

1

The Single-Person Decision Problem

Imagine yourself in the morning, all dressed up and ready to have breakfast. You might be lucky enough to live in a nice undergraduate dormitory with access to an impressive cafeteria, in which case you have a large variety of foods from which to choose. Or you might be a less-fortunate graduate student, whose studio cupboard offers the dull options of two half-empty cereal boxes. Either way you face the same problem: what should you have for breakfast?

This trivial yet ubiquitous situation is an example of a *decision problem*. Decision problems confront us daily, as individuals and as groups (such as firms and other organizations). Examples include a division manager in a firm choosing whether or not to embark on a new research and development project; a congressional representative deciding whether or not to vote for a bill; an undergraduate student deciding on a major; a baseball pitcher contemplating what kind of pitch to deliver; or a lost group of hikers confused about which direction to take. The list is endless.

Some decision problems are trivial, such as choosing your breakfast. For example, if Apple Jacks and Bran Flakes are the only cereals in your cupboard, and if you hate Bran Flakes (they belong to your roommate), then your decision is obvious: eat the Apple Jacks. In contrast, a manager's choice of whether or not to embark on a risky research and development project or a lawmaker's decision on a bill are more complex decision problems.

This chapter develops a *language* that will be useful in laying out rigorous foundations to support many of the ideas underlying strategic interaction in games. The language will be formal, having the benefit of being able to represent a host of different problems and provide a set of tools that will lend structure to the way in which we think about decision problems. The formalities are a vehicle that will help make ideas precise and clear, yet in no way will they overwhelm our ability and intent to keep the more practical aspect of our problems at the forefront of the analysis.

In developing this formal language, we will be forced to specify a set of assumptions about the behavior of decision makers or players. These assumptions will, at times, seem both acceptable and innocuous. At other times, however, the assumptions will be almost offensive in that they will require a significant leap of faith. Still, as the analysis unfolds, we will see the conclusions that derive from the assumptions

4 • Chapter 1 The Single-Person Decision Problem

that we make, and we will come to appreciate how sensitive the conclusions are to these assumptions.

As with any theoretical framework, the value of our conclusions will be only as good as the sensibility of our assumptions. There is a famous saying in computer science—“garbage in, garbage out”—meaning that if invalid data are entered into a system, the resulting output will also be invalid. Although originally applied to computer software, this statement holds true more generally, being applicable, for example, to decision-making theories like the one developed herein. Hence we will at times challenge our assumptions with facts and question the validity of our analysis. Nevertheless we will argue in favor of the framework developed here as a useful benchmark.

1.1 Actions, Outcomes, and Preferences

Consider the examples described earlier: choosing a breakfast, deciding about a research project, or voting on a bill. These problems all share a similar structure: an individual, or player, faces a situation in which he has to choose one of several alternatives. Each choice will result in some outcome, and the consequences of that outcome will be borne by the player himself (and sometimes other players too).

For the player to approach this problem in an intelligent way, he must be aware of three fundamental features of the problem: What are his possible choices? What is the result of each of those choices? How will each result affect his well-being? Understanding these three aspects of a problem will help the player choose his best action. This simple observation offers us a first working definition that will apply to any *decision problem*:

The Decision Problem A **decision problem** consists of three features:

1. **Actions** are all the alternatives from which the player can choose.
2. **Outcomes** are the possible consequences that can result from any of the actions.
3. **Preferences** describe how the player ranks the set of possible outcomes, from most desired to least desired. The **preference relation** \succsim describes the player's preferences, and the notation $x \succsim y$ means “ x is at least as good as y .”

To make things simple, let's begin with our rather trivial decision problem of choosing between Apple Jacks and Bran Flakes. We can define the set of actions as $A = \{a, b\}$, where a denotes the choice of Apple Jacks and b denotes the choice of Bran Flakes.¹ In this simple example our actions are practically synonymous with the outcomes, yet to make the distinction clear we will denote the set of outcomes by $X = \{x, y\}$, where x denotes eating Apple Jacks (the consequence of *choosing* Apple Jacks) and y denotes eating Bran Flakes.

