3)
5
> (* 3 4 5)
60
> (+ 2 (* 3 10))
32
> (+ 2 (* 3 (+ 4 8 6)))
56

```

*In each of these examples, the expression entered by the programmer is a legal Scheme expression that denotes a Scheme list. (You should convince yourself of this.) In addition, the evaluation of each list appears to result in an arithmetic computation—in fact, the kind of arithmetic computations you’ve seen in math classes over the years. In each case, the list is being evaluated according to the Default Rule.*

### Example 6.2.2

*Consider the expression  $(+ 2 3)$ , which denotes a list containing three items: the  $+$  symbol, the number two, and the number three. The Default Rule for evaluating such lists has two steps. The first step is to evaluate each item in the list. Now, the  $+$  symbol evaluates to the built-in addition function because the Global Environment is guaranteed to contain an entry associating the  $+$  symbol with the addition function.*




The remaining items in the list are numbers; thus, they trivially evaluate to themselves. The results of the first step are summarized below:

*the + symbol*  $\implies$  *the addition function*

*the number two*  $\implies$  *the number two*

*the number three*  $\implies$  *the number three*

Okay, so after evaluating all of the items in the list, we have the addition function and two numbers. The second step in the Default Rule involves applying that function to the remaining items (i.e., feeding the remaining items as input into that function), as illustrated below:



The resulting output datum is what we take to be the result of evaluating the original non-empty list! Thus, the result of evaluating the list containing the + symbol, the number two, and the number three, is (not surprisingly perhaps) the number five, which DrScheme reports in the Interactions Window using the character sequence 5. Here's a summary of this example:

$(+ \ 2 \ 3) \longrightarrow [ \text{list containing } + \text{ symbol, number two, number three} \implies \text{number five} ] \longrightarrow 5$

where the evaluation is explained by:

*First Step of the Default Rule:*

*+ symbol*  $\implies$  *addition function*

*number two*  $\implies$  *number two*

*number three*  $\implies$  *number three*

*Second Step of the Default Rule:*

*addition function applied to two and three yields output of five*

The evaluation of this list is illustrated in Fig. 6.1.

### Example 6.2.3

Although the Default Rule is not trivial, there are several advantages to it. First, it has only two steps, and they are always the same. Second, it can be used on arbitrarily complex lists without requiring any modifications. For example, recall the interaction:

```
> (+ 2 (* 3 10))
32
```

If we follow the rules we already know, we will see that nothing new is needed to explain this interaction. First, the character sequence  $(+ \ 2 \ (* \ 3 \ 10))$  is a legal Scheme expression that denotes a list. The denoted list contains three items: the + symbol, the number two, and a subsidiary list. The subsidiary list

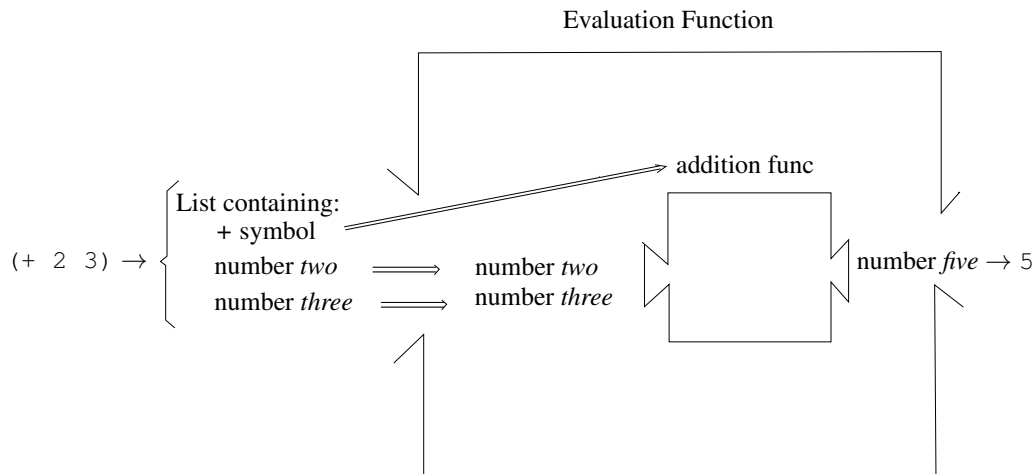


Figure 6.1: The evaluation of the list denoted by (+ 2 3)

contains three items: the \* symbol, the number three, and the number ten. (You should convince yourself of all of this before proceeding.) Okay, so far so good: we have seen that our input expression denotes a particular list. That list, which happens to be a list of lists, shall be the *Input Datum* for the evaluation function.

To evaluate this list, we need to use the *Default Rule*. The first step of the *Default Rule* requires us to evaluate each item in the list:

the + symbol ⇒ the addition function  
 the number two ⇒ the number two  
 the subsidiary list ⇒ **oops!**

Before we can complete the first step of the *Default Rule*, we must evaluate the subsidiary list (i.e., the list containing the \* symbol, the number three, and the number ten). Okay, so we pause for a moment, collect our thoughts, and then proceed.

To evaluate the subsidiary list, we need to use ... the *Default Rule*! The first step of the *Default Rule* requires us to evaluate each item in the list:

the \* symbol ⇒ the multiplication function  
 the number three ⇒ the number three  
 the number ten ⇒ the number ten

The second step of the *Default Rule* requires us to apply the first item (i.e., the function) to the rest of the items. In other words, we need to apply the multiplication function to the numbers three and ten. The result is the number thirty.

Now that we know that the subsidiary list evaluates to thirty, we can pick up from where we left off when evaluating the original list. The first step of the *Default Rule* (for evaluating the original list) requires us to evaluate each item in the list:

the + symbol ⇒ the addition function  
 the number two ⇒ the number two  
 the subsidiary list ⇒ the number thirty

The second step of the *Default Rule* then requires us to apply the first item (i.e., the addition function) to the rest of the items (i.e., the numbers two and thirty). The result is the number thirty-two. And that is the *Output Datum* that results from evaluating the original list! Phew! Of course, *DrScheme* reports this result using the character sequence 32.

### 6.2.1 A More Formal Description of the Default Rule

Consider a list  $\mathcal{L}$  that contains  $n$  data items,  $D_1, D_2, \dots, D_n$ . The evaluation of the list  $\mathcal{L}$  is derived as follows:

- First, evaluate each of the data items,  $D_1, D_2, \dots, D_n$ . The result will be  $n$  (possibly different) data items,  $K_1, K_2, \dots, K_n$ :
 
$$D_1 \implies K_1$$

$$D_2 \implies K_2$$

$$\dots$$

$$D_n \implies K_n$$
- Now, for the Default Rule to work,  $K_1$  *must* be a function. (If  $K_1$  is some other kind of datum, then DrScheme will report an error.)
- The second step is to apply the function  $K_1$  to the rest of the items,  $K_2, \dots, K_n$ . In other words, the items  $K_2, \dots, K_n$  are fed as input to the function  $K_1$ . (If the function  $K_1$  cannot accept that number of inputs, or if those items have the wrong data type, then DrScheme will report an error.) The resulting output will be some datum,  $P$ .
- The evaluation of the list  $\mathcal{L}$  is defined to be that datum  $P$  (i.e.,  $\mathcal{L} \implies P$ ).

As indicated by the parenthetical comments, it is possible for some things to go wrong in the process of evaluating a non-empty list. For example, the function  $K_1$  might expect a different number of inputs than are present in the rest of the original list. Or the attempt to evaluate one of the data  $D_i$  might be undefined. Or the application of the function  $K_1$  to the inputs  $K_2, \dots, K_n$  might be undefined because, for example, the function expects numbers and it gets something else. In any of these cases, the result is undefined and DrScheme would report an error. Thus, none of the following lists can be evaluated:

- a list containing the numbers *one*, *two* and *three*
- a list containing two instances of the *empty list*
- a list containing the + symbol, followed by the boolean *true* and the boolean *false*

It is important to understand that each of the above lists is a valid Scheme datum: each one is a *list*. It's just that these lists cannot be evaluated.

**Example 6.2.4**

*Here's an example of the default case of evaluating a non-empty list where things work out. Let  $\mathcal{L}$  be the list containing the following data:*

$D_1$ : the + symbol,     $D_2$ : the number one,     $D_3$ : the number two,     $D_4$ : the number three

*These Scheme data evaluate to the following:*

$K_1$ : the addition function,     $K_2$ : the number one,     $K_3$ : the number two,     $K_4$ : the number three

*Since the first of these,  $K_1$ , is in fact a function, it can be applied to the inputs  $K_2, K_3$  and  $K_4$  (i.e., the numbers one, two and three). This results in the output six, which is itself a Scheme datum. The number six is the result of evaluating the original list  $\mathcal{L}$ , as illustrated below.*

```
> (+ 1 2 3)
6
```

*Notice that because the addition function is a primitive function, its operation is invisible to us. We observe the inputs going in and the output coming out, but we do not get to see how the output is generated.*

The Default Rule for evaluating non-empty lists is how function application is made available to the Scheme programmer. In particular, if you want to apply a given function to a bunch of inputs, you create an expression that denotes the appropriate list and feed it to DrScheme.

The Default Rule has two steps. The first step involves evaluating each item in the original list, resulting in a bunch of new items. The second step involves applying the first new item—which must be a function—to the rest of the new items—which are the inputs to that function. The output value obtained by applying that function to those inputs is taken to be the output of evaluating the original list.

Scheme is called a *functional* programming language because function application is the central part of the computational model of Scheme. And the Default Rule is how the programmer gets function application to happen.

At this point, you should be able to write arbitrarily complex expressions that, when fed to DrScheme, cause correspondingly complex arithmetic computations to happen. That's pretty good. However, we'll have much more fun when we can design our own functions to do whatever we want them to do. For that, we'll need the `define` and `lambda` special forms, which shall be described in the next chapter.

#### Example 6.2.5

*The fact that 17 divided by 3 yields an answer (i.e., quotient) of 5 with a remainder of 2 can be confirmed by applying the built-in `quotient` and `remainder` functions:*

```
> (quotient 17 3)
5
> (remainder 17 3)
2
```

#### Example 6.2.6

*According to the contract for (the simplest use of) the built-in `printf` function (cf. Section 5.5.1), if the `printf` function is applied to a single input that is a string, then it will display the contents of that string—without the double-quotes—in the Interactions Window as side-effect printing, but the output value will be the void datum, as illustrated below.*

```
> (printf "hi there")
hi there
> (printf "this is a long string!")
this is a long string!
```

*Note that the textual information displayed by DrScheme in each case is side-effect printing, not a Scheme output value. More interesting uses of the `printf` function will be described in Chapter 10.*

★ By default, DrScheme clearly distinguishes side-effect printing from Scheme output values by displaying side-effect printing in one color, and output values in another, as illustrated in Fig. 6.2.

#### Example 6.2.7

*According to the contract for the built-in `void` function (cf. Section 5.5.2), the `void` function can be applied to any number of inputs, but invariably returns the void datum as its output, as illustrated below.*

```
> (void)
>
```

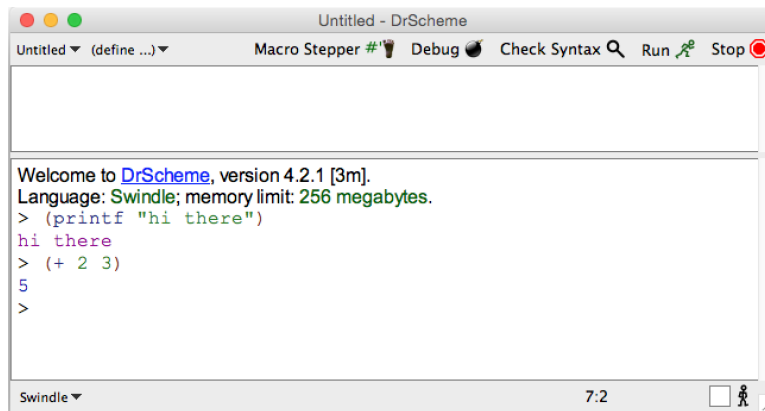


Figure 6.2: DrScheme’s use of different colors to distinguish side-effect printing from output values

```
> (void 3 #t () "xyz")
>
```

*Note that DrScheme does not display anything (other than the prompt) in the Interactions Window when the output value is the void datum.*

### Example 6.2.8

*We can use the Default Rule to explicitly apply the evaluation function to some inputs, as demonstrated below:*

