Appendix E
The Theory of Choice

In this appendix, we briefly review the theory of consumer choice. It is provided both as a background and reference on the core concepts of choice theory.

The most widely used theories of choice assume customers are *rational* decision makers who intelligently alter when, what, and how much to purchase to achieve the best possible outcome for themselves. This is a quite plausible assumption. Moreover, an important consequence of this rationality assumption is that customer behavior can be "predicted" by treating each customer as an agent that optimizes over possible choices and outcomes. Optimization theory can then be used to model their behavior. Indeed, for these reasons rational-customer models are the basis of most economic theory.

Yet despite the theoretical and intuitive appeal of the rationality assumption, instances of deviations from rational behavior are observed in experiments and in real life. Alternative theories of choice have emerged to explain such behavior. These models assume customers are not perfectly rational—that there are limits to how cleverly they behave or that they exhibit irrational biases in their choice decisions. These so-called *behavioral theories* are surveyed below as well.

Choice and Preference Relations

Given two alternatives, a *choice* corresponds to an expression of preference for one alternative over another. Here, "alternatives" may refer to different products, different quantities of the same product, bundles of different products or various uncertain outcomes (such as buying a house at the asking price versus waiting and bidding in an auction against other buyers). Similarly, given n alternatives, choice can be defined in terms of the preferences expressed for all pairwise comparisons between the n alternatives.

The mathematical construct that formalizes this notion of choice and preference is a *preference relation.* Customers are assumed to have a set of *binary preferences* over alternatives in a set *X.* That is, given any two alternatives x and y in *X,* customers can rank them and clearly say they prefer one over the other. This is represented by the notation $x \succeq y$. A customer strictly prefers x to y, denoted $x \succ y$, if he prefers x to y, but does not prefer y to x (that is, he is not indifferent between the two alternatives).

Consider a complete set of all such pairwise binary preferences between alternatives in X. The following two properties might be reasonably assumed about "rational" preferences:

- **Asymmetry** If x is strictly preferred to y, then y is not strictly preferred to x.
- **Negative transitivity** If x is not strictly preferred to y and y is not strictly preferred to z, then x is not strictly preferred to z.

Asymmetry and negative transitivity can be considered as "minimal consistency properties" for an expression of preference among a set of alternatives. A binary relation $\succ$ on a set X is called a *preference relation,* if it is asymmetric and negatively transitive. While asymmetry is quite plausible, negative transitivity is not a completely innocuous assumption, as illustrated by the following example:

Example E.1 Suppose you are choosing among jobs in three different cities. Suppose the two factors that matter most to you are income and the climate. The job in city x has a high salary of \$100,000, and the climate is average. The job in city y offers a salary of only \$50,000, but the climate is terrific. The job in city z offers a moderate salary of \$70,000 and the climate is poor. You might not strictly prefer x to y because although x offers a great salary, y offers a great climate. Likewise, you might not strictly prefer y to z because again, while y offers a great climate, z offers a higher salarv. However, you may very well prefer x to z, since x has both a higher salary and a better climate than does z. These preferences would violate negative transitivity.

Despite such shortcomings, the properties of asymmetry and negative transitivity form the classical basis for modeling customer preferences. The following are some examples of preference relations:

Example E.2 (LEXICOGRAPHIC MODEL) This model of preferences, due to Tversky [521], assumes customers rank order various attributes of a product and then evaluate them using a lexicographic rule. For example, a tennis racquet comes in three models *A, B,* and *C* with the following features:

Product	*Wide Body?*	*Graphite?*	*Black?*
A	Yes	No	Yes
B	Yes	Yes	No
C	No	Yes	Yes

The customer's decision rule is to rank all attributes from most important to least important and then eliminate alternatives which do not possess the most important attributes. If more than one alternative remains, the next most important attribute is chosen as a criterion for elimination of alternatives, and so on.