1. More on the concept of a set and the appropriate notation can be found in Section 19.1 of the mathematical appendix.

1.1.1 Preference Relations

Turning to the less familiar notion of a **preference relation**, imagine that you prefer eating Apple Jacks to Bran Flakes. Then we will write $x \succeq y$, which should be read as “ x is at least as good as y .” If instead you prefer Bran Flakes, then we will write $y \succeq x$, which should be read as “ y is at least as good as x .” Thus our preference relation is just a shorthand way to express the player’s ranking of the possible outcomes.

We follow the common tradition in economics and decision theory by expressing preferences as a “weak” ranking. That is, the statement “ x is at least as good as y ” is consistent with x being *better* than y or *equally as good as* y . To distinguish between these two scenarios we will use the **strict preference relation**, $x \succ y$, for “ x is strictly better than y ,” and the **indifference relation**, $x \sim y$, for “ x and y are equally good.”

It need not be the case that actions are synonymous with outcome, as in the case of choosing your breakfast cereal. For example, imagine that you are in a bar with a drunken friend. Your actions can be to let him drive home or to order him a cab. The outcome of letting him drive is a certain accident (he’s *really* drunk), and the outcome of ordering him a cab is arriving safely at home. Hence for this decision problem your actions are physically different from the outcomes.

In these examples the action set is *finite*, but in some cases one might have infinitely many actions from which to choose. Furthermore there may be infinitely many outcomes that can result from the actions chosen. A simple example can be illustrated by me offering you a two-gallon bottle of water to quench your thirst. You can choose how much to drink and return the remainder to me. In this case your action set can be described as the interval $A = [0, 2]$: you can choose any action a as long as it belongs to the interval $[0, 2]$, which we can write in two ways: $0 \leq a \leq 2$ or $a \in [0, 2]$.² If we equate outcomes with actions in this example then $X = [0, 2]$ as well. Finally it need not be the case that more is better. If you are thirsty then drinking a pint may be better than drinking nothing. However, drinking a gallon may cause you to have a stomachache, and you may therefore prefer a pint to a gallon.

Before proceeding with a useful way to represent a player’s preferences over various outcomes, it is important to stress that we will make two important assumptions about the player’s ability to think through the decision problem.³ First, we require the player to be able to rank *any two outcomes* from the set of outcomes. To put this more formally:

The Completeness Axiom The preference relation \succeq is **complete**: any two outcomes $x, y \in X$ can be ranked by the preference relation, so that either $x \succeq y$ or $y \succeq x$.

At some level the completeness axiom is quite innocuous. If I show you two foods, you should be able to rank them according to how much you like them (including being indifferent if they are equally tasty and nutritious). If I offer you two cars, you should be able to rank them according to how much you enjoy driving them, their safety

2. The notation symbol \in means “belongs to.” Hence “ $x, y \in X$ ” means “elements x and y belong to the set X .” If you are unfamiliar with sets and these kinds of descriptions please refer to Section 19.1 of the mathematical appendix.

3. These assumptions are referred to as “axioms,” following the language used in the seminal book by von Neumann and Morgenstern (1944) that laid many of the foundations for both decision theory and game theory.

6 • Chapter 1 The Single-Person Decision Problem

specifications, and so forth. If I offer you two investment portfolios, you should be able to rank them according to the extent to which you are willing to balance risk and return. In other words, the completeness axiom *does not let you be indecisive between any two outcomes*.⁴

The second assumption we make guarantees that a player can rank *all* of the outcomes. To do this we introduce a rather mild consistency condition called *transitivity*:

The Transitivity Axiom The preference relation \succsim is **transitive**: for any three outcomes $x, y, z \in X$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.

Faced with several outcomes, completeness guarantees that any two can be ranked, and transitivity guarantees that there will be no contradictions in the ranking, which could create an indecisive cycle. To observe a violation of the transitivity axiom, consider a player who strictly prefers Apple Jacks to Bran Flakes, $a \succ b$, Bran Flakes to Cheerios, $b \succ c$, and Cheerios to Apple Jacks, $c \succ a$. When faced with any two boxes of cereal, say $A = \{a, b\}$, he has no problem choosing his preferred cereal a . What happens, however, when he is presented with all three alternatives, $A = \{a, b, c\}$? The poor guy will be unable to decide which of the three to choose, because for any given box of cereal, there is another box that he prefers. Therefore, by requiring that the player have complete and transitive preferences, we basically guarantee that among *any set of outcomes*, he will always have at least one *best outcome* that is as good as or better than any other outcome in that set.

