

Neuroeconomics: Formal Models of Decision Making and Cognitive Neuroscience

Aldo Rustichini

OUTLINE

Introduction	33	The Computation of Utility	40
Axiomatic Decision Theory	34	<i>A Synthesis</i>	40
<i>The Method of Revealed Preferences</i>	34	Factors Affecting the Decision Process	40
Axioms	34	<i>A Simple Example</i>	41
Representation of Preferences	35	<i>Quality of the Signal and Response Time</i>	41
Cardinal and Ordinal Utilities	35	Cognitive Abilities and Preferences	42
Static Stochastic Choice	36	Appendix: Random Walk with Endogenous Barriers	43
<i>Economic Theories of Static Stochastic Choice</i>	36	Optimal Policy	44
Random Utility Models	37	Value and Quality of Signals	44
Stochastic Choice Models	37	References	45
Dynamic Stochastic Choice	38		
<i>The Random Walk Model</i>	38		
Decision in Perceptual Tasks	39		
<i>Formal Model</i>	39		
<i>Decision in Economic Choices</i>	39		

INTRODUCTION

We provide here a link between the formal theory of decision making and the analysis of the decision process as developed in neuroscience, with the final purpose of showing how this joint analysis can provide explanation of important elements of economic behavior.

The methodological standpoint we present is that experimental economics, including neuroeconomics, establishes relationships among variables, some inferred from observed behavior. In particular, a fundamental component of the neuroeconomics project is to establish connections between variables derived from observed behavior and psycho-physiological quantities. For example, the derived variables can be

utility, or parameters like risk aversion. If a researcher claims that the relationship between utility or value and a psycho-physiological quantity (like firing rate of a neuron) is linear, then one has to be sure that the derived variable is uniquely defined, up to linear transformations. If it is not, the statement is meaningless.

We first review the basic concepts and results in decision theory, focusing in particular on the issue of cardinal and ordinal utility, remembering that within the von Neumann-Morgenstern framework the utility function on lotteries (the only observable object) is only defined up to monotonic transformations.

We then show how, under well-specified assumptions, it is possible to identify a unique ordinal object, and how this is based on stochastic choice models. These are, however-static models, so they do not give an account of how the choice is reached. We show that the static models have a dynamic formulation, which extends the static one.

Once the main features of the decision process have been established, we can show how they explain important features of the choice – even some that had been ignored so far. For example, we show how risk aversion, impatience, and cognitive abilities are related.

AXIOMATIC DECISION THEORY

In economic analysis, decision theory is developed with a purely axiomatic method. The theory proceeds by first defining a set of choices that a subject (the decision maker, DM) faces. A choice is a finite set of options that are offered to the DM; a decision is the selection of one of these options. The observed data are pairs of choices offered and decisions taken: it is possible to collect these data experimentally asking a real DM to pick one out of two options, under the condition that the object selected is actually delivered to her.

The method and the main results of the theory are best illustrated in a simple and concrete example of choice environment, choice under risk. In this environment, the options are lotteries. A common lottery ticket provides an example of the abstract concept of lottery: a winning number is drawn at random, and with such a ticket, a person is entitled to a payment if the winning number is the one she has, and she receives no payment otherwise. In general, a lottery is a contract specifying a set of outcomes (the payments made to the subject in our example) and a probability for each of these outcomes. The probability is specified in advance and known to the subject, so in this model

there is only objective uncertainty, as opposed to the subjective uncertainty analyzed in [Savage \(1954\)](#) and [Anscombe and Aumann \(1963\)](#).

A lottery with two outcomes can be formally described with a vector $(x, p, y, 1 - p)$, to be interpreted as: this lottery gives the outcome x with probability p , and the outcome y with probability $1 - p$. For example, the lottery $(\$10, 1/2, \$0, 1/2)$ where outcomes are monetary payments gives a 50-50 chance of a payment of \$10, and nothing otherwise. Lotteries do not need to be a monetary amount, but for simplicity of exposition we confine ourselves to this case.

The Method of Revealed Preferences

We can observe the decisions made by our subject, while we do not observe her preferences directly. However, we may interpret her choices as a “revelation” that she makes of her preferences. Suppose that when she is presented with a choice between lottery L_1 and L_2 she chooses L_1 : we may say that she reveals she prefers L_1 to L_2 . Within economic analysis, it is in this sense, and in this sense only, that we can say that the DM prefers something. The two descriptions of her behavior, one with the language of decisions and the other with that of preferences, are, by the definition we adopt, perfectly equivalent. Since the language of preferences seems more intuitive, it is the one used typically by decision theory, and is the one used here. But how do we describe the behavior, or preferences, of our subject?

Axioms

Even with simple lotteries with two monetary outcomes, by varying the amounts and the probabilities we can obtain an infinite set of possible lotteries, and by taking all the possible pairs of these lotteries we can obtain infinitely many choices. To describe the behavior of a subject completely, we should in principle list the infinite set of decisions she makes. To be manageable, a theory needs to consider instead subjects whose decisions can be described by a short list of simple principles, or axioms.