```
> (eval +)
#<procedure:+>
```

*In this example, the list contains two items: the eval symbol and the + symbol. To evaluate this list using the Default Rule, we first evaluate each item in the list:*

*eval symbol  $\implies$  the evaluation function*

*+ symbol  $\implies$  the addition function*

*The second step of the Default Rule requires us to apply the first item (i.e., the evaluation function) to the second item (i.e., the addition function). Since Scheme functions always evaluate to themselves, the result is simply the addition function. DrScheme reports this result to us, in effect, the function associated with the + symbol.*

## 6.3 Summary

The evaluation of non-empty lists plays a critical role in Scheme’s computational model. By default, non-empty lists are evaluated using the Default Rule. The Default Rule has two steps:

- (1) evaluate each element of the non-empty list; and
- (2) apply the result of evaluating the first element to the results of evaluating all of the rest of the elements.

The result from Step Two is taken to be the result of evaluating the original non-empty list.

The Default Rule enables a Scheme programmer to apply a function to any desired inputs: just ask DrScheme to evaluate a list whose first element evaluates to the desired function, and the rest of whose elements evaluate to the desired inputs, as illustrated below:

```
> (+ 3 (* 4 10))  
43
```

As this example demonstrates, the evaluation of a list containing other lists is handled quite naturally: during the first step, when each element of the list must be evaluated, any subsidiary lists are evaluated by ... the Default Rule!

Later on, when you create functions of your own (cf. Chapter 9) you will give each new function a name (cf. Chapter 7). By doing so, you will then be able to apply your new function to whatever inputs you wish, courtesy of the Default Rule.

The evaluation of non-empty lists is only defined when the first element of the list evaluates to a function; and the rest of the elements evaluate to appropriate inputs for that function. Asking DrScheme to evaluate non-empty lists that do not meet these criteria typically results in an error. (The *special forms* introduced in Chapter 7 are exceptions to this.)

## Chapter 7

# Special Forms

In DrScheme, there is a special class of symbol expressions called *keywords*. Examples of keywords include: `and`, `cond`, `define`, `dotimes`, `if`, `lambda`, `let`, `or` and `quote`. Each of these keywords is a legal Scheme expression that denotes a symbol. For example, `quote` denotes the *quote* symbol, and `lambda` denotes the *lambda* symbol. For expository convenience, we may refer to expressions such as `quote` and `lambda` as keyword expressions, and the corresponding symbols (i.e., the *quote* symbol and the *lambda* symbol) as keyword symbols. However, that is not the interesting thing about keywords. The interesting thing about keywords is this:

- ★ When the first element of a non-empty list is a keyword symbol, then that list is a *special form*; and each kind of special form has its own special mode of evaluation.

For example, each of the following expressions denotes a list that is a special form:

```
(define x 3)
(quote (3 4 5))
(if condition then-clause else-clause)
(let ((x 4)) (+ x 8))
```

The important thing about special forms is that they are *not* evaluated according to the Default Rule introduced in Chapter 6. Instead, a special form is evaluated according to a special rule that is specific to the type of that special form—which is determined by the keyword symbol. Thus, there is one rule for evaluating `define` special forms, another rule for evaluating `quote` special forms, and so on. Importantly, each `define` special form is evaluated in the same way, just as each `quote` special form is evaluated in the same way. However, the rule for evaluating `define` special forms is very different from the rule for evaluating `quote` special forms.

Over the next several chapters, you will be introduced to about a dozen different kinds of special form. For each kind of special form, you will learn both the syntax and the semantics. The syntax of special forms is always in terms of a list whose first element is a keyword symbol; the rest of the list can be simple or complex, depending on the kind of special form. The semantics of a special form has two parts: (1) the list that is denoted by the special form expression, and (2) the special mode of evaluation for that kind of special form. As time goes on, you will use these special forms so often that their special modes of evaluation will become second nature to you. And, once you get the hang of it, learning the syntax and semantics for each new kind of special form will get easier and easier.

**Note.** In the Default Rule for evaluating non-empty lists, the first thing that happens is that each element of the list is evaluated, one after the other. In contrast, when evaluating a special form, which is also a non-empty list, some of the elements of that list may *not* be evaluated. Indeed, the *first* element of a special form (i.e., the keyword symbol) is *never* evaluated. (If DrScheme attempted to evaluate a keyword symbol, it would cause an error because the Global Environment typically does not contain entries corresponding to keyword symbols.)

The next sections introduce the `define` and `quote` special forms that you will use every day for the rest of your Scheme-programming life!

## 7.1 The `define` Special Form

The `define` special form is signaled by the `define` keyword. Its purpose is to insert a new symbol/value pair into the Global Environment.<sup>1</sup>

### 7.1.1 The Syntax of the `define` Special Form

A `define` special form expression is any character sequence of the form

$$(\text{define } C_1 C_2)$$

where  $C_1$  is an expression denoting some Scheme symbol  $s$ , and  $C_2$  can be any expression denoting any Scheme datum,  $e$ , as illustrated below.

$$C_1 \longrightarrow s \quad \text{and} \quad C_2 \longrightarrow e$$

Therefore:

$$(\text{define } C_1 C_2) \longrightarrow \text{List containing the } \text{define} \text{ symbol, the } s \text{ symbol, and the datum } e$$

For example, `(define x (+ 3 4))` is a `define` special form expression that denotes a list containing:

- (1) the `define` keyword symbol,
- (2) the symbol `x`, and
- (3) the list denoted by `(+ 3 4)`.

Some more examples of `define` special form expressions are given below.

```
(define addn-func +)
(define zero 0)
(define empty-list ())
```

### 7.1.2 The Semantics of the `define` Special Form

Each special form denotes a list; the `define` special form is no exception. More interesting is what happens when a `define` special form is *evaluated*. The special rule for evaluating `define` special forms is illustrated below:

$$(\text{define } C_1 C_2) \longrightarrow [\text{ List containing } \text{define}, s \text{ and } e \implies \text{ } ] \longrightarrow \text{ }$$

where the gray boxes are used to highlight the following facts:

- ★ The evaluation of a `define` special form does *not* generate any output value. (Well, technically, it generates the *void* datum as its output. Recall from Section 2.1.4 that the *void* datum is used to represent “no value”.)

Instead:

- ★ The purpose of the `define` special form is not to compute an output value, but to generate a very important *side effect*—namely, to insert a new entry into the Global Environment.

DrScheme evaluates a `define` special form by taking the following steps, in order:

- (1) Insert a new entry, 

$s$	<i>void</i>
-----	-------------

, into the Global Environment, where *void* is a temporary placeholder representing that there is not yet any value associated with the symbol  $s$ .
- (2) Evaluate the datum  $e$ , yielding some (usually different) datum  $E$ :  $e \implies E$ .
- (3) Insert  $E$  as the value for  $s$  in the Global Environment: 

$s$	$E$
-----	-----

.



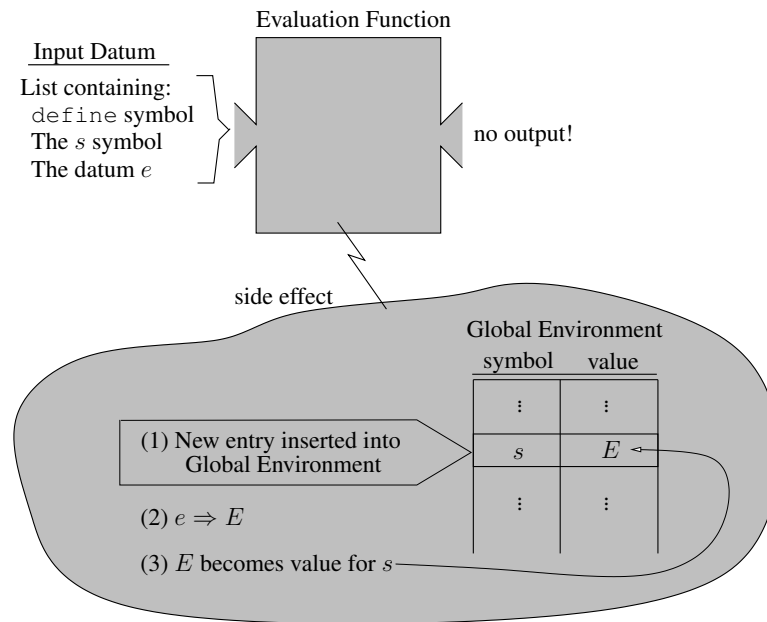


Figure 7.1: The side effect of `define`: inserting a new entry into the Global Environment

This process, except for the part about the use of *void* as a temporary placeholder, is illustrated in Fig. 7.1.

The purpose of evaluating a `define` special form is its side effect: to create a new entry in the Global Environment. Since it does *not* generate any output value—or, rather, since it generates the *void* datum as its output—DrScheme does not display anything in the Interactions Window in response to `define` special forms, as illustrated below:

```
> (define x 6)
> (define y 3)
> (define z 34)
>
```

Of course, something *has* happened!

#### Example 7.1.1

Typing the character sequence, `(define x (+ 1 2 3))`, into the Interactions Window and hitting the Enter key would result in the number six being associated with the symbol `x` in the Global Environment, as illustrated below.

`x`  $\longrightarrow$  the symbol `x`

`(+ 1 2 3)`  $\longrightarrow$  [a list containing the `+` symbol and the numbers one, two and three]  $\implies$  the number six

Side Effect: New Global Environment Entry: 

the symbol <code>x</code>	the number six
---------------------------	----------------

As noted above, DrScheme does not report any output value when evaluating a `define` special form. However, after evaluating it, subsequent attempts to evaluate the symbol `x` result in the value 6, as illustrated below:

<sup>1</sup>The `define` special form can also be used to insert entries into a local environment, but we shall not explore this capability.

```

> x
BUG! reference to undefined identifier: x
> (define x (+ 1 2 3))
> x
6
> (* x 100)
600
> x
6
> (* x 1000)
6000
> (* x x)
36
> x
6

```

Notice that the first attempt to evaluate the symbol `x` resulted in an error; however, after the `define` special form has been evaluated, each time the symbol `x` needs to be evaluated, the result is the value six. The subsequent expressions can be evaluated using what we have learned in previous chapters. We need the Default Rule and we need to know how to evaluate symbols. No new rules are needed. Part of the beauty of Scheme's computational model is that once it is learned, it can be used in an unbelievably wide variety of circumstances.

#### Example 7.1.2: Confirming the semantics of `define`

The following admittedly unusual interactions confirm the semantics of the `define` special form.

```

> w
BUG! reference to undefined identifier: w
> (define w w)
> w
>

```

Prior to evaluating the `define` special form, attempting to evaluate the symbol `w` results in an error, because there is no entry (yet) for `w` in the Global Environment. However, evaluating the `define` special form inserts an entry for `w` into the Global Environment. In particular, as described earlier, the following three steps are taken by DrScheme in evaluating the expression `(define w w)`:

- (1) A new entry, 

<code>w</code>	<code>void</code>
----------------	-------------------

, is inserted into the Global Environment.
- (2) The expression `w` is evaluated, yielding the value `void`:  $w \implies \text{void}$ .  
(That's what's currently stored in the Global Environment as the value for `w`!)
- (3) That value (i.e., `void`) is inserted as the value for `w` in the Global Environment.

Of course, in this case, the third step is redundant, since `void` is already there as the value for `w`.

Afterward, when we ask DrScheme to evaluate `w`, it does so, coming up with the answer `void`. However, since `void` is used to represent “no value”, DrScheme does not display anything! Instead, it just skips to the prompt, awaiting further instructions.

**Note.** Since a keyword is a symbol, like any other Scheme symbol, you could use the `define` special form to assign some value to it in the Global Environment. However, this is a bad idea precisely because it would cause that symbol to lose its status as a keyword. Thereafter, you would not be able to use special forms relying on that

keyword. This is something you might want to do once, just for fun. Afterward, you'll want to hit DrScheme's *Run* button to reset the Global Environment (i.e., to erase what you've done and thereby restore that symbol's status as a keyword).

## 7.2 The `quote` Special Form

Recall that whenever we enter an expression into the Interactions Window, DrScheme invariably *evaluates* the corresponding Input Datum to generate an Output Datum. (You may wish to refer back to Fig. 4.1.) However, sometimes we are interested in data that *cannot* be evaluated (e.g., a list containing a bunch of Social Security numbers). Since attempting to evaluate such data would cause an error, and since DrScheme *always* performs an evaluation, we need some way of *shielding* data from DrScheme's evaluation. That is the purpose of the `quote` special form.

### 7.2.1 The Syntax of the `quote` Special Form

The `quote` special form is indicated by the `quote` keyword. As a character sequence, it has the form

```
(quote C)
```

where  $C$  can be any legal Scheme expression. Below are listed several examples:

```
(quote x)
```

```
(quote (1 2 3))
```

```
(quote (hi there + #t ()))
```

```
(quote (1 (2 (3))))
```

### 7.2.2 The Semantics of the `quote` Special Form

Each `quote` special form denotes a list. In particular, an expression of the form,  $(\text{quote } C)$ , denotes a list containing two items: the *quote* symbol and whatever  $C$  denotes. For example, the expression  $(\text{quote } x)$  denotes a list containing the *quote* symbol and the symbol  $x$ . Similarly,  $(\text{quote } (1\ 2\ 3))$  denotes a list containing the *quote* symbol and a subsidiary list of numbers. More formally, if  $C$  denotes some datum,  $D$ , then  $(\text{quote } C)$  denotes a list containing the *quote* symbol and  $D$ . Using the arrow notation, we can say:

If:  $C \rightarrow D$

Then:  $(\text{quote } C) \rightarrow$  a list containing the *quote* symbol and  $D$

**Evaluating `quote` special forms.** The evaluation of a `quote` special form does not use the Default Rule for evaluating non-empty lists. Instead, `quote` special forms are evaluated using the following special rule:

- ★ A list containing the *quote* symbol and  $D$  evaluates to ...  $D$ .

Notice that, according to this rule, neither the *quote* symbol nor the datum  $D$  are evaluated.<sup>2</sup> Instead,  $D$  is the *result* of evaluating the two-element list. Indeed, the whole point of the `quote` special form is to *shield*  $D$  from evaluation.

#### Example 7.2.1

*Each of the following is an example of a `quote` special form:*

<sup>2</sup>In fact, the keyword symbol is *never* evaluated in a special form of any kind. The purpose of the keyword symbol is simply to indicate that the given list is a special form, thereby requiring a special mode of evaluation.