To foreshadow what will be our premise for decision making, a preference relation that is complete and transitive is called a **rational preference relation**. We will be concerned only with players who have such rational preferences, for without such preferences we can offer neither predictive nor prescriptive insights.

Remark As noted by the Marquis de Condorcet in 1785, it is possible to have a group of rational individual players who, when put together to make decisions as a group, will become an “irrational” group. For example, imagine three roommates, called players 1, 2, and 3, who have to choose one box of cereal for their apartment kitchen. Player 1’s preferences are given by $a \succ_1 c \succ_1 b$, player 2’s are given by $c \succ_2 b \succ_2 a$, and player 3’s are given by $b \succ_3 a \succ_3 c$. Imagine that our three players make choices in a democratic way and use majority voting to reach a decision. What will be the resulting preferences of the group, \succ_G ? When faced with the pair a and c , players 1 and 3 will vote for Apple Jacks, hence $a \succ_G c$. When faced with the pair c and b , players 1 and 2 will vote for Cheerios, hence $c \succ_G b$. When faced with the pair a and b , players 2 and 3 will vote for Bran Flakes, hence $b \succ_G a$. As a result, our three rational players will not be able to reach a conclusive decision using the group preferences that result from majority voting! This type of group indecisiveness resulting from majority voting is often referred to as the *Condorcet Paradox*. Because we will not be analyzing group decisions, it is not something we will confront, but it is useful to be mindful of such phenomena, in which imposing individual rationality does not imply “group rationality.”

4. In other words, this axiom prohibits the kind of problem referred to as “Buridan’s ass.” One version describes a situation in which an ass is placed between two identical stacks of hay, assuming that the ass will always go to whichever stack is closer. However, since the stacks are both the same distance from the ass, it will not be able to choose between them and will die of hunger.

1.1.2 Payoff Functions

When we restrict attention to players with rational preferences, not only do we get players who behave in a consistent and appealing way, but as an added bonus we can replace the preference relation with a much friendlier, and more operational, apparatus. Consider the following simple example. Imagine that you open a lemonade stand on your neighborhood corner. You have three possible actions: choose low-quality lemons (l), which imply a cost of \$10 and a revenue from sales of \$15; choose medium-quality lemons (m), which imply a cost of \$15 and a revenue from sales of \$25; or choose high-quality lemons (h), which imply a cost of \$28 and a revenue from sales of \$35. Thus the action set is $A = \{l, m, h\}$, and the outcome set is given by net profits and is $X = \{5, 10, 7\}$, where the action l yields a profit of \$5, the action m yields a profit of \$10, and the action h yields a profit of \$7. Assuming that obtaining higher profits is strictly better, we have $10 \succ 7 \succ 5$. Hence you should choose alternative m and make a profit of \$10.

Notice that we took a rather obvious profit-maximizing problem and fit it into our framework for a decision problem. We derived the preference relation that is consistent with maximizing profit, the objective of any for-profit business. Arguably it would be more natural and probably easier to comprehend the problem if we looked at the actions and their associated profits. In particular we can define the **profit function** in the obvious way: every action $a \in A$ yields a profit $\pi(a)$. Then, instead of considering a preference relation over profit outcomes, we can just look at the profit from each action directly and choose an action that maximizes profits. In other words, we can *use the profit function to evaluate actions and outcomes*.

As this simple example demonstrates, a profit function is a more direct way for a player to rank his actions. The question then is, can we find similar ways to approach decision problems that are not about profits? It turns out that we can do exactly that if we have players with rational preferences, and to do that we define a payoff function.⁵

Definition 1.1 A **payoff function** $u : X \rightarrow \mathbb{R}$ represents the preference relation \succeq if for any pair $x, y \in X$, $u(x) \geq u(y)$ if and only if $x \succeq y$.

To put the definition into words, we say that the preference relation \succeq is represented by the payoff function $u : X \rightarrow \mathbb{R}$ that assigns to each outcome in X a real number, if and only if the function assigns a higher value to higher-ranked outcomes.