The first axiom requires that the preferences are complete: for every choice between the two lotteries L_1 and L_2 , either L_1 is preferred to L_2 , or L_2 is preferred to L_1 . The occurrence of both possibilities is not excluded: in this case, the subject is indifferent between the two lotteries. When the subject prefers L_1 to L_2 , but does not prefer L_2 to L_1 , then we say that she strictly prefers L_1 to L_2 . The second axiom requires

the preferences to be transitive: if the DM prefers L_1 to L_2 and L_2 to L_3 , then she prefers L_1 to L_3 . We define the preference order \succeq by writing $L_1 \succeq L_2$ when decision maker prefers L_1 to L_2 and we write $L_1 \succ L_2$ when decision maker strictly prefers L_1 to L_2 . Formally:

Axiom 1 (Completeness and transitivity)

For all lotteries L_1, L_2 and L_3

1. Either $L_1 \succeq L_2$ or $L_2 \succeq L_1$
2. If $L_1 \succeq L_2$ and $L_2 \succeq L_3$ then $L_1 \succeq L_3$.

The next two axioms are also simple, but more of a technical nature. Suppose we have two lotteries, $L_1 = (x, p, y, 1 - p)$ and $L_2 = (z, q, w, 1 - q)$. Take any number r between 0 and 1. Imagine the following contract. We will run a random device, with two outcomes, *Black* and *White*, the first with probability r . If *Black* is drawn, then you will get the outcome of the lottery L_1 ; if *White* is drawn you will get the outcome of the lottery L_2 . This new contract is a compound lottery. If you do not care about how you get the amounts of money, then this is the lottery with four outcomes described as $(x, rp, y, r(1 - p), z, (1 - r)q, w, (1 - r)(1 - q))$. We write this new lottery as $rL_1 + (1 - r)L_2$.

The next axiom requires that if you strictly prefer L_1 to L_2 , then for some number r , you strictly prefer $rL_1 + (1 - r)L_2$ to L_2 . This seems reasonable: when r is close to 1 the composite lottery is very close to L_1 , so you should strictly prefer it to L_2 just like you strictly prefer L_1 .

Axiom 2 (Archimedean continuity)

If $L_1 \succ L_2$ then for some number $r \in (0, 1)$,

$$rL_1 + (1 - r)L_2 \succ L_2.$$

Finally, suppose that you strictly prefer L_1 to L_2 . Then for any lottery L_3 , you also strictly prefer $rL_1 + (1 - r)L_3$ to $rL_2 + (1 - r)L_3$. Again, this seems reasonable. When, in the description we gave above, *White* is drawn, then in both cases you get L_3 ; when *Black* is drawn, in the first case you get L_1 and in the second L_2 . Overall, you should prefer the first lottery $rL_1 + (1 - r)L_3$.

Axiom 3 (Independence)

If $L_1 \succ L_2$ then for any number $r \in (0, 1)$ and any lottery L_3 ,

$$rL_1 + (1 - r)L_3 \succ rL_2 + (1 - r)L_3.$$

Representation of Preferences

A fundamental result in decision theory (due to [von Neumann and Morgenstern, \(vNM\), 1947](#)) is that subjects having preferences that satisfy these axioms

(completeness, transitivity, Archimedean continuity and independence) behave as if they had a simple numerical representation of their preferences – that is, a function that associates with a lottery a single number, called the utility of the lottery, that we can write as $U(L)$. This function is called a representation of the preferences if whenever L_1 is preferred to L_2 , then the utility of L_1 is larger than the utility of L_2 , that is $U(L_1) > U(L_2)$. (Note that here we use $>$ not \succ because $U(L_1)$ is a numerical property not a preference.)

The vNM theorem also states that the preference order satisfies the axioms above if, and only if, the numerical representation has a very simple form, equal to the expectation of the utility of each outcome, according to some function u of outcomes. For example, the expected utility of the lottery $L = (x, p, y, 1 - p)$ is:

$$U(L) = pu(x) + (1 - p)u(y) \quad (4.1)$$

Cardinal and Ordinal Utilities

For neuroeconomics, and any research program that tries to determine how decisions are implemented, the utility function is the most interesting object. This function ties observed behavior with a simple one-dimensional quantity, the utility of the option, and predicts that the decision between two options is taken by selecting the option with the highest utility. However, if we are interested in determining the neural correspondents of the objects we have introduced, we must first know whether these objects are unique. For example, we may formulate the hypothesis that the decision is taken depending on some statistics of the firing rate of a group of neurons associated with each of the options. We may also consider that this firing rate is proportional to the utility we determine from observed choice behavior. Then we need to know whether this utility is uniquely determined. This introduces us to a fundamental distinction in decision theory, between cardinal and ordinal representation.

An ordinal representation of a preference is any utility function such that $U(L_1) > U(L_2)$ if, and only if, L_1 is strictly preferred to L_2 . There are clearly many such functions. For example, if M is any strictly increasing function, then also $M(U(L_1)) > M(U(L_2))$ if, and only if, L_1 is strictly preferred to L_2 . So we say that an ordinal representation is only unique up to increasing (or monotonic) transformation, like the one we have used from U to $M(U)$. Consider now the function u is [equation \(4.1\)](#), and take two numbers $a > 0$ and b . Replace the u function in (4.1) with the new function v defined for any value z by $v(z) = au(z) + b$. If we replace the u in [equation \(4.1\)](#) we obtain a new function on lotteries

which also represents preferences, and has the form of expected utility. Since these transformations leave the observed choices and preferences unchanged, the u in (4.1) is not unique.