```

> (quote x)
x
> (quote (1 2 3))
(1 2 3)
> (quote (+ 2 3))
(+ 2 3)

```

In the first example, `(quote x)` denotes a list containing the quote symbol and the symbol `x`. That list is the Input Datum. The result of evaluating that list is the symbol `x`—that is the Output Datum. Notice that the list is evaluated, but its second element is not. We can abbreviate this evaluation as follows:

$$(\text{quote } x) \longrightarrow [ \{\text{list with symbols quote and } x\} \implies \text{the symbol } x ] \longrightarrow x$$

This is quite different from the Default Rule for evaluating non-empty lists. Well, that's to be expected: the Default Rule was not used!

In the second example, `(quote (1 2 3))` denotes a list containing the quote symbol and a subsidiary three-element list. The result of evaluating that list is its second element (i.e., the subsidiary three-element list). Notice that the list containing the numbers one, two and three has not been evaluated. Indeed, any attempt to evaluate such a list would cause DrScheme to report an error since the first element of that list does not evaluate to a function. This example illustrates the use of a list as a container for data rather than something we'd like to have evaluated. The `quote` special form comes in handy for such cases.

In general, if  $C$  is an expression denoting some datum  $D$ , then entering the expression, `(quote C)`, into DrScheme will cause the following to happen:

$$(\text{quote } C) \longrightarrow [ \{\text{list containing quote symbol and } D\} \implies D ] \longrightarrow C'$$

Notice that the Input Datum is the two-element list that contains the *quote* symbol and the datum  $D$ . The Output Datum is simply  $D$ . Notice, too, that DrScheme may use a different character sequence,  $C'$ , to describe  $D$  to us; however,  $C'$  must nonetheless denote  $D$ . (An example of this will be given shortly.)

### Example 7.2.2

Notice the difference between the evaluations of `x` and `(quote x)` below:

```

> (define x (+ 1 2 3))
> x
6
> (quote x)
x

```

### Example 7.2.3

Here, we use the `define` special form to create a variable named `my-list` whose value is a three-element list. Notice the use of the `quote` special form to shield the three-element list from evaluation.

```

> (define my-list (quote (1 2 3)))
> my-list
(1 2 3)

```

### 7.2.3 Alternate Syntax for `quote` Special Forms

Since `quote` special forms are used so frequently, there is an alternate syntax for them. In particular, if  $C$  is any Scheme expression denoting some datum  $D$ , then the expressions, `(quote C)` and `'C`, denote the *same* two-element list—namely, a list containing the *quote* symbol and the datum  $D$ :

$$\begin{aligned} (\text{quote } C) &\longrightarrow \text{list containing } \textit{quote} \text{ symbol and } D \\ 'C &\longrightarrow \text{list containing } \textit{quote} \text{ symbol and } D \end{aligned}$$

The two character expressions are quite different, but both represent the same list! (Syntax vs. Semantics!)

#### Example 7.2.4

The expressions, `'num` and `(quote num)`, each represent a list containing the *quote* symbol and the *num* symbol, as illustrated below:

```
> (quote num)
num
> 'num
num
```

Although the abbreviation for `quote` special forms is useful, it requires care to remember that such expressions denote *lists*—and that those lists are evaluated using the special rule for the `quote` special form.

#### Example 7.2.5

The following examples demonstrate the equivalence between the two kinds of syntax for the `quote` special form.

```
> (quote (quote x))
'x
> ''x
'x
> (quote 000)
0
> '000
0
```

In the first example, DrScheme has chosen a different character sequence for describing the Output Datum—in this case, a list containing the *quote* symbol and the *x* symbol. Similar remarks apply to the third and fourth examples, where the number zero has been shielded from evaluation, but DrScheme has chosen to report the result using a more compact character sequence.

## 7.3 Summary

This chapter introduced *special forms*. A special form is a non-empty list whose first element is one of Scheme's special *keyword* symbols (e.g., `define` or `quote`). The keyword symbol determines the kind of special form (e.g., a `define` special form or a `quote` special form). Although they are non-empty lists, special forms are *not* evaluated by the Default Rule; instead, each kind of special form is evaluated by its own special rule: one rule for `define` special forms, one rule for `quote` special forms, and so on. The rules for evaluating special forms are very different from the Default Rule. For example, the first element of a special form is *never* evaluated. And,

frequently, some or all of the other elements are not evaluated either. This chapter focused on the `define` and `quote` special forms.

- ★ The `define` special form generates no output value, but has a very useful side effect: it inserts a new entry into the Global Environment.
- ★ The `quote` special form is used to shield a datum from evaluation; it has no side effects.

The `define` special form enables us to use symbols as variables (i.e., names for pieces of data). Later on, when you create functions of your own design, you will typically use the `define` special form to give them names. In turn, this will enable you to apply your new functions to any desired inputs simply by asking DrScheme (and the Default Rule) to evaluate an appropriate non-empty list.

The `quote` special form is useful when treating symbols or non-empty lists as pieces of data, rather than using them as names of variables or vehicles for applying functions to inputs. For example, the Default Rule would have problems evaluating a list containing a bunch of student names, but the `quote` special form could be used to shield that list from evaluation, as illustrated below:

```
> (quote (john paul george ringo))
(john paul george ringo)
> '(john paul george ringo)
(john paul george ringo)
```

## Special Forms Introduced in this Chapter

<code>define</code>	For inserting a new entry in the Global Environment
<code>quote</code>	For shielding a Scheme datum from evaluation

# Chapter 8

## Predicates

A function whose output is always a boolean (i.e., *true* or *false*) is called a *predicate*. (This is just convenient terminology; there is no *predicate* type in Scheme.) This chapter describes some of the commonly used, built-in Scheme predicates and illustrates their use.

### 8.1 Type-Checker Predicates

Scheme includes a bunch of primitive data types, including: *number*, *boolean*, *symbol*, *null* and *function*. Scheme also includes non-primitive data types, including *strings* and *non-empty lists*. For each one of these data types, Scheme includes a primitive function called a *type-checker predicate*. When a type-checker predicate is applied to some Scheme datum, it outputs *true* if that datum belongs to the indicated data type; otherwise, it outputs *false*. Thus, the type-checker predicate associated with the *number* data type outputs *true* whenever the input belongs to the *number* data type. Similarly, the type-checker predicate associated with the *list* data type outputs *true* whenever the input datum belongs to the *list* data type.<sup>1</sup> And so on.

For convenience, each of these type-checker predicates has an easy-to-remember *name*. In other words, for each type-checker predicate there is an entry in the Global Environment that links a particular symbol with that predicate. Thus, those symbols can be used to refer to the type-checker predicates. For example, the symbol `number?` evaluates to the type-checker predicate for the *number* data type; the symbol `boolean?` evaluates to the type-checker predicate for the *boolean* data type; and so on.

#### Example 8.1.1

*The following Interactions Window session demonstrates the existence of some of the built-in type-checker predicates.*

```
> number?
#<procedure:number?>
> symbol?
#<procedure:symbol?>
> boolean?
#<procedure:boolean?>
> list?
#<procedure:list?>
> null?
#<procedure:null?>
> procedure?
#<procedure:procedure?>
```

<sup>1</sup>The *list* data type is a *compound data type* that includes both non-empty lists and the empty list.

```

> void?
#<procedure:void?>
> string?
#<procedure:string?>

```

Notice that the symbols mirror the names of the corresponding data types, except that the symbol associated with the type-checker predicate for functions is `procedure?`, not `function?`.<sup>a</sup>

<sup>a</sup>This text uses the terms, *function* and *procedure*, interchangeably; however, the term *function* seems better suited given that Scheme is typically referred to as a *functional* programming language.

Each type-checker predicate is a function that can be applied to a single input. That input can be any type of Scheme datum. A type-checker predicate returns *true* if that input datum is of the appropriate data type.

### Example 8.1.2

Here's a contract for the built-in `number?` type-checker predicate:

<i>Name:</i>	<code>number?</code>
<i>Input:</i>	<code>d</code> , any Scheme datum
<i>Output:</i>	<code>#t</code> if <code>d</code> is a number; otherwise, <code>#f</code>

The contracts for the other type-checker predicates are similar.

### Example 8.1.3

The following interactions illustrate the behavior of the type-checker predicates.

```

> (number? 3)
#t
> (number? #t)
#f
> (boolean? #f)
#t
> (boolean? 'x)
#f
> (symbol? +)
#f
> (symbol? '+)
#t
> (null? ())
#t
> (null? '(+ 1 2))
#f
> (procedure? +)
#t
> (procedure? '+)
#f
> (list? '(+ 1 2))
#t
> (list? ())
#t

```



```

> (list? +)
#f
> (void? (void))
#t
> (void? void)
#f
> (string? "abc")
#t
> (string? '("a" "b" "c"))
#f
> (string? #t)
#f

```

Each of these expressions denotes a non-empty list that is evaluated according to the Default Rule. In each case, the first element of the list is a symbol that evaluates to a function, which is then applied to whatever the second element evaluates to. Notice that the `+` symbol in `(procedure? +)` evaluates to the addition function, whereas the `'` + expression in `(procedure? '+)` evaluates to the `+` symbol. Notice too that the `list?` type-checker predicate returns true for any list, whether empty or non-empty. Finally, recall that `void` is a built-in function whose output is the void datum. Thus, `(void)` evaluates to the void datum, whereas the symbol `void` evaluates to the built-in function.

## 8.2 Comparison Predicates

In addition to the primitive arithmetic functions for addition, subtraction, multiplication and division, Scheme includes several predicates for comparing numbers. Examples include the *greater-than*, *less-than* and *equal* predicates.<sup>2</sup> To enable us to refer to such predicates, each is associated with a particular symbol in the Global Environment.

```

>    greater than
>=   greater than or equal to
=    equal to
<    less than
<=   less than or equal to

```

Each of these predicates, when applied to two numeric inputs, generates the expected boolean output, as illustrated below.<sup>3</sup>

### Example 8.2.1

```

> (> 3 4)
#f
> (> 4 3)
#t
> (>= 4 3)
#t
> (= 3 4)
#f
> (= 3 3)
#t

```

<sup>2</sup>In other contexts, these predicates are commonly called *relational operators*.

<sup>3</sup>These predicates can also be applied to more than two inputs; however, we shall postpone discussion of such things until Chapter ??.

DrScheme also provides a comparison predicate called `eq?` that is more general than the `=` predicate. Whereas the `=` predicate only works on numerical input, the `eq?` predicate can be used to test the equality of inputs that can be any combination of numbers, booleans, symbols or the empty list. Here's a contract for the `eq?` predicate.

Name:	<code>eq?</code>
Inputs:	$d_1$ , a number, boolean, symbol, or the empty list $d_2$ , a number, boolean, symbol, or the empty list
Output:	<code>#t</code> if $d_1$ and $d_2$ are the same; <code>#f</code> otherwise.

### Example 8.2.2

Here are some examples of the `eq?` predicate in action.