For example, a customer may care most about whether a racquet has a wide body, then whether it is graphite, and lastly whether it is black. He would then prefer racquets with a wide body to all others without a wide body (regardless of the other attributes). Among all those with wide bodies, he would then select those that have graphite construction; among the remaining, he may select only the ones that are black, and so on. So for our three products above, this customer would prefer product in the order B, A, C. One can verify that the lexicographic model generates a preference relation among the alternatives.

Example E.3 (ADDRESS MODEL) Address models link attributes to preference without imposing the restriction that some attributes strictly dominate others as in the lexicographic model. Suppose we have n alternatives and each alternative has m attributes that take on real values. Alternatives can then be represented as n points, $z_1, \ldots, z_n$, in $\Re^m$, which is called *attribute space.* For example, in a travel context attributes may include departure time, arrival time, and price.

Each customer has an ideal point ("address") $y \in \Re^m$, reflecting his most preferred combination of attributes (such as an ideal departure time, arrival time, and price). A customer is then assumed to prefer the product closest to his ideal point in attribute space, where distance is defined by a metric ρ on $\Re^m \times \Re^m$ (such as Euclidean distance). These distances define a preference relation, in which $z_i \succ z_j$ if and only if $\rho(z_i, y) < \rho(z_j, y)$; that is, if z_i is "closer" to the ideal point y of the customer.

Utility Functions

Preference relations are intimately related to the existence of utility functions. Indeed, we have the following theorem (See Kreps [313] for a proof.):

THEOREM E.3 *If X is a finite set, a binary relation $\succ$ is a preference relation if and only if there exists a function $u : X \to \Re$ (called a **utility function**), such that*

$$x \succ y \quad \textit{iff} \quad u(x) > u(y).$$

Intuitively, this theorem follows because if a consumer has a preference relation, then all products can be ranked (totally ordered) by his preferences; a utility function then simply assigns a numerical value corresponding to this ranking. Intuitively, one can think of utility as a measure of "value," though in a strict sense its numerical value need not correspond to any such tangible measure. Theorem E.3 applies to continuous sets X (such as travel times or continuous amounts of money) as well under mild regularity conditions, in which case the utility function $u(\cdot)$ is then continuous. The following examples illustrate the construct of utility:

Example E.4 A utility function corresponding to the lexicographic model of Example E.2 can be constructed as follows: Suppose there are n alternatives with m attributes each. Let the attributes be ordered so that 1 represents the highest-valued attribute and m the lowest. Let $a_k(x), k = 1, \ldots, m,$ be binary digits representing whether alternative x, possesses attribute k. Then a utility satisfying Theorem E.3 is the binary number,

$$u(x) = a_{x1} a_{x2} \cdots a_{xm}.$$

Maximizing over these utilities leads to the same customer decisions as the lexicographic model.

Example E.5 Consider the address model of Example E.3. Again, Theorem E.3 guarantees that an equivalent utility maximization model exits that generates the same choices. In this case, it is easy to see that for customer y the continuous utilities

$$u(z) = c - \rho(z, y),$$

where c is an arbitrary constant, produce the same decision rule as the address model.

Utility for Money and Consumer Budgets

It is often convenient to narrow the choice of utilities further and express utility in monetary terms. To do so, one can pose the question: Given the customer's preference for n goods (purchase alternatives), a vector of market prices $\mathbf{p} = (p_1, \ldots, p_n)$ for these goods, and a level of monetary wealth w, how would a customer "spend" his wealth? To make matters simpler, we assume quantities x_i of each good i are continuous and our customer has a continuous utility function $u(\mathbf{x})$. Let $\mathbf{x} = (x_1, \ldots, x_n)$. The *consumer budget problem* can then be formulated as[1]

$$\begin{aligned} v(w) = \quad & \max \quad u(\mathbf{x}) \\ & \text{s.t.} \quad \mathbf{p}^{\top}\mathbf{x} \leq w \\ & \qquad\quad \mathbf{x} \geq 0. \end{aligned} \tag{E.1}$$

In other words, customers purchase quantities x_i of each good i to maximize their total utility subject to the constraint that they can spend at most their total wealth w. The optimal solution gives the customer's utility for wealth (or money) $v(w)$; the optimal solution, $\mathbf{x}^*$, gives the customer's *demand* for each of the n goods.