However, these are the only transformations we can apply. A second remarkable part of the vNM theorem is that if two functions u and v represent the preferences of a subject as expected utility (that is, as in (4.1)), then it must be that $v(z) = au(z) + b$ for some positive number a and some number b . In this case the two functions are said to be linear transformations of each other, and representations like these are called cardinal representations. A different but equivalent way of saying this is that if we consider functions on a range of monetary prizes between a minimum of 0, say, and a maximum value M , and we agree to normalize the utility function u to $u(0) = 0$ and $u(M) = 1$, then there is a unique such function that, once substituted in equation (4.1), represents the preferences of the DM.

However, the observed decision between two choices is determined by the function U , and this is only unique up to *monotonic* (not necessarily linear) transformations. So even if we agree to normalize $U(0) = 0$ and $U(M) = 1$, there are still infinitely many such U s. If we are looking for a neural basis of choice, then the only sensible statements that involve the function U are those that remain true if we take monotonic transformations of that function. For example, statements like “the firing rate is a linear transform of the U ” are meaningless.

Can we do better than this? We can, if we agree to extend the set of observed data to include errors and time in the decision process. This will take us to the next topic of stochastic choices, and one step closer to the models of decision currently applied in neuroscience.

STATIC STOCHASTIC CHOICE

To illustrate and motivate this new point of view, we begin with a finding discovered in the 1940s by an Iowa researcher, D. Cartwright (Cartwright, 1941a, 1941b; Cartwright and Festinger, 1943). He asked subjects to pick one of two alternatives. By changing the parameter appropriately, the experimenter could make the choice more or less difficult – for example, setting the width of two angles closer would make the task of choosing the wider angle between the two a more difficult task. Also, by asking the subject to make the same choice repeatedly, at some distance in time, he could test the frequency of the choice of one

or the other of the alternatives in different decision problems. He could now construct what we can call the empirical random choice: for every set of options, the frequency of choice of each option out of that set.

He also measured the response time for each choice and then plotted the average response time for each decision problem against the minimum frequency of any of the two choices in that same problem. The key finding was that the longest response time was observed when the minimum frequency was approaching 50%; the problems in which the subject was more likely to select, in different trials, both options were also those in which she was taking more time to decide. A related result is the “symbolic distance” effect, first stated in Moyer and Landauer (1967).

The finding of Cartwright suggests a model of decision where two opposing forces push in the direction of each of the options. When the difference between these two forces is large, the decision is frequently in favor of the favored option, and the decision is taken quickly. When they are the same, the frequency of choice of the two options becomes closer, and the response time becomes longer.

For our purposes of outlining a theory of the decision process when the decision is among economic choices, it is important to note that for economic choices the same result holds. Suppose we determine the utility of a subject from the observed choices, that is, the quantity $U(L)$ for every lottery L . We can now measure the distance between the utility of any two lotteries in a choice, and conjecture that the analogue of the Cartwright results holds in this situation: the closer the two options in utility, the longer the time to decide, and the higher the minimum probability of choosing any of the two. This conjecture has been confirmed in several studies. There is one problem, however: what is the distance between the utilities? If the utility is unique up to monotonic affine transformations, then the distance is well defined up to re-scaling by a single number. But we have just seen that the U in (4.1) is not unique up to monotonic affine transformations, thus even after normalization we have infinitely many such functions. So how can we measure in a meaningful way the distance in utility between two options? The key to a solution is in the inconsistency of choice that we have just reported.

Economic Theories of Static Stochastic Choice

The experimental evidence reviewed in the previous section suggests that when repeatedly faced with a choice between the same two options, the subject may not always choose the same option in

each instance. In contrast, the utility theory we have reviewed so far predicts that if the utility of one of the two is larger, that should always be the chosen one. The key idea of the stochastic theory of choice is that the relative frequency of the choice of one option over the other gives a better measure of the utility of the two options.

There are two classes of models of stochastic choice in economic theory. Both address the following problem. Suppose that a DM is offered, in every period, the choice of a set of lotteries, a menu. We observe her choices over many periods. For a given menu, the choices may be different in different periods, but we can associate for every menu the frequency of choices over that menu – that is, a probability distribution over the set. Both classes of models want to determine the underlying preference structure that produces this observed frequency.

Let us state formally the problem that we have just described. For every nonempty set Y , let $\mathcal{P}(Y)$ be the set of all finite subsets of Y , and $\Delta(Y)$ be the set of all probability measures over Y . Let X be a set of options: for example, the set of lotteries that we have considered so far. A *random choice rule* (RCR) σ is a function from $\mathcal{P}(X)$ to $\Delta(X)$, mapping an element $D \in \mathcal{P}(X)$ to σ^D , such that for every such D , $\sigma^D(D) = 1$. The value $\sigma^D(x)$ is the observed frequency of the choice of x out of D .