```
> (eq? 3 3)
#t
> (eq? 3 'x)
#f
> (eq? 'x 'x)
#t
> (eq? 'x #t)
#f
> (eq? 'x ())
#f
> (eq? () ())
#t
```

The `eq?` predicate is most frequently used to compare whether two symbols are the same. If you know that the inputs will be numbers, then you should use the `=` function. And if you know that the inputs will be booleans ... stay tuned!

★ The `eq?` function does not work well when comparing non-empty lists! More on that later!

## 8.3 Summary

This chapter introduced *predicates*—that is, functions that generate boolean output values. DrScheme provides a wide variety of built-in predicates. Each built-in predicate has a corresponding entry in the Global Environment so that it can be used by a Scheme programmer. For example, the built-in *less-than* predicate is the value associated with the `<` symbol in the Global Environment. By taking advantage of the Default Rule for evaluating non-empty lists, the *less-than* function can be applied to inputs, as demonstrated below:

```
> (< 3 4)
#t
> (< (+ 2 3) (- 10 9))
#f
```

This chapter introduced two sets of built-in predicates: *type-checker predicates* and *comparison predicates*. Type-checker predicates simply check whether a given datum belongs to a specified data type. For example, the `number?` predicate checks whether its input is a number, and the `list?` predicate checks whether its input is a list, as demonstrated below:

```
> (number? 3)
#t
> (number? '(a b c))
#f
> (list? '(a b c))
#t
```

The `list?` predicate works for any kind of list: empty or non-empty. The `null?` predicate works only for the empty list. The `procedure?` predicate works for functions. The comparison predicates include the standard functions for comparing numbers (e.g., *less-than* and *greater-than-or-equal-to*), as well as the more general `eq?` predicate that works on any combination of numbers, booleans, symbols, or the empty list.

### Built-in Functions Introduced in this Chapter

Type-checker Predicates: `number?`, `symbol?`, `boolean?`, `list?`,  
`null?`, `procedure?`, `void?`, `string?`.

Comparison Predicates: `<`, `<=`, `=`, `>=`, `>` (these work only on numbers).  
`eq?` (this works on numbers, booleans, symbols or the empty list).



## Chapter 9

# Defining Functions

So far, what we know about Scheme is enough to enable us to use the Interactions Window like we would a glorified calculator. There are lots of built-in functions that we can apply to various kinds of input. Each built-in function has a more-or-less convenient *name* (i.e., for each built-in function there is an entry in the Global Environment that links a particular symbol to that function). However, the fun won't really begin until we can design our own functions to do whatever we want them to do. This chapter describes how to do this in the Scheme programming language.

### 9.1 Defining Functions vs. Applying Them to Inputs

#### Example 9.1.1

*In a math class, you might see a function defined using an equation such as*

$$f(x) = x^2$$

*In this case, the name of the function is  $f$ , and we might casually describe it as the squaring function—because for each possible input value,  $x$ , the corresponding output value is the square of  $x$  (i.e.,  $x^2$ ). Notice that the mathematical definition,  $f(x) = x^2$ , gives a prescription for generating appropriate output values should  $f$  ever happen to be applied to any input values. In particular, the definition of  $f$  includes an input parameter,  $x$ , that is used to refer to potential input values. In addition, the expression,  $x^2$ , on the righthand side of the equation indicates how to compute the corresponding output value for any given value of  $x$ . (The expression on the righthand side is sometimes referred to as the body of the function.) For example, if we wanted to know the output value generated by  $f$  when given 3 as its input, we could get the answer by first substituting the value 3 for  $x$  in the expression,  $x^2$ , yielding  $3^2$ . Evaluating the expression,  $3^2$ , would then yield the desired output value, 9. Similarly, if we wanted to know the output value generated by  $f$  when given the input value 4, we would first substitute the value 4 for  $x$  in the expression,  $x^2$ , yielding  $4^2$ , which evaluates to 16.*

#### Example 9.1.2

*In the preceding example, the function  $f$  took a single input value. However, we can similarly define functions that take multiple inputs. For example, the function,  $g$ , defined below, takes two inputs, represented by the input parameters  $w$  and  $h$ :*

$$g(w, h) = wh$$

*This function can be used to compute the area of a rectangle whose width is  $w$  and height is  $h$ . To apply this function to the input values, 3 and 7, we first substitute 3 for  $w$ , and 7 for  $h$  in the expression,  $wh$ , yielding  $3 \cdot 7$ . Evaluating this expression results in the desired output value, 21.*

In general, the mathematical definition of a function specifies how to generate appropriate output values should the function ever be applied to any input values. A function definition includes a list of *input parameters* and a *body*. Once a function has been defined, it can be applied to appropriate input values as follows. First, the desired input values are substituted for the appropriate input parameters in the body of the function. Next, the resulting expression is evaluated, thereby yielding the desired output value.

### Example 9.1.3

The following defines a function,  $v$ , that can be used to compute the volume of a cone:

$$v(r, h) = \frac{1}{3}\pi r^2 h$$

It has two input parameters,  $r$  and  $h$ , that respectively represent the radius and height of the cone. To compute the volume of a cone of radius 3 and height 2, we apply the function  $v$  to the input values 3 and 2, as follows. First, we substitute the values 3 and 2 for  $r$  and  $h$ , respectively, in the body,  $\frac{1}{3}\pi r^2 h$ , yielding the expression,  $\frac{1}{3}\pi(3^2)(2)$ . Evaluating this expression yields the desired output value,  $6\pi$ .

## 9.2 The lambda Special Form

The Scheme programming language provides the `lambda` special form to enable us to specify functions of our own design.

- ★ The use of the `lambda` symbol in a `lambda` special form is a tip of the cap to the fact that the underlying mathematical theory, originally developed in the 1930s, is called the *Lambda Calculus*.

Like any special form in Scheme, the `lambda` special form is a list whose first element is a keyword symbol—in this case, the symbol `lambda`. The second element in a `lambda` special form is used to specify the input parameter(s) for the function being defined. The rest of the elements in the `lambda` special form constitute the *body* of the function being defined. If you're wondering where the *name* of the function is specified, recall that the `define` special form is used to assign names to things in Scheme. Furthermore, a single function could have several different names. Thus:

- ★ The `lambda` special form specifies everything about a function *except its name*.

### Example 9.2.1: The Squaring Function in Scheme

Recall the mathematical definition of the squaring function:

$$f(x) = x^2$$

This mathematical definition does three things:

- It specifies a single input parameter,  $x$ , for the function being defined;
- It specifies a body,  $x^2$ , for the function being defined; and
- It specifies a name,  $f$ , for the function being defined.

In Scheme, the first two jobs are handled by the `lambda` special form. For example, the following `lambda` expression can be used to specify a squaring function in Scheme:

```
(lambda (x) (* x x))
```

This `lambda` expression denotes a `lambda` special form (i.e., a Scheme list whose first element happens to be the `lambda` symbol). Like any special form, a `lambda` special form has its own, special rule for being evaluated. For now, suffice it to say that:

★ *The evaluation of a lambda special form always results in a function.*

Thus, if the expression, `(lambda (x) (* x x))`, is typed into the Interactions Window, DrScheme will report that its evaluation yields a function, as illustrated below:

```
> (lambda (x) (* x x))
#<procedure>
```

Admittedly, the character sequence generated by DrScheme is not very descriptive. It simply says that the evaluation of the corresponding lambda special form has resulted in a function.

★ *At this point, it is important to stress that the function has been created; however, it has not yet been applied to any inputs!*

We can demonstrate that the function created above behaves like a squaring function by first giving it a name and then applying it to a variety of input values. The following Interactions Window session demonstrates how to name our function:

```
> (define square (lambda (x) (* x x)))
>
```

The `define` special form is used to create an entry in the Global Environment that associates the `square` symbol with the function specified by the lambda expression. Recall that when a `define` special form is evaluated, the given symbol—in this case, `square`—is not evaluated; however, the given expression—in this case, `(lambda (x) (* x x))`—is evaluated. Thus, the value associated with the `square` symbol is the function that results from evaluating the given lambda special form, as demonstrated below.

```
> square
#<procedure:square>
```

Once we have given a name to our function, we can then use it like any of the built-in functions, as demonstrated below:

```
> (square 3)
9
> (square 4)
16
> (square -8)
64
```

Each of the above expressions is evaluated using the Default Rule for evaluating non-empty lists. In each case, the `square` symbol evaluates to the function that we defined earlier, which is then applied to the desired input value.

### Example 9.2.2

Incidentally, it is possible to define and apply a function without ever having given it a name, as the following Interactions Window session demonstrates:

```
> ((lambda (x) (* x x)) 4)
16
```

The Default Rule for evaluating non-empty lists is used to evaluate the above expression. In the process, each element of the list is evaluated. The first element of the list is the `lambda` special form, which evaluates to the (unnamed) squaring function. The second element of the list evaluates to the number four. The result of applying that function to that input yields the desired output, sixteen. Later on, we shall encounter situations where it is convenient to use functions without bothering to name them.

### Example 9.2.3

The following Interactions Window session demonstrates how to define, name, and apply functions analogous to the functions,  $g(w, h) = wh$  and  $v(r, h) = \frac{1}{3}\pi r^2 h$ , seen earlier:

```
> (define rect-area (lambda (w h) (* w h)))
> (rect-area 2 3)
6
> (rect-area 3 8)
24
> (define cone-volume (lambda (r h) (* 1/3 3.14159 r r h)))
> (cone 1 3)
3.14159
> (cone-volume 10 1)
104.71966666666665
```

In the `cone-volume` function, 3.14159 is used as an approximation of  $\pi$ , and the expression, `(* 1/3 3.14159 r r h)`, takes advantage of the fact that the built-in multiplication function can be applied to any number of input values.

## 9.3 The Syntax and Semantics of Lambda Expressions

This section presents the syntax and semantics of lambda expressions. Initially, it restricts attention to those in which the *body* consists of a single expression; later, it addresses those in which the body consists of multiple expressions.

### 9.3.1 The Syntax of a Lambda Expression

A lambda expression has the following syntax:

$$(\text{lambda } (C_1 C_2 \dots C_n) B)$$

where:

- each  $C_i$  is a character sequence denoting some Scheme symbol,  $s_i$ ;
- the symbols,  $s_1, s_2, \dots, s_n$ , are distinct (i.e., there are no duplicates); and
- $B$  is a character sequence denoting a Scheme datum,  $\mathcal{D}$ , of any kind.

Thus,  $C_1, C_2, \dots, C_n$  specify  $n$  distinct *input parameters* for the lambda expression, and  $B$  specifies the *body* of the lambda expression.

### Example 9.3.1

The following are examples of well-formed lambda expressions:

- `(lambda () 44)`



- `(lambda (x) (* x x))`
- `(lambda (w h) (* w h))`
- `(lambda (r h) (* 1/3 3.14159 r r h))`
- `(lambda (x y z) (* x (- y z)))`

For the last expression, `(x y z)` specifies the parameter list and `(* x (- y z))` specifies the body.

### Example 9.3.2

In contrast, the following are examples of malformed lambda expressions:

- `(lambda (x y x) (* x y))`
- `(lambda (x 10) (* x 10))`
- `(lambda x)`

## 9.3.2 The Semantics of a Lambda Expression

The semantics of a lambda expression stipulates the Scheme datum that the lambda expression denotes, as well as how that Scheme datum is evaluated. As suggested by the preceding examples, a lambda expression invariably denotes a list—called a `lambda special form`—and the evaluation of that list invariably results in a Scheme function. The semantics of the lambda expression also includes a description of the subsequent behavior of that function should it ever be applied to any input(s).

**The list denoted by a `lambda special form`.** Assuming that

- each  $C_i$  denotes a Scheme symbol,  $s_i$ ;
- the symbols,  $s_1, s_2, \dots, s_n$ , are distinct; and
- $B$  denotes some Scheme datum  $\mathcal{D}$ ,

then a lambda expression of the form

$$(\text{lambda } (C_1 C_2 \dots C_n) B)$$

denotes a Scheme *list* whose elements are as follows:

- the `lambda` symbol;
- a list containing  $n$  distinct symbols,  $s_1, s_2, \dots, s_n$ ; and
- the Scheme datum,  $\mathcal{D}$

This list is referred to as a `lambda special form`.

**Note.** By now, you should be getting used to the fact that a piece of syntax, such as `(lambda (x) (* x x))`, denotes a Scheme datum—in this case, a Scheme *list* containing the `lambda` symbol and two subsidiary lists. Although it is important to be able to distinguish expressions from the Scheme data they denote, doing so can get quite tedious in chapter after chapter. Therefore, for the sake of expository convenience, the rest of this book shall frequently blur this distinction. Thus, we may talk of the list, `(1 2 3)`, even though we really mean the list *denoted by the expression* `(1 2 3)`. Similarly, we may say that the expression `(lambda (x) (* x x))` evaluates to a function, when we really mean that the *list* denoted by the expression `(lambda (x) (* x x))` evaluates to a function.

### The Evaluation of a `lambda` Special Form

- ★ The most important thing to know about the evaluation of a `lambda` special form is that the result is invariably a *function*; however, the evaluation of a `lambda` special form only *creates* the function; it does not *apply* it to any input(s).

For convenience, we shall refer to such functions as `lambda` functions. Thus, a `lambda` function is a function that resulted from having evaluated a `lambda` special form.

Although evaluating a `lambda` special form only creates the corresponding function, it is necessary to describe what that function would do *if* it ever were applied to input values. That is the subject of the next section.

### 9.3.3 Applying a `lambda` Function to Input Values

Up to this point, the only environment that we have considered has been the Global Environment. However, when a `lambda` function is applied to inputs, the expressions in the function's body are evaluated with respect to an automatically-created *local* environment. As will be seen, the relationship between the Global Environment and the new local environment is one of *inclusion*: the local environment can be thought of as a smaller room that sits inside the Global Environment.

- ★ For the purposes of this chapter, it is assumed that the `lambda` function was created by evaluating its `lambda` special form with respect to the Global Environment, as has been the case in all of the preceding examples.

#### Example 9.3.3: Applying the Squaring Function

Consider the expression, `(lambda (x) (* x x))`. As noted above, it evaluates to a Scheme function. When this `lambda` function is applied to some input value, say 4, the following things happen:

- A local environment is created that contains a single entry in which the symbol `x` has the value 4.
- The expression, `(* x x)`, which constitutes the body of the function, is evaluated with respect to the newly created local environment. This means that: (1) any occurrence of the symbol `x` is evaluated using the entry for `x` in the local environment, ignoring any entry for `x` that might exist in the Global Environment; and (2) all other symbols are evaluated with respect to the Global Environment. The evaluation of `(* x x)` therefore yields the result 16, because `x` evaluates to 4 in the local environment, and `*` evaluates to the built-in multiplication function in the Global Environment.
- That value, 16, is taken to be the output value that results from applying the `lambda` function to the input value 4.

This process is illustrated in Fig. 9.1.

Notice that expressions in the body of a function can refer to data that are stored in one of two places:

- (1) the environment within which the function was created—in this case, the Global Environment; or
- (2) the local environment that contains entries associated with the input parameters.

#### Example 9.3.4: Computing the Volume of a Sphere

You may recall that the volume of a sphere of radius,  $r$ , is given by the function  $f(r) = \frac{4}{3}\pi r^3$ . Thus, for example, the volume of a sphere of radius 1 is  $\frac{4}{3}\pi$ ; and the volume of a sphere of radius 2 is  $\frac{32}{3}\pi$ .

The following Interactions Window session first creates a global variable, `pi`, to hold the value 3.14159. It then defines a function, named `sphere-volume`. Finally, it applies this function to some sample input values.

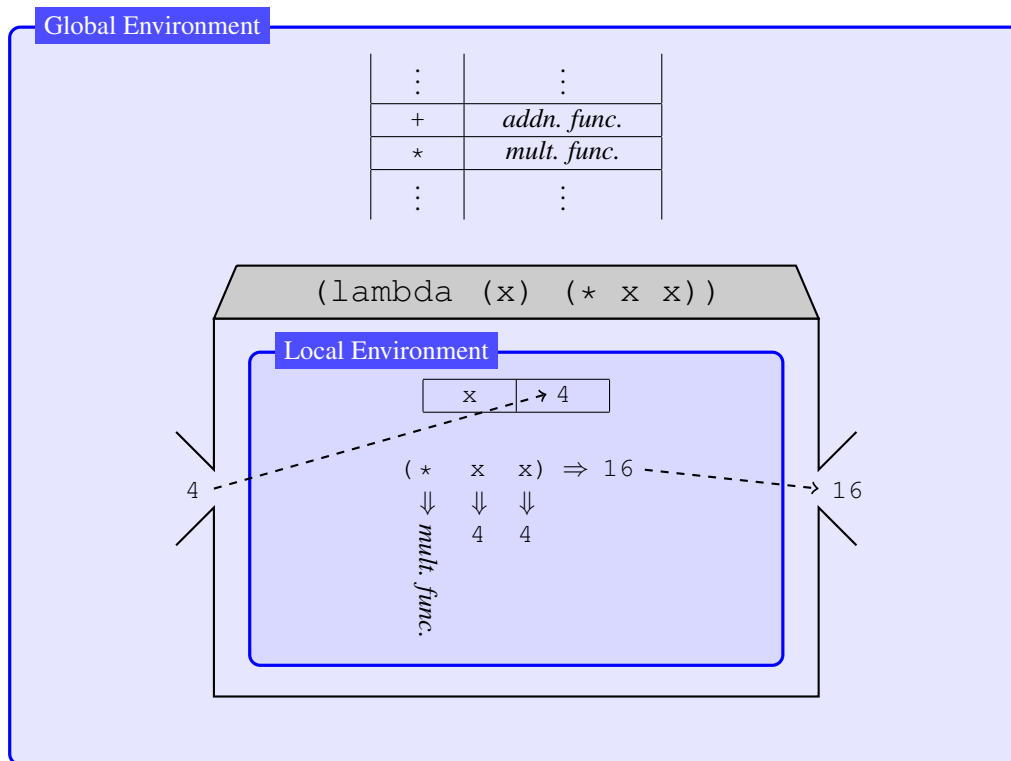


Figure 9.1: Applying the function corresponding to `(lambda (x) (* x x))` to the value 4