Utility for money is increasing in w since one can always "not spend" the wealth w. Also, since the utility for money depends on the prices of goods, if prices change, both the demand x^* and the utility for money may change. The *marginal utility of money* $u'(w)$ also depends on the customer's wealth w. The utility for money $v(w)$ is concave if $u(x)$ is concave,[2] in which case the consumer has decreasing marginal utility for money. Intuitively, this is because at low levels of wealth only highly essential goods are purchased (food, water, clothing, shelter)—all of which have very high utility to most of us. As wealth rises, each marginal dollar is allocated to somewhat less important purchases.

If the function $u(\mathbf{x})$ is continuously differentiable and we let π denote the optimal Lagrange multiplier on the budget constraint in (E.1), then the marginal value of money is

$$v'(w) = \pi.$$

We can use this fact to redefine utilities in monetary terms. Indeed, since our customer's monetary utility for an additional dollar should be one dollar, we should have $v'(w) = \pi = 1$ if utilities are measured in dollars. This change of units can be accomplished by rescaling the customer's utility functions by $v'(w) = \pi$ to form the modified utilities

$$\tilde{u}(\mathbf{x}) = \frac{u(\mathbf{x})}{\pi}. \tag{E.2}$$

[1]Dynamic versions of this consumer budget problem can also be formulated by allowing customers to purchase over multiple periods and invest money at a given interest rate for future consumption. Other variations introduce wages and a utility for leisure time and allow customers to increase their monetary wealth by varying their time allocated to labor, and so on.

[2]This follows easily from the convexity of the budget constraint and the fact that (E.1) is a maximization problem. Concavity of the utility function corresponds to having decreasing marginal utility of consumption for goods, which is a natural assumption.

Reservation Prices

A *reservation price* is the monetary amount a consumer is willing to give up to acquire an extra marginal unit of some good. Reservation prices are also referred to as the customer's *willingness to pay.* Formally, if x^* denotes the optimal solution to (E.1), the reservation price, denoted v_i, for an additional unit of good i is given by

$$v_i \equiv \frac{\partial \tilde{u}(\mathbf{x}^*)}{\partial x_i}, \tag{E.3}$$

where $\tilde{u}(\mathbf{x}^*)$ is the monetary utility (E.2). The first-order conditions of the budget problem imply $\frac{\partial \tilde{u}(\mathbf{x}^*)}{\partial x_i} = p_i$ since $\tilde{u}'(w) = \pi = 1$ when utilities are measured in dollars. Combining this with (E.3) implies that $v_i = p_i$. Thus, a customer's reservation price for goods that are *currently consumed* is simply the current market price. The reason for this equivalence, intuitively, is that if our customer valued another unit of good i at strictly more than its market price, then he would be able to increase his utility by reducing consumption of other goods and increasing his consumption of good i. Since our customer is assumed to be maximizing utility, this cannot occur.

On the other hand, for goods i that are not being consumed, so $x_i^* = 0$, the first-order conditions to (E.1) imply $\frac{\partial \tilde{u}(\mathbf{x}^*)}{\partial x_i} < p_i$, or equivalently $v_i < p_i$. In other words, by (E.3) the customer's reservation price for the first unit of good i is strictly less than its current market price. Moreover, the customer would change only his allocation and buy good i if its price p_i dropped below his reservation price v_i.

This formal analysis of reservation price is arguably less important in practice than the informal concept—namely, that the reservation price is the maximum amount a customer is willing to pay for an additional unit of good i. And to entice a customer to buy good i, the price must drop below his reservation price. Still, the analysis highlights the important fact that reservation prices are not "absolute" quantities. Like utility for money, they depend on customers' preferences, wealth, their current consumption levels, and the prices of other goods the customers may buy; change one of these factors, and customers' reservation price may change.