Random Utility Models

In random utility models (see [McFadden and Richter, 1991](#), for an early axiomatic analysis, and [Gul and Pesendorfer, 2003](#), for a very recent development) the subject has a set of different potential utility functions (almost different selves), and only one of them is drawn every time she has to make a decision. This momentarily dominant utility decides the choice for that period. Since utilities are different, the choices from the same set of options may be different in different times, although in every period the DM picks the best option.

The hypothesis that random choice is produced by random utilities imposes restrictions on observed behavior. For example, in this class of models choices are made from a set of lotteries, called a menu. Since each of these utility functions is linear, the choice is always in a special subset of the menu (technically, its boundary). A representation of the random-choice rule in random utility models is a probability distribution over utilities such that the frequency of the choice of x out of D , $\sigma^D(x)$ is equal to the probability of the set of utilities that have the element x as a best choice out of D .

Stochastic Choice Models

In stochastic choice models, the utility function is the same in every period. The DM does not always choose the option with the highest utility, but she is more likely to choose an option, the higher its utility is compared to that of the other options. The power of these models is based on two ideas. The first is the decomposition of the decision process in two steps; evaluation and choice. The second is that frequency of choice gives a measure of the strength of preferences. Together, they give a way to identify a cardinal utility. Early axiomatic analysis of this problem is in [Davidson and Marschak \(1959\)](#) and in [Debreu \(1958\)](#). A set of axioms that characterize RCRs which have a stochastic choice representation and that separate these two ideas is presented in, [Maccheroni et al. \(2007\)](#). We examine both ideas in detail.

Utility Function and Approximate Maximization

A representation in stochastic choice models has two elements. The first is the evaluation, which is performed by a utility function that associates a real number with each option in the available set. The second is an approximate maximization function associating to each vector of utilities the probability of choosing the corresponding option.

The utility function is naturally determined on the basis of the random choice rule σ . Write $\sigma(x, y) = \sigma^{\{x,y\}}(x)$ and consider the relation defined by

$$x \succeq y \text{ if and only if } \sigma(x, y) \geq \sigma(y, x).$$

As usual, a function u on X represents the order \succeq if $x \succeq y$ if, and only if, $u(x) \geq u(y)$.

To define the second element, fix u and let U be the range of this function: $U = u(X)$. An approximate maximum selection is a function p from $\mathcal{P}(U)$ to $\Delta(U)$, associating with set A a probability p^A which is concentrated on A (that is, $p^A(A) = 1$) and is monotonic (that is, for every $a, b \in A$ if $a \geq b$ then $p^A(a) \geq p^A(b)$).

A representation of the RCR σ in stochastic choice models is given by a pair (u, p) of a utility function u on X representing \succeq and an approximate maximization function p such that

$$\sigma^D(x) = p^{u(D)}(u(x)) \tag{4.2}$$

In [Maccheroni et al. \(2007\)](#) give a set of axioms that characterize RCR with such a representation.

Moreover, a pair (v, q) represents σ if, and only if, there exists an increasing function $g : u(X) \rightarrow R$ such that

$$v = g \circ u \text{ and } q^B(b) = p^{g^{-1}(B)}(g^{-1}(b)) \quad \forall B \in \mathcal{P}(v(X)) \tag{4.3}$$

In other words, the function u is only determined up to monotonic (not just affine) transformations, so it is still an ordinal, not a cardinal object. Stochastic choice, by itself, does not imply the existence of and does not reveal a cardinal utility.

Strength of Preferences

A measure of the *strength of the preferences* of the DM indicates, for any x, y, z , and w , whether she prefers x to y more than she does z to w . As a special case, it also indicates whether she prefers x to y more than she does z to z itself – that is, whether she prefers x to y , so strength of preferences contains implicitly a preference order. How do we access this measure? One way is through verbal statements made by the DM: she introspectively evaluates the strength and communicates it to the experimenter, with words, not with choice.

Stochastic choice provides us with a second, objective way of measuring the strength of preferences. The value $\sigma(x, y)$ describes how frequently the option x is chosen instead of y . If we compare the frequency of choices out of two other options $\{z, w\}$, and we observe that $\sigma(x, y) > \sigma(z, w)$, then we may say that the DM likes x more than y with stronger intensity than she likes z more than she likes w : we write $(x, y) \geq (z, w)$ to indicate this order over pairs. A random choice rule as characterized in representation (4.2) is a measure of the strength preferences.

Representation (4.2) shows clearly that knowing the strength of preference does not by itself determine the utility function as a cardinal object. We can always introduce a monotonic transformation of the u function, provided we undo this transformation with an appropriate transformation of the approximate maximization function p .