```
> (define pi 3.14159)
> (define sphere-volume (lambda (r) (* 4/3 pi r r r)))
> (sphere-volume 1)
4.1887866666666666
> (sphere-volume 2)
33.510293333333333
```

Consider the evaluation of the expression, `(sphere-volume 2)`. It involves the following steps:

- First, a local environment is set up containing a single entry in which the symbol `r` has the value 2.
- Next, the expression, `(* 4/3 pi r r r)`, which constitutes the body of the function, is evaluated with respect to that local environment. In the process, the `*` symbol evaluates to the built-in multiplication function, `4/3` evaluates to itself, the symbol `pi` evaluates to 3.14159, and the symbol `r` evaluates to 2. Applying the multiplication function to the values `4/3`, 3.14159, 2, 2 and 2 yields the result: 33.510293333333333.
- Finally, the value 33.510293333333333 is reported as the output value generated by applying the `sphere-volume` function to the input value 2.

Notice that in the second step, the value for `r` came from the local environment, whereas the values for `*` and `pi` came from the Global Environment.

- ★ When evaluating a symbol such as `r` or `pi` with respect to a local environment, if the symbol has an entry in the local environment, that entry is used; otherwise, the symbol's value is derived from the Global Environment.

The evaluation of `(sphere-volume 2)` is illustrated in Fig. 9.2.

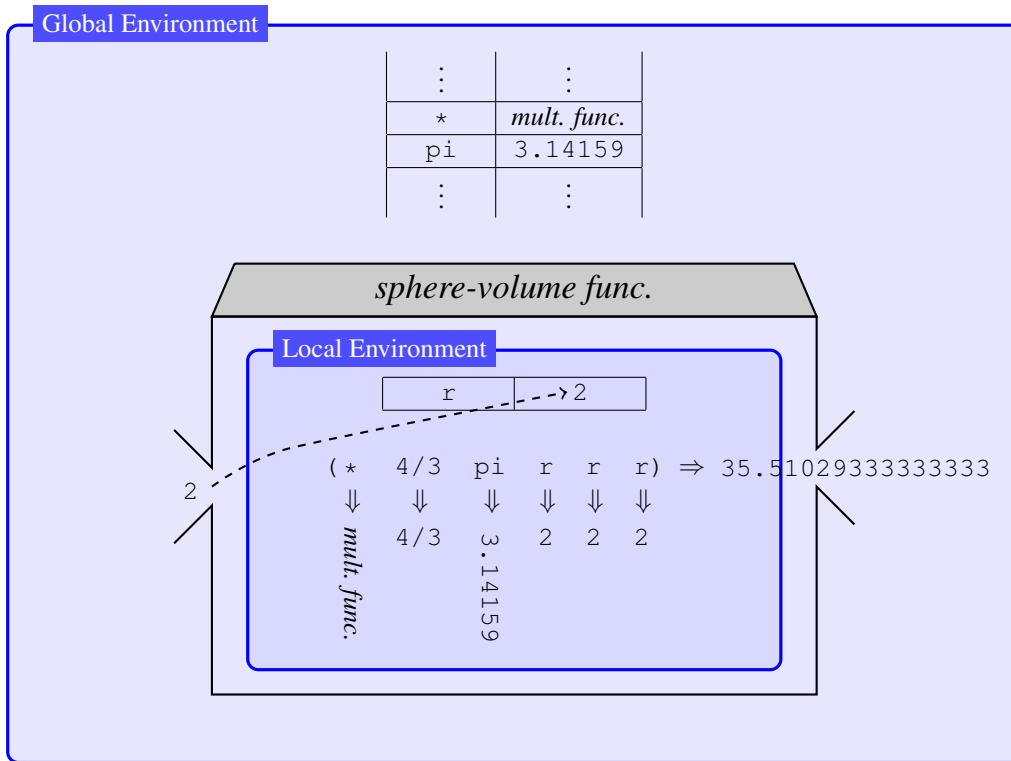


Figure 9.2: Applying the sphere-volume function to the value 2

The following Interactions Window session (continuing from the one given above) illustrates that the existence of a global variable named `r` has no effect on the local variable that also happens to be named `r`. In contrast, changing the value of the global variable, `pi`, has disastrous effects! (That is one of many reasons why the use of global variables should be very carefully restricted!)

```
> (define r 55)
> (sphere-volume 1)
4.1887866666666666
> (sphere-volume 2)
33.510293333333333
> (define pi 100)           ← Yikes!!
> (sphere-volume 1)       ← Yikes!!
400/3
```

### Example 9.3.5: More Complex Input Expressions

So far, the examples have involved simple input expressions such as `1` or `2`. This example demonstrates that complex input expressions can be handled without requiring any new evaluation tools. Consider the following Interactions Window session:

```
> (define square (lambda (x) (* x x)))
> (square (+ 2 3))
```

```
25
> (square (- 8 5))
9
> (square (square 10))
10000
```

The evaluation of the first expression simply defines a squaring function, as seen in previous examples. The evaluation of the expression, `(square (+ 2 3))`, is done according to the Default Rule for evaluating non-empty lists. In particular:

- The `square` symbol evaluates to the squaring function;
- The expression, `(+ 2 3)`, evaluates to 5;
- The squaring function is applied to the input value 5, generating the output value 25.

Similar remarks apply to the evaluation of `(square (- 8 5))` and `(square (square 10))`. In each case, the input expressions, no matter how complex, are evaluated first to generate the corresponding input values. For example, the evaluation of `(square (square 10))` involves the following steps:

- The `square` symbol evaluates to the squaring function;
- The expression, `(square 10)`, evaluates to 100;
- The squaring function is applied to 100, yielding the output value, 10000.

Notice that the evaluation of the input expression, `(square 10)`, itself required using the Default Rule for evaluating non-empty lists. In particular:

- The `square` symbol evaluates to the squaring function;
- The expression, `10`, evaluations to 10; and
- The squaring function is applied to 10, yielding the output value 100.

### Example 9.3.6

Here's an example of a function that takes more than one input (i.e., parameter).

```
> (define discriminant
  (lambda (a b c)
    (- (* b b) (* 4 a c))))
> (discriminant 1 2 -4)
20
> (discriminant 1 0 -3)
12
```

Notice that the syntax of Scheme allows expressions to occupy multiple lines. This is quite useful when writing longer expressions. DrScheme automatically indents sub-expressions to make longer expressions easier to read. Hitting the tab key will automatically cause the current line to snap to the appropriate amount of indentation.

**Differences Between Mathematical Notation and Lambda Notation**

Recall that in a math class, you might define a function using an equation such as  $f(x) = x^2$ . Later on, you might apply that function to various inputs, using expressions such as  $f(3) = 9$  or  $f(5) = 25$ .

In Scheme, we can use a lambda special form to define a function without giving it a name. For example, we might evaluate `(lambda (x) (* x x))` to create a squaring function. However, we cannot replace the parameter `x` in that lambda expression by arbitrary expressions. For example, `(lambda (3) (* 3 3))` is malformed in Scheme. (Recall the rules of syntax for lambda expressions.) But we can see a similarity to the common mathematical notation for applying functions to inputs as follows.

**Example 9.3.7: lambda functions vs. mathematical functions**

```
> (define f (lambda (x) (* x x)))
> (f 3)
9
> (f (+ 2 3))
25
> (f (f 10))
10000
```

*The corresponding mathematical equations/expressions would be:*

$$f(x) = x^2$$

$$f(3) = 9$$

$$f(2 + 3) = 25$$

$$f(f(10)) = 10000$$

**Example 9.3.8: A Lambda Expression with a Bigger Body**

*The following illustrates that a lambda expression can have more than one expression in its body.*

```
> (define useless-function
  (lambda (input)
    input
    (* input input)
    (* input input input)
    input
    ()))
> (useless-function 35)
()
> (useless-function 888)
()
```

*In this case, the body of the function includes five expressions (i.e., everything after the parameter list).*

- \* *The semantics of Scheme stipulates that when a lambda function having multiple expressions in its body is subsequently applied to input(s), the expressions in the body are evaluated sequentially, one after the other.*
- \* *Furthermore, the value of the last expression in the body is taken to be the output value for the function.*

*Thus, in the above example, each of the expressions in the body is evaluated in turn, and the value of the last expression (i.e., `()`) serves as the output value. This function is kind of silly since the results of evaluating the first four expressions in its body are thrown away.*

*★ The only way that intermediate expressions in the body of a function could have any impact is if they caused side effects.*

Up to this point, the only function that we have seen that has side effects is the built-in `printf` function. It displays the contents of a string in the Interactions Window. This is a harmless side effect that can be very useful.

## 9.4 Summary

This chapter introduced the `lambda` special form whose purpose is to enable a Scheme programmer to specify functions. A `lambda` special form includes:

- (1) the `lambda` symbol;
- (2) a list of input parameters; and
- (3) one or more expressions constituting the *body* of the function.

The result of evaluating a `lambda` special form is always a function. For example, the result of evaluating `(lambda (x) (* x x))` is a function whose sole input parameter is `x`, and whose body is `(* x x)`.

In Scheme, the following are distinct:

- The function that is generated by evaluating a `lambda` special form;
- Any name(s) that might be given to that function; and
- The process of applying that function to input(s).

The `define` special form is used to give names to things, including functions. For example, the following expression associates the squaring function with the name `square`.