Lotteries and Stochastic Outcomes

Many choices in life involve uncertain outcomes, such as buying insurance, making investments or eating at a new restaurant. How do customers respond to these sorts of uncertainties? The theory of choice under uncertainty is a deep and extensive topic. Here, we outline the basic ideas and highlight the main concepts.

Consider again a discrete, finite set of n alternatives, $X = \{x_1, \ldots, x_n\}$. Let $\mathcal{P}$ be the class of all probability distributions $P(\cdot)$ defined on X. That is, $P \in \mathcal{P}$ is a function satisfying $\sum_i P(x_i) = 1$ and $P(x_i) \geq 0$ for $i = 1, \ldots, n$. One can think of each P as a "lottery," the outcome of which is that the customer is left with one of the alternatives x_i according to the distribution P.

What can we say about the customer's preference for these various lotteries? Specifically, when can we say that for any two lotteries P_1 and P_2, customers "prefer" one over the other (denoted by $P_1 \succ P_2$)?

To answer this question we again need to make some assumptions on customer preferences. First, we will assume there exists a preference relation $\succ$ on the n different outcomes x_i as before. Second, for any two lotteries P_1 and P_2, consider a compound lottery parameterized by α as follows:

STEP 1: A coin is flipped with probability of heads equal to α.

STEP 2: If the coin comes up heads, the customer enters lottery P_1; otherwise, the customer enters lottery P_2.

Denote this compound lottery by $\alpha P_1 + (1-\alpha)P_2$. Note this compound lottery is also contained in the set $\mathcal{P}$ (i.e., $\mathcal{P}$ is a convex set). We then require the following consistency properties on a customers preference for lotteries:

- **Substitution axiom** For all P_1, P_2, and P_3 in $\mathcal{P}$ and all $\alpha \in (0,1]$, if $P_1 \succ P_2$, then $\alpha P_1 + (1-\alpha)P_3 \succ \alpha P_2 + (1-\alpha)P_3$.
- **Continuity axiom** For all P_1, P_2, and P_3 in $\mathcal{P}$ with $P_1 \succ P_2 \succ P_3$, there exist values $\alpha \in (0,1)$ and $\beta \in (0,1)$ such that $\alpha P_1 + (1-\alpha)P_3 \succ P_2 \succ \beta P_1 + (1-\beta)P_3$.

Roughly, the first axiom says that if one gamble produces strictly preferred outcomes for any realization of uncertainty, then the customer should strictly prefer it. The second axiom says that if a customer strictly prefers one gamble to another, then he should be willing to accept a sufficiently small risk of an even worse outcome to take the preferred gamble. Both are reasonable assumptions.

Under these two axioms, there exist utilities on outcomes such that the expected utility of each lottery defines a customer's preference relation among lotteries. Specifically,

THEOREM E.4 *A preference relation on the lotteries $\mathcal{P}$ exists that satisfies the substitution and continuity axioms if and only if there exists a utility function $u(\cdot)$ such that $P_1 \succ P_2$ if and only if*

$$\sum_{i=1}^{n} u(x_i)P_1(x_i) > \sum_{i=1}^{n} u(x_i)P_2(x_i).$$

That is, if and only if the expected utility from lottery P_1 exceeds the expected utility of lottery P_2. In addition, any two utility functions u and u' satisfying the above must be affine transformations of each other; that is,

$$u(x) = cu'(x) + d,$$

for some real $c > 0$ and d.

This result is due to von Neumann and Morgenstern [541] and is known as the *von Neumann-Morgenstern expected-utility theory.* Essentially, it allows us to extend utility as a model of customer preference to the case of uncertain outcomes, with expected utility replacing deterministic utility as the criterion for customer decision making. Since the original deterministic outcomes (e.g., outcome x_i occurs with probability $P(x_i) = 1$) are included in $\mathcal{P}$, the von Neumann-Morgenstern expected utilities also help us "narrow down" the list of possible utility functions for the customer.