To obtain u as a cardinal object, a specific and strong condition on the random choice rule is needed. The nature of the condition is clear: u is a cardinal object if the strength of preference only depends on the difference in utility, namely if the following difference representation holds:

$$\begin{aligned} \sigma(x, y) \geq \sigma(z, w) \text{ if, and only if,} \\ u(x) - u(y) \geq u(z) - u(w). \end{aligned} \quad (4.4)$$

Debreu (1958) investigates conditions insuring that condition (4.4) is satisfied. A necessary condition for the existence of a u as in (4.4) is clearly:

$$\begin{aligned} \sigma(x, y) \geq \sigma(x', y') \text{ and } \sigma(y, z) \geq \sigma(y', z') \\ \text{imply } \sigma(x, z) \geq \sigma(x', z') \end{aligned} \quad (4.5)$$

(see also Krantz *et al.*, 1971; Shapley, 1975; Köbberling, 2006).

Together with an additional technical axiom (solvability), axiom (4.5) is all is needed for the existence of a function u that is a cardinal object: that is, if a function v also satisfies (4.4) then v is an affine monotonic transformation of u ; that is, there are two numbers $a > 0$ and b such that $v = au + b$.

This opens the way for a complete stochastic choice representation of the random choice rule, with the additional condition that the utility u is cardinal. In a complete model of stochastic choice, if we introduce the additional axiom (4.5) then the approximate maximization function p depends only on the differences, that is:

$$p^{(r,s)}(r) = P(r - s) \quad (4.6)$$

for some function P . The question is now: how is that probability P implemented?

DYNAMIC STOCHASTIC CHOICE

In the plan of determining the neural basis of decision, we have two final steps. First, we have to produce a model of the decision process that produces a stochastic choice as described in the previous section. Second, we have to specify and test the neural basis implementing this process. Let us begin with the first.

The Random Walk Model

The model's original formulation is in Ratcliff (1978). As the title indicates, the theory was originally developed for memory retrieval, where the task is as follows. A subject has to decide whether an item that is in front of her is the same as one she has seen sometimes in the past, or not. She has the following information available. First, she has the visual evidence of the object in front of her. This object can be described abstractly as a vector of characteristics – the color, the smoothness of the surface, the width, the length, and so on. The subject also has some memory stored of the reference object, which can again be described by a vector of the same characteristics as the first one. If the description of the object is very detailed, the vector is a high-dimensional vector. The subject has to decide whether the object in front of her is the same as the object stored in memory, so she has a simple binary (*yes*, it is the same object, or *no*) decision to take.

In an experimental test, we can measure the time the subject takes to decide, her error rate, and how these

variables depend on some parameter that we control – for example, how different the two objects are.

A plausible model of the process is as follows. The subject compares, one by one, each coordinate in the vector of characteristics of the real and recalled object. She may find that, to the best of her recollection, they coincide, or they do not. She proceeds to count the number of coincidences: an agreement of the features is taken as evidence in favor of “yes,” a disagreement as evidence of “no”. If the vector of evidence is very long, the subject may decide to stop before she has reviewed all the characteristics, according to a simple stopping rule: decide in favor of “yes” the first time the number of agreements minus the number of disagreements is larger than a fixed threshold; decide in favor of “no” when a similar lower barrier is reached. The general form of a decision process based on this idea is the random walk of decision. The model has been presented in a discrete or continuous time version. In the continuous time formulation, the process is typically assumed to be a Brownian motion, or at least a time homogeneous stochastic process.

The model has several parameters: first, those describing the process. For example, if the process is in continuous time and is a Brownian motion, the process is described by the mean and the variance. The second parameters are the barriers. There are at least two important observed variables: the probability that one of the two decisions is taken, and the time needed to reach the decision. The model has sharp predictions on the two variables: for example, if the drift in favor of one of the two options is stronger, then the probability of that option being chosen increases. Also, when the difference in drift between the two choice is small, then the time to take a decision increases.

DECISION IN PERCEPTUAL TASKS

Intense research regarding the neural foundation of the random walk model of decision has been undertaken in the past few years. To illustrate the method and the findings, we begin again with a classical experiment (Shadlen and Newsome, 1996, 2001; Schall, 2001).

In the experiment, the subject (for example, a rhesus monkey) observes a random movement of dots. A fraction of the dots is moving in one of two possible directions, left or right, while the others move randomly. The monkey has to decide whether the fraction of dots moving coherently is moving to the left or to the right, being instructed to do this after intensive training. If the monkey makes the right choice, it is compensated by a squirt of juice. Single

neuron recording of neurons shows that the process of deciding the direction is the outcome of the following process: some neurons are associated with the movement to the left, and others to the right. The overall firing rate of the “left” and “right” neurons is, of course, roughly proportional to the number of dots moving in the two directions. The decision is taken when the difference between the cumulative firing in favor on one of the two alternatives is larger than a critical threshold.

Formal Model

A key feature of the information process described above is that each piece of information enters additively into the overall evaluation. This has the following justification. Suppose that information is about a state that is affecting rewards. A state is chosen by the experimenter, but is unknown to the subject. Information is provided, in every period, in the form of signals drawn independently in every period, from a distribution over signals that depends on the state. How is the information contained in the signal observed in every period aggregated?