```
(define square
  (lambda (x)
    (* x x)))
```

The application of this function to an input is handled by the evaluation of an expression such as `(square 10)`, which is carried out by the Default Rule for evaluating non-empty lists.

The application of a `lambda` function involves the creation of a local environment that contains one entry for each input parameter. The input values to which the function is being applied become the values associated with the corresponding input parameters in the local environment. For example, when applying the squaring function to the input value `10`, the input parameter `x` receives the value `10` in the local environment. Next, each expression in the body of the function is evaluated with respect to that local environment. In particular, any symbol `s` that must be evaluated is evaluated by looking first for a corresponding entry in the local environment; if no entry for `s` is found there, then the Global Environment is checked. In other words, the local environment has higher priority when evaluating symbols in the body of a `lambda` function. Thus, when evaluating `(* x x)` in the body of the squaring function, `x` evaluates to `10`, courtesy of the local environment, whereas `*` evaluates to the built-in multiplication function courtesy of the Global Environment. Finally, the output obtained by evaluating the last expression in the body of the function is taken to be the result of applying the function to the given input(s). Thus, the output `100`, obtained by evaluating `(* x x)`, is taken to be the output value for the application of the squaring function to the input value `10`.

The parameter lists in a `lambda` special form may specify zero or more parameters, each represented by a Scheme symbol. And the body of a `lambda` special form may include one or more expressions. However, it is only reasonable to include more than one expression in the body of a function if the evaluation of those expressions cause some side effects.

### **Special Forms Introduced in this Chapter**

lambda Used to specify functions of our own design.



# Chapter 10

## Some practicalities

This chapter introduces the following practicalities:

- Further capabilities of the built-in `printf` function. This function, which takes a *string* as one of its inputs, can be used to display nicely formatted information in DrScheme's Interactions Window. Its functionality is similar to that of the `format/print` functions found in many programming languages.
- The built-in `load` function. This function causes the Scheme expressions in a specified file to be evaluated as though they had been manually typed into the Interactions Window. As such, these expressions are evaluated with respect to the Global Environment. In this way, a library of useful Scheme definitions can be incorporated into your own program quite easily. The name of the file is specified by a *string*.
- Comments. A comment is a piece of syntax that DrScheme completely ignores. Comments are used by programmers to help clarify—for people—what the program/code is supposed to do.

### 10.1 More Fun with the Built-in `printf` Function

Recall from Section 5.5 that the built-in `printf` function can be applied to strings to generate *side-effect printing* in the Interactions Window, as illustrated below.

```
> (printf "you are amazing!")  
you are amazing!
```

Unlike the input string `"you are amazing!"`, which is a Scheme datum, the text displayed in the Interactions Window is not a Scheme datum; instead, it is merely something that happens on the side. The output value generated by the `printf` function is the *void* datum.

**Escape sequences.** In the above example, the `printf` function effectively copied the contents of the input string into the Interactions Window *verbatim*. However, the `printf` function sometimes deviates from this simple behavior. In particular, as the `printf` function walks through the input string, it reacts to a few special character sequences in special ways. For example, it reacts to the character sequence `~%`, by moving to a new line in the Interactions Window (i.e., it interprets `~%` as a *newline* character). It also interprets `\n` as a newline character. In addition, whenever it encounters the character sequence, `~s`, in the input string, the `printf` function treats it as a place-holder for a piece of data to be displayed, as discussed in Example 10.1.1 below. Because the character sequences `~%`, `\n`, and `~s` are not interpreted literally, but involve the `printf` function *escaping* from a literal interpretation, they are frequently called *escape sequences*. (And the characters `~` and `\` that introduce escape sequences are sometimes called escape characters.) Although the `printf` function can deal with a variety of other escape sequences, these are the only ones that we'll need for this course. Their use enables the `printf` function to generate nicely formatted text in the Interactions Window. For this reason, the input string is frequently called a *format string*—which explains the `f` in `printf`.

In summary, the `printf` function causes the contents of the format string (i.e., its first input) to be displayed *verbatim* in the Interactions Window, except that:

- the quotation marks are omitted;
- each instance of `~%` or `\n` is interpreted as a newline character and, thus, causes subsequent text to be displayed on the next line in the Interactions Window; and
- each instance of the escape sequence, `~s`, is replaced by a character sequence representing the *value* of the corresponding input expression.

Notice that if the format string contains  $n$  instances of `~s`, then there must be  $n$  input expressions following the format string, as follows:

```
(printf format-string expr1 ... exprn)
```

#### Example 10.1.1: Formatted printing with `printf`

The following Interactions Window session illustrates the use of the escape sequences `~%`, `\n` and `~s` by the `printf` function.

```
> (printf "Hi there!\nBye there!")
Hi there!
Bye there!
> (printf "Oh, I get it!~%This sentence begins on a new line!")
Oh, I get it!
This sentence begins on a new line!
> (printf "First thing: ~s, second thing: ~s~%" (+ 2 3) (* 6 7))
First thing: 5, second thing: 42
> (printf "Line One!~% Line Two!!~% Line Three!!!~%")
Line One!
  Line Two!!
    Line Three!!!
> (printf "First ===> ~s, Second ===> ~s, Third ===> ~s~%"
        (+ 4 2) (- 9 6.3) (* 4 100))
First ===> 6, Second ===> 2.7, Third ===> 400
> (printf "A symbol: ~s, a string: ~s, a boolean: ~s~%"
        'I-am-a-symbol
        "I am a String!"
        (> 4 2))
A symbol: I-am-a-symbol, a string: "I am a String!", a boolean: #t
```

#### Example 10.1.2: The `printf` function and the *void datum*

The following interaction demonstrates that the `printf` function generates the *void datum* as its output:

```
> (void? (printf "hi\n"))
hi
#t
```

In this example, the Default Rule for evaluating non-empty lists is used to evaluate the expression, `(void? (printf "hi\n"))`. First, each element of the list is evaluated:

- the `void?` symbol evaluates to the built-in `void?` function; and
- `(printf "hi\n")` evaluates to the *void datum*—while causing `hi` to be displayed in the Interactions Window as a side effect.

Next, the `void` datum is fed as input into the `void?` type-checker predicate, resulting in the output value `#t`. Thus, `hi` is side-effect printing, while `#t` is the output value.

Although DrScheme does not normally display the `void` datum, we can force it to do so, as follows:

```
> (printf "Show us void: ~s~%" (void))
Show us void: #<void>
```

However, keep in mind that `#<void>` is not legal Scheme syntax. If you enter `#<void>` into the Interactions Window, you'll get a red error message!

**Including multiple expressions within the body of a lambda function.** Recall that the body of a lambda function may contain multiple expressions. When such a function is called, each of the expressions in the body is evaluated in turn. However, it is only the value of the *last* expression in the body that determines the output value for the function call. Since the output values of earlier expressions are ignored, it only makes sense to include multiple expressions in the body of a function if some of those expressions generate side effects. The following example considers a function whose body contains expressions that generate side-effect printing.

### Example 10.1.3

The following lambda function, called `verbose-func`, contains multiple expressions in its body. When the `verbose-func` is called, each expression in its body is evaluated. The first four expressions cause the built-in `printf` function to be called, thereby generating several lines of side-effect printing in the Interactions Window. However, it is the evaluation of the last expression in the function's body that generates an output value for the function call.

```
> (define verbose-func
  (lambda (a b)
    (printf "Hi. This is verbose-func!~%" )
    (printf "The value of the first input is: ~s~%" a)
    (printf "The value of the second input is: ~s~%" b)
    (printf "Their product is:~%" )
    (* a b)))
> (verbose-func 3 4)
Hi. This is verbose-func!
The value of the first input is: 3
The value of the second input is: 4
Their product is:
12
>
```

In this case, the output value of the function call is twelve, which DrScheme displays in one color; the previous four lines of text are just side-effect printing, which DrScheme displays in a different color.

★ Part 1 of this book explores how much can be accomplished *without* using side effects. Therefore, most of the functions we write will include only a single expression in the body. However, we will sometimes use the `printf` function to generate useful side-effect printing in the Interactions Window.

**Example 10.1.4: Defining a useful `tester` function**

The `printf` function can be used to define a `tester` function that will greatly facilitate the testing of whatever Scheme function we happen to be creating. The `tester` function can also be used to test our understanding of how arbitrary Scheme data get evaluated.

```
(define tester
  (lambda (datum)
    (printf "~s ==> " datum)
    (eval datum)))
```

The `tester` function takes any Scheme datum as its input. As a side effect, it prints out a representation of that datum in the Interactions Window. For its output value, it simply evaluates the input datum. The following Interactions Window session demonstrates its use.

```
> (tester '(+ 1 2))
(+ 1 2) ==> 3
> (tester (+ 1 2))
3 ==> 3
> (tester '+)
+ ==> #<primitive:+>
> (tester +)
#<primitive:+> ==> #<primitive:+>
```

These examples demonstrate that the `tester` function is most useful when the quote special form is used to shield the desired input expression from evaluation. For example, notice the difference between the evaluations of `(tester '(+ 1 2))` and `(tester (+ 1 2))`. In the first case, `(+ 1 2)` is shielded from evaluation by the quote special form; thus, the list `(+ 1 2)` is fed as input to the `tester` function. That is why `(+ 1 2)` is printed out in the Interactions Window before the arrow. After that side-effect printing, the `eval` function is then used to explicitly evaluate the list `(+ 1 2)`, generating the output value 3. Since the formatting string given to `printf` does not include a newline character, the side-effect printing and the output value are both displayed on the same line.

★ The `tester` function is one of the rare cases where the built-in `eval` function is explicitly invoked.

## 10.2 The Built-in `load` Function

Scheme includes a built-in `load` function that causes all of the Scheme expressions in a specified file to be evaluated in an Interactions Window session. Here's the contract:

Name:	<code>load</code>
Input:	<i>filename</i> , a string
Output:	The result of evaluating the <i>last</i> expression in the file named <i>filename</i> .
Side Effect:	Whatever side effects result from evaluating <i>all</i> of the expressions in the file named <i>filename</i> .

**Example 10.2.1**

Suppose the file `"test.txt"` contains the following expressions:

```
(printf "Loading test.txt!!")

(define tester
```

```
(lambda (datum)
  (printf "~s ==> " datum)
  (eval datum)))

(define x 34)
```

Then the following Interactions Window session could ensue:

```
> x
BUG! reference to undefined identifier: x
> tester
BUG! reference to undefined identifier: tester
> (load "test.txt")
Loading test.txt!!
> x
34
> (tester 'x)
x ==> 34
```

Notice that the first attempts to evaluate `tester` and `x` generated errors because there were not yet any entries for these symbols in the Global Environment. However, after loading the file `test.txt`, subsequent attempts to evaluate `x` and to use `tester` succeed.

This example demonstrates that useful function definitions can be conveniently stored in a file, to be loaded whenever needed.

- ★ The *Run* button on DrScheme's toolbar is similar to the `load` function, except that it causes the Scheme expressions currently residing in the Definitions Window to be evaluated within a fresh Interactions Window session.

## 10.3 Comments

In Scheme programs, the semi-colon character is used to initiate *comments*. The text that constitutes a comment is ignored by DrScheme, as illustrated by the following example.

### Example 10.3.1

```
(define tester
  (lambda (datum)
    ;; Print (the value of) DATUM -- without a newline character
    (printf "~s ==> " datum)
    ;; Then explicitly evaluate (the value of) DATUM
    (eval datum)))

;; Sample TESTER expressions
;; -----

(tester '(+ 2 3))
(tester (+ 2 3))
```

Evaluating the above code in the Interactions Window would have the same result as evaluating the following, uncommented code:

```

(define tester
  (lambda (datum)
    (printf "~s ==> " datum)
    (eval datum)))

(tester '(+ 2 3))
(tester (+ 2 3))

```

The purpose of comments is to make a Scheme program easier for people to understand. DrScheme ignores the comments completely.

**Contracts in Scheme programs.** One of the most important uses of comments is to enable a Scheme program to include an explicit contract for each function it defines. The following example illustrates the format for contracts that will be used for the rest of the course.

### Example 10.3.2: A contract for the squaring function

*The following comment block constitutes a contract for the squaring function seen in Example 9.2.1.*