It is important to notice that representing preferences with payoff functions is convenient, but that payoff values by themselves have no meaning whatsoever. Payoff is an *ordinal* construct: it is used to order the alternatives from most to least desirable. For example, if I like Apple Jacks more than Bran Flakes, then I can construct the payoff function $u(\cdot)$ so that $u(a) = 5$ and $u(b) = 3$. I can also use a different payoff function $\tilde{u}(\cdot)$ that represents the same preferences as follows: $\tilde{u}(a) = 100$ and $\tilde{u}(b) = -237$. Just as Fahrenheit and Celsius are two different ways to describe hotter and colder temperatures, there are many ways to represent preferences with payoff functions.

Using payoff functions instead of preferences will allow us to operationalize a theory of how decision makers with rational preferences ought to behave, and how they often will behave. They will choose actions that maximize a payoff function that

5. Recall that a function relates each of its inputs to exactly one output. For more on this see Section 19.2 of the mathematical appendix.

8 • Chapter 1 The Single-Person Decision Problem

represents their preferences. One last question we need to ask is whether we know for sure that this method will work: is it true that players will surely have a payoff function representing their preferences? One case is easy and worth going through briefly. In what follows, we provide a formal proposition and a formal, yet fairly easy to follow, proof.

Proposition 1.1 *If the set of outcomes X is finite then any rational preference relation over X can be represented by a payoff function.*

Proof The proof is by construction. Because the preference relation is complete and transitive, we can find a least-preferred outcome $\underline{x} \in X$ such that all other outcomes $y \in X$ are at least as good as \underline{x} , that is, $y \succeq \underline{x}$ for all other $y \in X$. Now define the “worst outcome equivalence set,” denoted X_1 , to include \underline{x} and any other outcome for which the player is indifferent between it and \underline{x} . Then, from the remaining elements of $X \setminus X_1$,⁶ define the “second worst outcome equivalence set,” X_2 , and continue in this fashion until the “best outcome equivalence set,” X_n , is created. Because X is finite and \succeq is rational, such a finite collection of n equivalence sets exists. Now consider n arbitrary values $u_n > u_{n-1} > \dots > u_2 > u_1$, and assign payoffs according to the function defined by: for any $x \in X_k$, $u(x) = u_k$. This payoff function represents \succeq . Hence we have proved that such a function exists. ■

This proposition is useful: for many realistic situations, we can create payoff functions that work in a similar way as profit functions, giving the player a useful tool to see which actions are best and which ought to be avoided. We will not explore this issue further, but payoff representations exist in many other cases that include infinitely many outcomes. The treatment of such cases is beyond the scope of this textbook, but you are welcome to explore one of the many texts that offer a more complete treatment of the topic, which is referred to under the title “representation theorems.” (See, e.g., Kreps [1990a, pp. 18–37, and 1988] for an in-depth treatment of this topic.)

As we have seen so far, the formal structure of a decision problem offers a coherent framework for analysis. For decades, however, teachers, students, and practitioners have instead used the intuitive and graphically simple tool of *decision trees*.

Imagine that, in addition to Apple Jacks (a) and Bran Flakes (b), your breakfast options include a muffin (m) and a scone (s). Your preferences are given as $s > a > m > b$. (Recall that we now consider preferences over outcomes as directly over actions.) Consider the following payoff representation: $v(s) = 4$, $v(a) = 3$, $v(m) = 2$, and $v(b) = 1$. We can write down the corresponding decision tree, which is depicted in Figure 1.1.

To read this simple decision tree, notice that the player resides at the “root” of the tree on the left, and that the tree then branches off, each branch representing a possible action. In the example of choosing breakfast, each action results in a final payoff, and these payoffs are written to correspond to each of the action branches. Our rational decision maker will look down the tree, consider the payoff from each branch, and choose the branch with the highest payoff.

The node at which the player has to make a choice is called a **decision node**. The nodes at the end of the tree where payoffs are attached are called **terminal nodes**. As

6. The notation $A \setminus B$ means “the elements that are in A but are not in B ,” or sometimes “the set A less the set B .”