Risk Preferences

An important special case of expected-utility theory is when outcomes represent different monetary amounts, so alternatives correspond to different levels of wealth and lotteries correspond to different gambles on a customer's ending wealth level. For this discussion, we assume the wealth levels are continuous and that the customer has preferences for wealth that satisfy the conditions of Theorem E.4. Also, assume the lotteries are now continuous distributions F on $\Re$.[3]

Consider now any given lottery F (a distribution on possible wealth outcomes) and μ_F denote the mean of the distribution. A customer is said to have *risk-averse preferences* if he prefers the certain wealth μ_F to the lottery F itself for all possible lotteries F. That is, the customer always prefers the certainty of receiving the expected wealth rather than a gamble with the same mean. The customer is said to have *risk-seeking preferences* if he prefers the gamble F to the certain outcome μ_F for all F. Finally, he has *risk-neutral preferences* if he is indifferent between the lottery F and the certain reward μ_F.[4] We then have the following result:

THEOREM E.5 *A customer's preference $\succ$ for lotteries exhibits risk-aversion (risk-seeking) behavior if and only if their von Neumann-Morgenstern utility function $u(w)$ is concave (convex). Their preference is risk-neutral if and only if $u(w)$ is affine.*

Thus, risk preferences are linked directly to concavity or convexity of the customer's utility function. The reason is quite intuitive; with a concave utility function for wealth, a customer gains less utility from a given increase in wealth than he loses in utility from the same decrease in wealth. Hence, the upside gains produced by the volatility in outcomes do not offset the downside losses, and customers therefore prefer the certain average to the uncertain outcomes of the lottery. Since most customers have a decreasing marginal utility for wealth, risk aversion is a good assumption in modeling customer behavior.

Still, the concept of risk aversion has to be addressed with care in operational modeling. While it is true that most customers are risk-averse when it comes to *large* swings in their wealth, often the gambles we face as consumers have a relatively small range of possible outcomes relative to our wealth. For example, a customer may face a price risk in buying a CD or book online. However, the differences in prices for such items are extremely small compared to his total wealth. In such cases, the utility function is "almost linear" in the range of outcomes affecting the decision and the customer tends to behave "as if" he were risk-neutral.[5] Similar statements apply to firms. Generally, they are risk-averse too, but for decisions and gambles that involve "small" outcomes relative to their total wealth and income, they tend to be approximately risk-neutral. Hence, risk-neutrality is a reasonable assumption in operational models and, indeed, is the standard assumption in RM practice.

[3]The extension of Theorem E.4 to the continuous case requires some additional technical conditions that are beyond the scope of this chapter. See Kreps [313].

[4]Note that a customer's preferences may not fall into any of these three categories. For example, many consumers take out fire insurance, preferring a certain loss in premium payments every year to the gamble between making no payments but potentially loosing their house, yet simultaneously play their local state lottery, which has an expected loss but provides a small probability of a large wealth pay-off. Such behavior violates a strict risk preference.

[5]Formally, one can see this by taking a Taylor series approximation of the utility function about the customer's current wealth to; the first-order approximation is affine, corresponding to risk-neutrality.

Information Asymmetry

Another important fact related to customer choice is that normally much of a customer's information is private, information that only the customer "knows" and information that cannot be directly observed by a firm. Normally, both a customer's preferences and wealth are private information. One can perhaps gain clues to a customer's preference by observing their purchase behavior over time (their so-called *revealed preferences*), and partial information on their wealth may be garnered from surveys and transactional data. But in general, much of the data affecting customers' choice behavior remains hidden.

This "information asymmetry" between customers and firms has implications for pricing and RM as discussed in detail in Chapters 6 and 8. To give a quick sense of the effect, consider how customers react to a posted price. Due to information asymmetry, the selling firm rarely knows a customer's true reservation price for their product. If they did, they could potentially offer the customer a price only marginally less than their reservation price and maximize the revenue obtained from each customer. Instead, most firms have to guess at each customer's reservation price. As a result, sometimes they price too high, and the customer does not purchase at all; other times they price too low, and although the customer may decide to purchase, they lose an opportunity for a revenue gain as the customer would have been willing to pay more. In this way, the private information of customers often allows them to retain some surplus, even from a monopoly seller.