In a simple formal example, suppose that the decision maker has to choose between two actions: left (l) and right (r). She receives a payment depending on the action she chooses and an unobserved state of nature $s \in \{L, R\}$; this is equal to \$1 if, and only if, she chooses the correct action l if the state is L . Her utility is a function defined on the set $A \equiv \{l, r\}$ of actions and set of states $S \equiv \{L, R\}$ by $u(l, L) = u(r, R) = 1$, $(l, R) = u(r, L) = 0$. She has an initial subjective probability p that the state is R , and can observe a noisy signal on the true state of nature, according to the probability $P_s(x)$ of observing x at s .

The posterior odds ratio of L versus R with a prior P , after the sequence (x_1, x_2, \dots, x_n) is observed, is given by:

$$\frac{P(L|x_1, x_2, \dots, x_n)}{P(R|x_1, x_2, \dots, x_n)} = \frac{P(L)}{P(R)} \prod_{i=1}^n \left(\frac{P_L(x_i)}{P_R(x_i)} \right)$$

so that the log of the odds ratio are simply the sum of the log of the odd ratios of the signal

$$\log \frac{P(L|x_1, x_2, \dots, x_n)}{P(R|x_1, x_2, \dots, x_n)} = \log \frac{P(L)}{P(R)} + \sum_{i=1}^n \log \frac{P_L(x_i)}{P_R(x_i)}.$$

Decision in Economic Choices

We suggest that the mental operation that is performed when the subject has to choose between two

economically valuable options consists of two steps. First, the individual has to associate a utility with each of the two options. Second, she then has to decide which of these two computed quantities is larger. This second step is a simple comparison of quantities. The first is completely new, and is specific to economic analysis. Note two important features of this model: first, even if the decision maker assigns (somewhere in her brain) a strictly larger utility to one of the two options, she still does not choose for sure that option: she only has a larger probability of doing so. Second, the decision maker has a *single* utility or preference order over outcomes. The choice outcome is not deterministic, because the process from utility evaluation to choice is random.

What is the evidence supporting this view? Let us begin from the step involving comparison of quantities. Experiments involving comparison of numbers, run with human subjects (see [Sigman and Dehaene, 2005](#)), confirm the basic finding that the response time is decreasing with the distance between the two quantities that are being compared. For example, if subjects have to decide whether a number is larger or smaller than a reference number, then the response time is decreasing approximately exponentially with the distance between the two numbers. So there is experimental evidence that suggests that the operation of comparing quantities follows a process that is close to that described by the random walk model. The last missing element is: do we have evidence that there are areas of the brain where neurons fire in proportion to the utility of the two options?

THE COMPUTATION OF UTILITY

In this experiment, a monkey is offered the choice between two quantities of different food or juices: for example, 3 units of apple, or 1 unit of raisin.

By varying the quantities of juice of each type offered, the experimenter can reconstruct, from “revealed preferences,” the utility function of the monkey. This function can be taken to be, for the time being, an artificial construct of the theorist observing the behavior. The choices made by the subjects have the typical property of random choice: for example, between any amount less or equal to 2 units of apple and 1 unit of raisin, the monkey always chose the raisin. With 3 units of apple and 1 of raisin, the frequency of choice was 50/50 between the two. With 4 or more units of apple, the monkey always went for the apple. This is the revealed-preference evidence.

At the same time, experimenters can collect single neuron recording from areas that are known to be

active in evaluation of rewards (for example, area 13 of the orbito-frontal cortex). They can then plot the average firing rate over several trials (on the y -axis) against the estimated utility of the option that was eventually chosen on the x -axis, thus obtaining a clear, monotonic relationship between the two quantities. These results are presented in detail in Chapter 29.

A Synthesis

We have now the necessary elements for an attempt to provide a synthesis of the two approaches, one based on economic theory and the other on neuroscience.

Consider a subject who has to choose between two lotteries. When considering each of them, she can assign to it an estimate of the expected utility of each option. This estimate is likely to be noisy. When she has to choose between the two lotteries, she can simply compare the (possibly noisy) estimate of the two utilities: thus the choice between the two lotteries is now determined by the comparison of these two values. At this stage, the choice is reduced to the task of comparing two numerical values, just as the task that the random walk model analyzes.

In summary, this model views the decision process as the result of two components: the first reduces the complex information describing two economic options to a numerical value, the utility of each option. The second performs the comparison between these two quantities, and determines, possibly with an error, the larger of the two. The comparison in this second step is well described by a random walk of decision.

FACTORS AFFECTING THE DECISION PROCESS

In the standard random walk model, the barrier that the process has to hit is fixed. Suppose now that the information available to the decision maker in two tasks is different, and is of better quality in one of them.

For example, in a risky choice the DM has a precise statement on the probability of the outcomes in the lotteries she has to choose from. In the ambiguous choice, on the contrary, she has only limited information on these probabilities. She must provide an estimate on the likelihood of different outcomes on the basis of some reasonable inference. Similarly, in a choice of lotteries that are paid at different points in time, lotteries paid in the current period are easier

to analyze than those paid further in, say, 1 month, because the decision maker has to consider which different contingencies may occur in the next month, and how they might affect the outcome and the utility for her of different consequences. Consider now the prediction of this model on the response time and error rate in the two cases. Intuitively, a harder task should take longer. This is what the random walk model predicts: if the distance from the initial point that the process has to cover is the same, and the process is slower when the information is worse, then the response time should be longer in the harder process. However, we observe the opposite: the response time in the ambiguous choice is consistently shorter than in the risky choice.