```

;; SQUARE
;; -----
;; INPUT:   X, a number
;; OUTPUT:  The value X*X (i.e., X squared)

```

*My personal convention is to use upper-case letters for the names of the function and its inputs, while the actual Scheme code uses lower-case letters.*

*★ Aside from this difference, the names of the function and its inputs in the contract should match the corresponding names in the actual function definition.*

*By convention, if a function does not generate any side effects, then the contract need not mention side effects.*

### Example 10.3.3: A contract for the `tester` function

*The following code fragment includes a contract for the `tester` function followed by the actual function definition. Note that a blank line should separate the contract from the function definition.*

```

;; TESTER
;; -----
;; INPUT:   DATUM, any Scheme datum
;; OUTPUT:  The result of evaluating (the value of) DATUM
;; SIDE EFFECT: Displays (the value of) DATUM *before*
;;              evaluating it

(define tester
  (lambda (datum)
    ;; Display (the value of) DATUM
    (printf "~s ==> " datum)
    ;; Evaluate (the value of) DATUM
    (eval datum)))

```

- ★ To avoid being overly cumbersome, contracts may intentionally blur the distinction between the *names* of input parameters—which are symbols—and their *values*—which can be anything.

#### Example 10.3.4: Revised contract for `tester`

*Instead of (correctly) saying that the `tester` function displays (the value of) `datum` before evaluating (the value of) `datum`, a typical contract might say that the `tester` function displays `datum` before evaluating it. (Even though the symbol `datum` is not what is displayed by `tester`!) In effect, the contract is using the symbol `datum` to refer to its value in the local environment, much as a person uses the name Barack Obama to refer to the 44<sup>th</sup> president of the United States. Of course, you should never let the true distinction between a symbol and its value stray too far from conscious awareness!*

```
;; TESTER
;; -----
;; INPUT:  DATUM, any Scheme datum
;; OUTPUT: The result of evaluating DATUM
;; SIDE EFFECT: Displays DATUM *before* evaluating it

(define tester
  (lambda (datum)
    ;; Display DATUM
    (printf "~s ==> " datum)
    ;; Evaluate DATUM
    (eval datum)))
```

## 10.4 Summary

This chapter introduced the built-in `printf` function, the built-in `load` function, and *comments*.

Almost any character sequence that begins and ends with double quotes denotes a *string* datum in Scheme. (The exceptions (e.g., `"hi\"`) involve escape sequences (e.g., `\`) that effectively *capture* the final double quote. They need not concern us.) For example, `"the brown dog\n"` and `"i am a fox"` both denote strings in Scheme.

The built-in `printf` function has the useful side effect of displaying text in the Interactions Window. The `printf` function takes a string—sometimes called a *formatting string*—as its first input. That string may include escape sequences such as `~%`, `\n` and `~s` that are interpreted in special ways by the `printf` function. In particular, the `printf` function interprets each character of the formatting string literally, except that `~%` and `\n` are interpreted as *newline* characters, and `~s` is interpreted as a placeholder for a piece of data. For each occurrence of `~s` in the formatting string, there must be a corresponding additional input to `printf`. Thus, if the formatting string includes *n* occurrences of `~s`, then there must be *n* additional inputs to `printf` after the formatting string, as illustrated below:

```
> (printf "One: ~s, Two: ~s, Three: ~s~%" 1 2 (+ 1 2))
One: 1, Two: 2, Three: 3
```

Notice that the double quotes from the formatting string are not displayed in the Interactions Window.

The `tester` function was defined to use `printf` to display a datum before evaluation, and then to explicitly use the built-in `eval` function to evaluate that datum. When using the `tester` function, input expressions are typically quoted to shield them from evaluation by the Default Rule, as illustrated below:

```
> (tester '(+ 1 2))
(+ 1 2) ==> 3
```

The built-in `load` function can be used to *load* the contents of a file automatically, instead of having to manually type its contents directly into the Interactions Window. The input to the `load` function is a string representing the name of the file. For example, if `myfile.txt` contains a bunch of function definitions, then the expression `(load "myfile.txt")` would cause those function definitions to be evaluated by DrScheme just as though they had been manually typed into the Interactions Window. Those functions could then be used during the remainder of the Interactions Window session. DrScheme's *Run* button is similar, except that it loads the expressions currently residing in the Definitions Window into the Interactions Window.

Finally, the semi-colon is a character that is used to introduce *comments* in Scheme. In particular, any sequence of characters that starts with a semi-colon and continues to the end of the line is completely ignored by DrScheme. An effective programmer uses concise comments to explain what their code is (supposed to be) doing. One important use of comments is to provide a *contract* for each function that is defined in a given program.

### Built-in Functions Introduced in this Chapter

<code>printf</code>	To do side-effect printing in the Interactions Window
<code>load</code>	To load the contents of a file



# Chapter 11

## Conditional Expressions I

Solving problems often involves making decisions or choosing from a set of alternatives. For example, to determine the appropriate letter grade for a given exam score, one might reason as follows: “If the grade is at least 90, output *A*; otherwise, if the grade is at least 80, output *B*, and so on.” The simplest kind of programming decision is a *binary* decision (i.e., choosing one of two alternatives). In English, a binary decision can be represented by a sentence of the form, “If some condition holds, then do one thing; otherwise, do something else.” For example: “If it is raining, take an umbrella; otherwise, wear sunglasses.” In this case, the *condition* is whether or not it is raining; the *then* clause is “take an umbrella”; and the *else* clause is “wear sunglasses”.

In Scheme, programmers use the `if` special form to make binary decisions. The `if` special form is an example of a *conditional expression*. Like many conditional sentences in English, an `if` special form has a *condition*, a *then* clause, and an *else* clause. For example, `(if (> x y) x y)` is an instance of an `if` special form, where the condition is `(> x y)`, the *then* clause is `x`, and the *else* clause is `y`. Like any special form, an `if` special form is evaluated in its own special way. Importantly, the evaluation of an `if` special form depends upon whether the condition evaluates to true.

Although a single `if` special form can only make a binary decision, multiple `if` special forms can be *nested* to, in effect, make an *n-ary* decision (i.e., a decision to select one from among *n* choices), as in: “If the grade is at least 90, give an *A*; otherwise, if the grade is at least 80, give a *B*; otherwise, if the grade is at least 70, ...”

- ★ The evaluation of the `if` special form is *lazy* in the sense that only the computations needed to ascertain the final value are actually performed.

This chapter also introduces the `when` special form, which is useful in cases where an *else* expression is not needed. Chapter ?? introduces the `cond` special form, which facilitates making *n-ary* decisions.

### 11.1 The `if` Special Form

We begin by introducing the `if` special form under the assumption that its condition evaluates to an actual boolean value (i.e., `#t` or `#f`). Afterward, we will relax that assumption.

The syntax of an `if` special form is as follows:

```
(if condExpr thenExpr elseExpr)
```

where:

- *condExpr* is a condition (i.e., an expression that evaluates to `#t` or `#f`); and
- *thenExpr* and *elseExpr* are any Scheme expressions.

**Example 11.1.1**

The following expressions are legal examples of the `if` special form:

```
(if (> 2 4) (* 8 2) (* 6 5))

(if (> 4 2) 'then 'else)

(if #f "then" "else")
```

The semantics of the `if` special form stipulates that it is evaluated as follows.

- First, the condition, *condExpr*, is evaluated. Then, depending on its value, one of the following will happen.
  - **Case 1: *condExpr* evaluates to #t.** In this case, *thenExpr* is evaluated—and the value of the `if` special form is whatever *thenExpr* evaluates to.
  - **Case 2: *condExpr* evaluates to #f.** In this case, *elseExpr* is evaluated—and the value of the `if` special form is whatever *elseExpr* evaluates to.

Notice that the condition, *condExpr*, is *always* evaluated; however, after that, *one and only one* of the remaining expressions, *thenExpr* or *elseExpr*, is evaluated. We say that the evaluation of the `if` special form is *lazy*, in the sense that it only evaluates the expressions needed to compute the value of the entire `if` expression. For example, if the condition evaluates to #t, then the value of the *else* expression is not needed and, thus, it is not computed. This kind of selective evaluation is available to special forms because each special form specifies its own mode of evaluation; however, this kind of selective evaluation is not available to the Default Rule, which always evaluates every item in a non-empty list.

**Example 11.1.2**

The following interactions demonstrate the evaluation of the `if` special forms seen earlier.

```
> (if (> 2 4) (* 8 2) (* 6 5))
30
> (if (> 4 2) 'then 'else)
then
> (if #f "then" "else")
"else"
```

In the first expression, the condition, `(> 2 4)`, evaluates to #f. Thus, the *else* expression, `(* 6 5)`, is evaluated. Its value, 30, is the value of the entire `if` expression.

In the second expression, the condition, `(> 4 2)`, evaluates to #t. Thus, the *then* expression, `'then`, is evaluated. Its value, `then`, is the value of the entire `if` expression.

In the third expression, the condition, `#f`, evaluates to #f. Thus, the *else* expression, `"else"`, is evaluated. Its value, `"else"`, is the value of the entire `if` expression. (Recall that strings evaluate to themselves.)

Although the preceding examples illustrate the semantics of the `if` special form, they are kind of silly because in each case the condition has a determined value and, therefore, the entire `if` expression seems unnecessary. That is true. However, as the following example demonstrates, an `if` expression that appears in the body of a function can involve conditions that depend on the values of one or more input parameters—and those values are not known until the function is applied to inputs.

**Example 11.1.3: Using an `if` expression in the body of a function**

Below, a function, `how-big`, is defined. If given a number less than 10, its output is the symbol, `small`; otherwise, its output is the symbol, `big`.

```
;; HOW-BIG
;; -----
;; INPUT:  NUM, a number
;; OUTPUT: The symbol SMALL, if NUM is less than 10;
;;         Otherwise, the symbol BIG.

(define how-big
  (lambda (num)
    (if (< num 10)
        'small
        'big)))
```

The following interactions demonstrate its behavior:

```
> (how-big 5)
small
> (how-big 102)
big
```

Notice that the result of evaluating the condition, `(< num 10)`, depends on the value of `num` in the local environment, which is not known at the time the function is specified by the programmer; instead, the value of `num` is known only when the function `how-big` is eventually applied to some input.

- \* The values of the input parameters for a function cannot be known when the programmer is writing the body of the function. Therefore, if the programmer wants the function to do different things for different inputs, the `if` special form can be quite useful.

**In-Class Problem 11.1.1**

Define a function, called `sign`, that satisfies the following contract.

```
;; SIGN
;; -----
;; INPUT:  X, a number
;; OUTPUT: 1, if X > 0; 0, if X = 0; -1, if X < 0
```

Here are some examples of the desired behavior:

```
> (sign 3)
1
> (sign 0)
0
> (sign -4.2)
-1
```

*Hint: Start by defining a function that outputs 1 if  $x > 0$ , and 0 in all other cases.*

**The non-strict version of the `if` special form.** In the strict version of the `if` special form, the condition must be an expression that evaluates to a boolean (i.e., either `#t` or `#f`). In the non-strict version, the condition can be *any* Scheme expression, as illustrated below.

#### Example 11.1.4

The following are legal instances of the `if` special form:

```
(if 72 "yup" "nope")
(if "condie" "yup" "nope")
(if (* 3 4) 'hello 'goodbye)
```

The semantics of the non-strict version of the `if` special form is governed by the following rule:

- ★ When interpreting the value of the condition, anything other than `#f` counts as *true* (i.e., `#f` is the only Scheme datum that counts as *false*).

#### Example 11.1.5

The following Interactions Window session demonstrates the evaluation of the non-strict `if` expressions seen earlier.

```
> (if 72 "yup" "nope")
"yup"
> (if "condie" "yup" "nope")
"yup"
> (if (* 3 4) 'hello 'goodbye)
hello
```

In each case, the condition being tested evaluates to a non-boolean value. Since `#f` is the only thing that counts as false, the conditions in these examples all count as true. Thus, in each case, the `then` expression is evaluated—and the value of the `then` expression is the value of the entire `if` expression.

## 11.2 Simplifying Conditional Expressions

Conditional expressions can be used in many ways to enable Scheme functions to make finely tuned decisions amongst any number of cases. Although conditional expressions stated in English can guide your programming efforts, they can sometimes lead to solutions that are more complex than they need to be. That's okay! Once your function is working, you can focus attention on how to simplify the expressions it uses. In addition, as you gain more practice, the simpler expressions may come to mind sooner in the programming process.

At first, we restrict attention to expressions that evaluate to boolean values—that is, either `#t` or `#f`. Afterward, we consider expressions that may evaluate to any type of Scheme data, but subject to the interpretation that anything other than `#f` counts as *true*, while only `#f` counts as *false*.

#### Definition 11.1: Equivalent boolean conditions

Suppose that `boolOne` and `boolTwo` are two boolean conditions (i.e., expressions that evaluate to booleans no matter what environment they are evaluated in). The expressions, `boolOne` and `boolTwo` are called equivalent if, whenever they are evaluated with respect to the same environment, the resulting boolean values are the same. In other words, `boolOne` evaluates to `#t` if and only if `boolTwo` evaluates to `#t`.

**Example 11.2.1: Simplifying expressions involving `eq?`**

Suppose that `boolie` is a symbol whose value is a boolean (i.e., either boolean `true` or boolean `false`). Then, according to the above definition, the expression, `(eq? boolie #t)`, is equivalent to the much simpler expression, `boolie`, as demonstrated below.