Deviations from Rational Behavior

While rational behavior is the standard assumption underlying most of the theory and practice of RM, it is far from being completely accepted as a model of how an actual customer behaves. Indeed, much of the recent work in economics and customer behavior has centered on explaining observed, systematic deviations from rationality on the part of customers.

The seminal work in this area is that of Kahneman and Tversky [278, 277], who showed that customers often exhibit consistent biases when faced with simple choices in an experimental setting. Their key insight is that most individuals tend to evaluate choice in terms of losses and gains from their status quo wealth, rather than evaluating choices in terms of their terminal wealth as in classical utility theory. People also show a tendency toward "loss aversion" rather than risk aversion, and they have a strong preference for certainty of outcomes when evaluating choices. Finally, how gains and losses are expressed matter as well.

They showed that how questions of choice are "framed" have a large impact on customer choice. When choices are framed in terms of gains versus losses, customers typically care more about avoiding losses than about making gains. This is true even if the "gains" and "losses" amount to exactly the same choice. For example, if a public health policy choice is framed as a gain (200 of 800 diseased people will be saved) or as a loss (600 of 800 diseased people will die), most people respond differently, even though the outcomes are identical.

Other experiments revealed that people put a much higher value on a product they already own than one that they don't own because giving up a product they have feels like a loss. This behavior is part of the rationale behind the common marketing strategy of offering products on a "free 30-day trial"–that customers are much more willing to pay to "avoid losing" the trial product than they are willing to pay to acquire that same product initially. (Of course, other simpler explanations—such as reassuring the customer of the quality of the product—can also explain such guarantees.)

Another bias people exhibit is due to what is called *mental accounting,* in which customers tend to evaluate gains and losses for different categories of goods differently because they have "mental budgets" for each category of goods. For example, suppose you purchase a $1,000 watch and then immediately lose it. You might then be reluctant to replace it because in some sense your "budget" for purchasing watches has been exhausted. However, suppose you lost $1,000 in the stock market and you did not own a watch at the time. Then you might be willing to buy a new $1,000 watch because there is no direct association between the $1,000 dollar loss and the amount you might have "mentally allocated" to spend on a watch (for example, you might account for this as "an investment loss" not a "expensive-watch loss"). Such heuristic accounting again violates the rationality assumptions of classical consumer behavior.

Kahneman and Tversky [278] developed what they termed *prospect theory* to explain such effects. Prospect theory differs from expected-utility theory in several respects. For one, it handles the probabilities of outcomes differently, treating them as "decision weights" that may or may not correspond to actual probabilities. Indeed, prospect theory postulates that the subjective decision weights used by most customers tend to overweigh small probabilities and underweigh high probabilities. Prospect theory also uses the notion of "value" rather than "utility," where value is defined in terms of deviations from a reference point (the customer's status quo wealth). They postulate an S-shaped curve for the value function, which is convex for losses below the reference point and concave for gains above the reference point. Using this construct, Kahneman and Tversky [278] are able to model and explain many observed deviations from rational behavior.

Do such findings mean that expected utility theory is "dead"? Not really. In a gross sense, people do tend to behave in accordance with rationality assumptions. However, what this behavioral theory shows quite clearly is that the axioms of rational behavior, plausible as they are, do not apply uniformly and that there are situations in which deviations from rational behavior are systematic and substantial.

The main consequence of these findings for RM practice is that one should always understand the "environment" in which choices are made; the details of the buying situation matter in terms of customers' responses. How prices are presented, what "reference point" the customer perceives, the framing of the choice decision, their sense of "ownership" over the product—all can potentially influence their responses. While many of the tactics used to influence these factors lie in the domain of general marketing and are beyond the scope of this book, the general message that the choice environment matters is nevertheless an important one for RM practitioners and researchers to heed. Indeed, we expect these behavioral theories of demand to influence RM practice more directly in the years ahead.