A consideration of the extreme case in which the signal that is observed is completely non-informative reveals what might be the missing step. Suppose that indeed the signal provides no information. In this case, waiting to observe the signal provides no improvement over the immediate decision. Since waiting typically implies a cost (at least an opportunity cost of time that could be better used in other ways), the decision in this case should be immediate, because delay only produces a waste of time. So, in the case of the worse possible signal, the response time is the shortest. This conclusion seems to contradict the prediction of the random walk, but instead it contradicts only the assumption that the barrier the process has to hit is fixed. The distance from the initial point at which the process stops should instead depend on the quality of the signal: everything else being equal, a better signal is worth being observed for a longer time.

In the next section we make this informal argument more formal, by showing precisely that when the quality of the signal is better, two opposing factors are active: first, the quality of the signal advises to wait and get better information. This counteracts the second, direct effect (proceeding with a better signal is faster), and may produce what we observe: longer response times with the better, more informative signal.

A Simple Example

The intuitive content of the model can be appreciated better if we consider first the very simple decision problem already introduced in the Formal model section above. If the decision maker receives no additional information, the value for her problem is

$$v(p) = \max\{p, 1 - p\}$$

with the optimal choice of r if $p \leq 0.5$, and l otherwise.

Suppose now that the decision maker can observe instances of an informative signal on the state: the function from the true state to a signal space is called, using a term of statistical theory, an experiment. She can observe the signal produced by the experiment for as many periods as she wants, but the final utility will then be discounted by a factor δ . Now, it is no longer necessarily optimal to choose immediately on the basis of the prior belief; rather, it may be better to wait, observe the signal, update the belief, and make a better choice. Since the value of the reward is discounted, the decision maker has a genuine problem: she has to decide between collecting information, and choosing immediately.

Assume for the moment that an optimal policy, for a given initial belief p , exists. The value of the problem computed at the optimal policy for any such initial belief defines the value function for the problem, which we denote by V . This function is obviously larger than v , since the decision maker has the option of stopping immediately. It is known that the optimal policy for this decision maker can be described as a function of the belief she has regarding the state – that is, on the current value of p . The way in which this dependence works is clear. For a belief p at which $V(p) = v(p)$, the optimal policy is to do what yields $v(p)$; namely, to stop. For the values for which $V(p) > v(p)$, since stopping would only give $v(p)$, the decision maker has to continue experimenting. It turns out in this simple example that there is a cutoff belief, call it p^* , such that it is optimal to stop if, and only if, $p \geq p^*$ or (symmetrically) $p \leq 1 - p^*$.

Consider now the effect on the decision to stop when the quality of the signal provided to the decision maker improves. Introducing a notation used later, we denote the experiments P and Q , with P more informative than Q . Note that the function v does not depend on the experiment, but the value function V and the cutoff belief p^* depend on it, and we write, for example, $V(P, \cdot)$ and $p^*(P)$ to make this dependence explicit. When P replaces Q , the value function V becomes larger, because the information is better (this is intuitively clear, and is proved formally below). Therefore, the set of beliefs at which V is equal to v becomes smaller; that is, the critical belief p^* becomes larger: $p^*(P) \geq p^*(Q)$. Note for future reference that this value also depends on the other parameters of the problem, in particular the discount factor δ , although we do not make this dependence explicit in the notation.

Quality of the Signal and Response Time

What is the effect of this change on the response time? An increase in the value of p^* tends, everything

else being equal, to make the response time longer: it takes more observation to reach a cutoff which is farther from the initial belief. Since an improvement of the signal increases p^* , this direct effect would by itself produce a longer response time. However, a better signal also reduces the time needed to reach a fixed cutoff belief, since the information is more effective.

The net effect is studied below for a more general class of problems, but it is easy to see intuitively what it is. Consider first the case in which the signal provides no information at all. In this case there is no point in waiting and experimenting, and therefore the optimal policy is to stop immediately. Consider now the case in which the experiment provides complete information: as soon as the signal is observed, the state is known for sure. In this case, the optimal waiting time is at most one period: if the decision maker decides to experiment at all, then she will not do it for longer than one period, since in that single period she gets all the information she needs, and additional signals are useless. Note that these two conclusions are completely independent of the value of the discount, since our argument has never considered this value.

Consider now the case of an experiment of intermediate quality between the two extremes just considered: the experiment provides some information, so the posterior belief is more accurate, but the information is never enough to reach complete certainty. If the discount factor becomes closer to 1, then the opportunity cost of gathering additional information becomes smaller. The value of a utility at T is scaled down by a factor δ^T , which is close enough to 1. So if we keep the information fixed, and consider larger and larger values of δ , we see that the cutoff belief p^* increases. Since the experiment is fixed, the effect on the time to reach this cutoff now is unambiguous. Note that in fact the time to stop is a function of the history of signals observed. The probability distribution on this set is given by the experiment. Since the cutoff is higher, for any history the time to reach this cutoff increases, and it is easy to see that we can make it arbitrarily large.