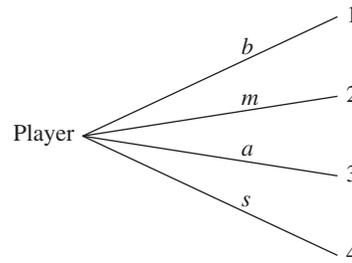


FIGURE 1.1 A simple breakfast decision tree.

the next chapter demonstrates, the structure of a decision tree will become slightly more involved and useful to capture more complex decision problems. We will return to similar trees in Chapter 7, where we consider the strategic interaction between many possible players, which is the main focus of this book.

1.2 The Rational Choice Paradigm

We now introduce *Homo economicus* or “economic man.” *Homo economicus* is “rational” in that he chooses actions that maximize his well-being as defined by his payoff function over the resulting outcomes.⁷ The assumption that the player is rational lies at the foundation of what is known as the **rational choice paradigm**. Rational choice theory asserts that when a decision maker is choosing between potential actions he will be guided by rationality to choose his best action. This can be assumed to be true for individual human behavior, as well as for the behavior of other entities, such as corporations, committees, or nation-states.

It is important to note, however, that by adopting the paradigm of rational choice theory we are imposing some implicit assumptions, which we now make explicit.

Rational Choice Assumptions The player fully understands the decision problem by knowing:

1. all possible actions, A ;
2. all possible outcomes, X ;
3. exactly how each action affects which outcome will materialize; and
4. his rational preferences (payoffs) over outcomes.

Perhaps at a first glance this set of assumptions may seem a bit demanding, and further contemplation may make you feel that it is impossible to satisfy for most decision problems. Still, it is a benchmark for a world in which decision problems are completely understood by the player, in which case he can approach the problems in a systematic and structured way. If we let go of any of these four knowledge

7. A naive application of the *Homo economicus* model assumes that our player knows what is best for his long-term well-being and can be relied upon to always make the right decision for himself. We take this naive approach throughout the book, though we will sometimes question how appropriate this approach is.

requirements then we cannot impose the notion of rational choice. If (1) is unknown then the player may be unaware of his best course of action. If (2) or (3) are unknown then he may not correctly foresee the actual consequences of his actions. Finally if (4) is unknown then he may incorrectly perceive the effect of his choice's consequence on his well-being.

To operationalize this paradigm of rationality we must choose among *actions*, yet we have defined preferences—and payoffs—over *outcomes* and not actions. It would be useful, therefore, if we could define preferences—and payoffs—over actions instead of outcomes. In the simple examples of choosing a cereal or how much water to drink, actions and outcomes were synonymous, yet this need not always be the case. Consider the situation of letting your friend drive drunk, in which the actions and outcomes are not the same. Still each action led to one and only one outcome: letting him drive leads to an accident, and getting him a cab leads to safe arrival. Hence, even though preferences and payoff were defined over outcomes, this *one-to-one correspondence*, or function, between actions and outcomes means that we can consider the preferences and payoffs to be over actions, and we can use this correspondence between actions and outcomes to define the payoff over actions as follows: if $x(a)$ is the outcome resulting from action a , then the payoff from action a is given by $v(a) = u(x(a))$, the payoff from $x(a)$. We will therefore use the notation $v(a)$ to represent the payoff from action a .⁸ Now we can precisely define a rational player as follows:

Definition 1.2 A player facing a decision problem with a payoff function $v(\cdot)$ over actions is rational if he chooses an action $a \in A$ that maximizes his payoff. That is, $a^* \in A$ is chosen if and only if $v(a^*) \geq v(a)$ for all $a \in A$.

We now have a formal definition of *Homo economicus*: a player who has rational preferences and is rational in that he understands all the aspects of his decision problem and always chooses an option that yields him the highest payoff from the set of possible actions.