```
> (define boolie #t)
> boolie
#t
> (eq? boolie #t)
#t
> (define bully #f)
> bully
#f
> (eq? bully #t)
#f
```

A similar simplification also works for any expression that evaluates to a boolean value:

```
> (eq? (> 5 2) #t)
#t
> (> 5 2)
#t
> (eq? (= 3 2) #t)
#f
> (= 3 2)
#f
```

Similarly, if `boolie` is any expression that evaluates to a boolean, then an `if` expression whose condition is `(eq? boolie #t)` can also be simplified.

```
> (define xyz #t)
> (if (eq? xyz #t) 'yes 'no)
yes
> (if xyz 'yes 'no)
yes
> (define abc #f)
> (if (eq? abc #t) 'yes 'no)
no
> (if abc 'yes 'no)
no
> (if (= 3 2) 'yes 'no)
no
> (if (> 5 2) 'yes 'no)
yes
```

★ In summary, if `boolie` is any expression that evaluates to a boolean, then the following simplifications yield equivalent expressions:

$$(eq? \textit{boolie} \textit{\#t}) \rightsquigarrow \textit{boolie}$$

$$(if (eq? \textit{boolie} \textit{\#t}) \textit{thenExpr} \textit{elseExpr}) \rightsquigarrow (if \textit{boolie} \textit{thenExpr} \textit{elseExpr})$$

Furthermore, in view of the non-strict version of the `if` special form, the latter simplification works for *any* expression *boolie*, whether it evaluates to a boolean or not.

### Example 11.2.2: Another way of simplifying `if` expressions

According Defn. 11.1, the expression, `(if (> x y) #t #f)`, is equivalent to the simpler expression, `(> x y)`. The following interactions demonstrate the equivalence in two different environments: one where  $x > y$ , and one where  $x < y$ .

```

> (define x 32)                ←— Setting up an environment where  $x > y$ 
> (define y 4)
> (if (> x y) #t #f)
#t
> (> x y)
#t

> (define x 32)                ←— Setting up an environment where  $x < y$ 
> (define y 1000)
> (if (> x y) #t #f)
#f
> (> x y)
#f

```

★ More generally, if *boolCond* is any boolean condition, then the following simplification yields an equivalent expression:

$$(\text{if } \text{boolCond } \#t \#f) \rightsquigarrow \text{boolCond}$$

So, if you ever find yourself writing an `if` expression whose *then* and *else* clauses are `#t` and `#f`, respectively, consider making the above simplification.

Next, we consider the same simplification, but applied to conditions whose evaluations do not necessarily yield boolean values. In such cases, the simplification yields equivalent expressions—as long as we consider anything other than `#f` to count as true, and `#f` to be the only thing that counts as false.

### Example 11.2.3: Simplifying `if` expressions: non-strict truth values

```

> (if 'happy #t #f)           ←— 'happy ⇒ happy, which counts as true
#t
> 'happy
happy
> (if #f #t #f)              ←— The condition #f ⇒ #f, which counts as false
#f
> #f
#f

```

In the first example, the condition `'happy` evaluates to the symbol `happy`, which counts as true. Therefore, the `if` expression evaluates to `#t`, which is the result of evaluating its *then* expression. In the second example, the expression `'happy` evaluates to the symbol `happy`, which counts as true. Thus, the expressions, `(if 'happy #t #f)` and `'happy`, are equivalent in the sense that they both evaluate to things that count as true. The third and fourth lines demonstrate the case where the condition of an `if` expression evaluates to `#f`.

**In-Class Problem 11.2.1**

Define a function, called `convert-to-boolean`, that takes any Scheme datum as its input. It should return `#t` as its output if the input is anything that counts as true; otherwise, it should return `#f`, as illustrated below.

```
> (convert-to-boolean #t)
#t
> (convert-to-boolean (+ 3 2))
#t
> (convert-to-boolean #f)
#f
```

*Hint: Use an unsimplified conditional expression!*

**11.3 The `when` Special Form**

In certain programming circumstances, you may want an `if` special form that does not need an *else* case. The `when` special form is provided to handle such circumstances. In its simplest form, the `when` special form has the following syntax:

```
(when condExpr thenExpr)
```

Such an expression is evaluated as follows. First, the condition, *condExpr*, is evaluated. If it evaluates to `#t` (or something that counts as true), then *thenExpr* is evaluated, and its value is the value for the entire `when` expression. However, if *condExpr* evaluates to `#f`, then *thenExpr* is skipped, and the value of the entire `when` expression is *void*.

**Example 11.3.1**

The following interactions demonstrate the semantics of the simplest use of the `when` special form.

```
> (when #t 3)
3
> (when (> 3 2) (* 4 5))
20
> (when (> 2 3) (* 4 5))
> (void? (when #f 3))
#t
```

Like the `if` special form, the `when` special form is most useful when used within the body of a function.

**Example 11.3.2**

Consider the following version of the `how-big` function that takes an extra input, `verbose?`. When `verbose?` is true, the function prints out some information about the inputs; otherwise, it doesn't print out anything.

```
;; HOW-BIG-V2
;; -----
;; INPUTS:  NUM, a number
;;         VERBOSE?, a boolean
```

```

;; OUTPUT: A symbol, either SMALL or BIG, depending on
;;         whether NUM < 10
;; SIDE EFFECT: When VERBOSE? is true, it prints out
;;              information about the inputs.

(define how-big-v2
  (lambda (num verbose?)
    ;; Do some side-effect printing?
    (when verbose?
      (printf "Inside HOW-BIG-V2: NUM = ~s and VERBOSE? = ~s~%"
              num verbose?))
    ;; Output value
    (if (< num 10)
        'small
        'big)))

```

*Here are some examples of its behavior.*

```

> (how-big-v2 3 #t)
Inside HOW-BIG-V2: NUM = 3 and VERBOSE? = #t
small
> (how-big-v2 3 #f)
small
> (how-big-v2 15 #t)
Inside HOW-BIG-V2: NUM = 15 and VERBOSE? = #t
big
> (how-big-v2 15 #f)
big

```

★ Because the `when` special form can evaluate to *void* (e.g., when its condition evaluates to `#f`), `when` should *not* be used to generate *output values*. Instead, like in the preceding example, `when` should only be used to generate helpful side effects (e.g., side-effect printing).

Later on, in Part II of the book, when we discuss destructive programming, we will see additional uses of the `when` special form.

**More general version of the `when` special form.** Because the `when` special form never includes an *else* expression, it can include multiple *then* expressions in its body. In this way, the body of a `when` expression is similar to the body of a *lambda* function. In general, the `when` special form has the following syntax:

```

(when condExpr
  expr1
  expr2
  ...
  exprn)

```

The semantics of the `when` special form stipulates that it is evaluated as follows. First, the expression, *condExpr*, is evaluated. If it evaluates to something that counts as true, then the expressions, *expr*<sub>1</sub>, ..., *expr*<sub>*n*</sub>, are evaluated in turn, and the value of the last expression serves as the value of the entire `when` expression. However, if *condExpr* evaluates to `#f`, then the subsidiary expressions, *expr*<sub>1</sub>, ..., *expr*<sub>*n*</sub>, are skipped, and the entire `when` expression simply evaluates to *void*.



**In-Class Problem 11.3.1**

Modify the `how-big-v2` function so that it includes a `when` expression that has multiple `printf` expressions in its body.

## 11.4 Summary

This chapter introduced the `if` special form for making binary decisions. When evaluating conditions, the `if` special form accommodates *non-strict* truth values. In particular, anything other than `#f` counts as *true*. Equivalently, only `#f` counts as *false*.

An `if` special form has the form, `(if condExpr thenExpr elseExpr)`. An `if` special form is evaluated as follows. First, the condition, `condExpr`, is evaluated. If it evaluates to `#t` (or something that counts as *true*), then the *then* expression, `thenExpr`, is evaluated, and its value is taken to be the value of the entire `if` expression. However, if the condition evaluates to `#f`, then the *else* expression, `elseExpr`, is evaluated, and its value is taken to be the value of the entire `if` expression. Thus, either the *then* expression or the *else* expression is evaluated, but never both.

An `if` expression whose *then* expression is `#t`, and whose *else* expression is `#f` can be simplified. For example, `(if (> x y) #t #f)` is equivalent to `(> x y)`.

A `when` expression is useful in cases where no *else* expression is needed. For example, a `when` expression can be used to generate side-effect printing in certain cases, but not others. Because a `when` expression can evaluate to *void*, it should *not* be used to generate an output value!

### Special Forms Introduced in this Chapter

- |                   |   |
|-------------------|---|
| <code>if</code>   | For making binary decisions                             |
| <code>when</code> | For cases where an <i>else</i> expression is not needed |

# Index

- booleans, 7
- built-in functions
  - \*: multiplication function, 19
  - +: addition function, 19
  - : subtraction function, 19
  - /: division function, 19
  - <: less than, 43
  - <=: less than or equal to, 43
  - =: equal to, 43
  - >: greater than, 43
  - >=: greater than or equal to, 43
  - boolean?: type checker, 41
  - eq?, 44
  - eval: evaluation function, 21
  - integer?, 21
  - list?: type checker, 41
  - load: for loading files, 62
  - null?: type checker, 41
  - number?: type checker, 41
  - printf, 22
  - procedure?: type checker, 41
  - quotient: integer division, 21, 30
  - remainder: integer division, 21, 30
  - string?: type checker, 41
  - symbol?: type checker, 41
  - void, 22
  - void?: type checker, 41, 60
- character sequences, 7, 8
- comments, 63
- computational model, 2, 15, 30
- contracts, 20
  - as comments in a Scheme program, 64
- data, 7
  - non-primitive, 10
  - primitive, 7
- data type, 7
  - booleans, 8
  - functions, 19
  - non-empty lists, 25
  - null, 9
  - numbers, 7
  - strings, 10
  - symbols, 9
  - void, 9, 22
- debug, 3
- Default Rule for evaluating non-empty lists, 26, 29
  - applying functions to inputs, 27, 30
  - generating errors, 29
  - on lists of lists, 27
- denotation
  - single arrow notation ( $\longrightarrow$ ), 8, 16
- destructive functions, 3
- DrScheme, 15
  - interactions window, 15, 16
  - prompt, 16
  - the Run button, 63
- elements
  - of non-empty lists, 25
- empty list, 9
- environments, 12
- escape sequences, 59
- evaluation, 11, 15
  - double arrow notation ( $\implies$ ), 11, 16
  - of booleans, 11
  - of functions, 21
  - of non-empty lists, 26
  - of numbers, 11
  - of special forms, 33
  - of strings, 11
  - of symbols, 12
  - of the empty list, 11
  - of the *void* datum, 11
- evaluation function, 11
  - as a built-in function, 21
- expressions, 7
  - primitive data, 7
  - primitive data expressions, 7
- functions, 2, 11, 47
  - built-in, 9, 19
  - defining vs. applying, 47
  - function application using the Default Rule, 27
- Global Environment, 13, 16
  - and built-in functions, 19
- identity function, 11

- keyword symbols, 33, *see also* special forms
  - redefining, 36
- Lambda Calculus, 2
- lambda function
  - applying to inputs, 52
- lazy evaluation, 67
  - of the `if` special form, 68
- local environment, 52
- local vs. global environments, 54
- non-destructive functions, 3
- non-empty lists, 25
  - syntax and semantics, 25
- non-strict truth values
  - with `if`, 70
- numbers, 7
- predicates, 41
  - type-checker predicates, 41
- program, 7
- sample functions
  - `cone-volume`, 50
  - `convert-to-boolean`, 73
  - `discriminant`, 55
  - `f`, 56
  - `how-big`, 69
  - `rect-area`, 50
  - `sphere-volume`, 52
  - `square`, 49, 54, 57
  - `tester`, 62
    - and `quote`, 62
  - `verbose-func`, 61
- semantics, 1
- side effect
  - vs. output value, 9
- side effect, of the `printf` function, 22
- side effects, 3, 57
  - of `define` special form, 34
  - of `printf`, 30
- special forms, 33
  - `define`, 34
    - semantics of, 34
    - side effect of, 34
    - syntax of, 34
  - `if`, 67
  - `lambda`, 48
    - semantics of, 51
    - syntax of, 50
  - `quote`, 37
    - alternate syntax (`'`), 39
    - semantics of, 37
    - syntax of, 37
  - `when`, 73
  - syntax, 1
  - truth values, 7
  - variables, 12
  - `void datum`, 9, 22, 34