So far we have seen some simple examples with finite action sets. Consider instead an example with a continuous action space, which requires some calculus. Imagine that you're at a party and are considering engaging in social drinking. Given your physique, you'd prefer some wine, both for taste and for the relaxed feeling it gives you, but too much will make you sick. There is a one-liter bottle of wine, so your action set is $A = [0, 1]$, where $a \in A$ is how much you choose to drink. Your preferences are represented by the following payoff function over actions: $v(a) = 2a - 4a^2$, which is depicted in Figure 1.2. As you can see, some wine is better than no wine (0.1 liter gives you some positive payoff, while drinking nothing gives you zero), but drinking a whole bottle will be worse than not drinking at all ($v(1) = -2$). How much should you drink? Your maximization problem is

$$\max_{a \in [0,1]} 2a - 4a^2.$$

Taking the derivative of this function and equating it to zero to find the solution, we obtain that $2 - 8a = 0$, or $a = 0.25$, which is a bit more than two normal glasses of

8. To be precise, let $x : A \rightarrow X$ be the function that maps actions into outcomes, and let the payoff function over outcomes be $u : X \rightarrow \mathbb{R}$. Define the payoff over actions as the composite function $v = u \circ x : A \rightarrow \mathbb{R}$, where $v(a) = u(x(a))$.

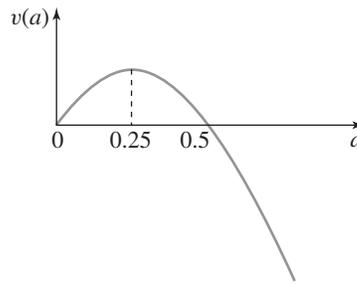


FIGURE 1.2 The payoff from drinking wine.

wine.⁹ Thus, by considering how much wine to drink as a decision problem, you were able to find your optimal action.

1.3 Summary

- A simple decision problem has three components: actions, outcomes, and preferences over outcomes.
- A rational player has complete and transitive preferences over outcomes and hence can always identify a best alternative from among his possible actions. These preferences can be represented by a payoff (or profit) function over outcomes and the corresponding payoffs over actions.
- A rational player chooses the action that gives him the highest possible payoff from the possible set of actions at his disposal. Hence by maximizing his payoff function over his set of alternative actions, a rational player will choose his optimal decision.
- A decision tree is a simple graphic representation for decision problems.

1.4 Exercises

- 1.1 **Your Decision:** Think of a simple decision you face regularly and formalize it as a decision problem, carefully listing the actions and outcomes without the preference relation. Then assign payoffs to the outcomes and draw the decision tree.
- 1.2 **Going to the Movies:** There are two movie theaters in your neighborhood: Cineclass, which is located one mile from your home, and Cineblast, located three miles from your home. Each is showing three films. Cineclass is showing *Casablanca*, *Gone with the Wind*, and *Dr. Strangelove*, while Cineblast is showing *The Matrix*, *Blade Runner*, and *Aliens*. Your problem is to decide which movie to go to.

9. To be precise, we must also make sure that first, the second derivative is negative for the solution $a = 0.25$ to be a local maximum, and second, the value of $v(a)$ is not greater at the two boundaries $a = 0$ and $a = 1$. For more on maximizing the value of a function, see Section 19.3 of the mathematical appendix.

12 • Chapter 1 The Single-Person Decision Problem

- a. Draw a decision tree that represents this problem without assigning payoff values.
 - b. Imagine that you don't care about distance and that your preferences for movies are alphabetic (i.e., you like *Aliens* the most and *The Matrix* the least). Using payoff values 1 through 6 complete the decision tree you drew in part (1). Which option would you choose?
 - c. Now imagine that your car is in the shop and that the cost of walking each mile is equal to one unit of payoff. Update the payoffs in the decision tree. Would your choice change?
- 1.3 **Fruit or Candy:** A banana costs \$0.50 and a piece of candy costs \$0.25 at the local cafeteria. You have \$1.25 in your pocket and you value money. The money-equivalent value (payoff) you get from eating your first banana is \$1.20, and that of each additional banana is half the previous one (the second banana gives you a value of \$0.60, the third \$0.30, and so on). Similarly the payoff you get from eating your first piece of candy is \$0.40, and that of each additional piece is half the previous one (\$0.20, \$0.10, and so on). Your value from eating bananas is not affected by how many pieces of candy you eat and vice versa.
- a. What is the set of possible actions you can take given your budget of \$1.25?
 - b. Draw the decision tree that is associated with this decision problem.
 - c. Should you spend all your money at the cafeteria? Justify your answer with a rational choice argument.
 - d. Now imagine that the price of a piece of candy increases to \$0.30. How many possible actions do you have? Does your answer to (c) change?
- 1.4 **Alcohol Consumption:** Recall the example in which you needed to choose how much to drink. Imagine that your payoff function is given by $\theta a - 4a^2$, where θ is a parameter that depends on your physique. Every person may have a different value of θ , and it is known that in the population (1) the smallest θ is 0.2; (2) the largest θ is 6; and (3) larger people have higher θ s than smaller people.
- a. Can you find an amount that no person should drink?
 - b. How much should you drink if your $\theta = 1$? If $\theta = 4$?
 - c. Show that in general smaller people should drink less than larger people.
 - d. Should any person drink more than one 1-liter bottle of wine?
- 1.5 **Buying a Car:** You plan on buying a used car. You have \$12,000, and you are not eligible for any loans. The prices of available cars on the lot are given as follows:

| Make, model, and year | Price |
|-----------------------|---------|
| Toyota Corolla 2002 | \$9,350 |
| Toyota Camry 2001 | 10,500 |
| Buick LeSabre 2001 | 8,825 |
| Honda Civic 2000 | 9,215 |
| Subaru Impreza 2000 | 9,690 |

For *any given year*, you prefer a Camry to an Impreza, an Impreza to a Corolla, a Corolla to a Civic, and a Civic to a LeSabre. For *any given year*, you are willing to pay up to \$999 to move from any given car to the next preferred one. For example, if the price of a Corolla is $\$z$, then you are willing to buy it rather than a Civic if the Civic costs more than $\$(z - 999)$, but you would prefer to buy the Civic if it costs less than this amount. Similarly you prefer the Civic at $\$z$ to a Corolla that costs more than $\$(z + 1000)$, but you prefer the Corolla if it costs less. For *any given car*, you are willing to move to a model a year older if it is cheaper by at least \$500. For example, if the price of a 2003 Civic is $\$z$, then you are willing to buy it rather than a 2002 Civic if the 2002 Civic costs more than $\$(z - 500)$, but you would prefer to buy the 2002 Civic if it costs less than this amount.

- a. What is your set of possible alternatives?
 - b. What is your preference relation between the alternatives in (a) above?
 - c. Draw a decision tree and assign payoffs to the terminal nodes associated with the possible alternatives. What would you choose?
 - d. Can you draw a decision tree with different payoffs that represents the same problem?
- 1.6 **Fruit Trees:** You have room for up to two fruit-bearing trees in your garden. The fruit trees that can grow in your garden are either apple, orange, or pear. The cost of maintenance is \$100 for an apple tree, \$70 for an orange tree, and \$120 for a pear tree. Your food bill will be reduced by \$130 for each apple tree you plant, by \$145 for each pear tree you plant, and by \$90 for each orange tree you plant. You care only about your total expenditure in making any planting decisions.
- a. What is the set of possible actions and related outcomes?
 - b. What is the payoff of each action/outcome?
 - c. Draw the associated decision tree. What will a rational player choose?
 - d. Now imagine that the reduction in your food bill is half for the second tree of the same kind. (You like variety.) That is, the first apple tree still reduces your food bill by \$130, but if you plant two apple trees your food bill will be reduced by $\$130 + \$65 = \$195$, and similarly for pear and orange trees. What will a rational player choose now?
- 1.7 **City Parks:** A city's mayor has to decide how much money to spend on parks and recreation. City codes restrict this spending to no more than 5% of the budget, and the yearly budget of the city is \$20,000,000. The mayor wants to please his constituents, who have diminishing returns from parks. The money-equivalent benefit from spending $\$c$ on parks is $v(c) = \sqrt{400c} - \frac{1}{80}c$.
- a. What is the action set for the city's mayor?
 - b. How much should the mayor spend?
 - c. The movie *An Inconvenient Truth* has shifted public opinion, and now people are more willing to pay for parks. The new preferences of the people are given by $v(c) = \sqrt{1600c} - \frac{1}{80}c$. What now is the action set for the mayor, and how much spending should he choose to cater to his constituents?