We can now conclude that the time to decide (the response time that we observe) is a hill-shaped function of the quality of information. This conclusion holds in a more general model, which is presented in the Appendix to this chapter.

COGNITIVE ABILITIES AND PREFERENCES

We present how the model we have developed so far can explain experimental as well as real-life choice

behavior of a large group of subjects, relating the choice made in different environments to cognitive abilities. Economic theory makes no statement regarding the correlation between characteristics of individual preferences in different domains. For example, the coefficient of risk aversion is considered independent of the impatience parameter. Also, no correlation is assumed between these preferences and the cognitive ability (CA) of the individual. The predictions of the theory of choice that we have presented are different.

How can cognitive abilities affect preferences? In the theory we have developed so far, the utility of an option is perceived with a noise. The more complex the option is, the larger the noise in the perception. For example, evaluating the utility of a monetary amount paid for sure is easy, and no one has any doubts when choosing between \$10 and \$15. Instead, evaluating a lottery giving on average \$10 is harder, and it is harder still to compare the choice between two lotteries. Similarly, the utility of \$15 to be paid on 10 days is not as sharply perceived as the same payment immediately: we have to consider several different possible intervening factors, such as the impossibility of getting or receiving the payment, other payments that can be received in the same interval, and so on.

Different degrees of CA make the perception of an option more or less sharp. Consider now the choice between a certain amount and a lottery. While the utility of the first is perceived with precision by every individual, the noise around the second one increases for individuals with a lower CA, and so that option is less likely to be chosen by those individuals: subjects with a lower CA make more risk-averse choices. Similarly, in the choice between a payment now and one in the future, they perceive the second more noisily than the first, and so they are less likely to choose it, and they make more impatient choices. The theory predicts that impatience and risk aversion are correlated, and these in turn are correlated with cognitive abilities.

Test of the Theory

We examined whether and how attitudes to risk, ambiguity, and inter-temporal choices are related in a large ($N = 1066$) sample of drivers in an important national (USA) company (see [Burks et al., 2007](#)). Thanks to an agreement with the company, we ran extensive (4 hours) laboratory experimental testing with the participating subjects on a battery of tasks involving choice under risk, ambiguity, choice over time delayed payments, as well as a variety of psychological measurements and cognitive tasks (see [Burks](#)

et al., 1943 for a detailed description of the experiment). Similar results, which confirm the robustness of ours, can be found in Benjamin *et al.* (2007); Dohman *et al.* (2007). From a different perspective, the issue of the connection between cognitive abilities (specifically numeracy) and decision making can be found in Peters *et al.* (2006).

We had three separate measures of CAs: a measure of the IQ (Raven's matrices), a measure of numerical ability (Numeracy) on tests provided by the ETS (Educational Testing Service), and the score on a simple game played against the computer (called *Hit 15*, because the game is a race between two players to reach position 15 on a gameboard) which measures the planning ability of the individual.

In the choice under uncertainty, subjects were asked to choose between a fixed lottery and a varying certain amount. The lottery was either risky (with known, equal probability of the two outcomes) or ambiguous (unknown probability of two colors, and the subject was free to pick the winning color). In choices of different profiles of payments, subjects had to choose between two different payments at two different points in time, a smaller payment being paid sooner.

A first clear effect due to CA was the number of errors the subject made, if we define error (as before) as the number of switches between certain amount and lottery above two. We found that inconsistency increases with our measures of CA, in particular IQ and *Hit 15* score.

The effect of CA on preferences was as predicted: the patience and the index of cognitive ability are positively correlated. Also, risk aversion and the index of cognitive ability are negatively correlated. As a result, there is a negative correlation between risk aversion and impatience.

The effect of the difference in cognitive ability extends to behavior in strategic environments. In our experiment, subjects played a discrete version of the trust game: both players were endowed with \$5; the first mover could transfer either \$0 or the entire amount, and the second player could return any amount between \$0 and \$5. Both amounts were doubled by the experimenter. Before the choice, subjects reported their belief on the average transfer of the participants in the experiment both as first and as second movers.

We found that a higher IQ score makes a subject a better predictor of the choice of the others as first movers: while the average underestimates the fraction of subjects making a \$5 transfer, subjects with higher IQ are closer to the true value. Similarly, they are closer to the true value of the transfers of second movers.

The behavior is also different. As second movers, subjects with higher IQ make higher transfers when they have received \$5, and smaller transfers in the opposite case.

The behavior as first movers is more subtle to analyze, since beliefs also enter into the choice: since subjects with higher IQ believe that a larger fraction of second movers will return money, they might be influenced by this very fact. In addition, the difference in risk aversion might affect choices. Once we control for these factors, however, subjects with higher IQ are more likely to make the \$5 transfer.

We also followed the performance on the workplace in the months following the initial collection of experimental data; in particular, the length of time the subject remained with the company, and, when relevant, the reason for quitting the job. In the training offered by the company, quitting before a year can be safely considered to be evidence of poor planning: trainees leave the company with a large debt (for the training costs have to be paid back to the company if an employee quits before the end of the first year), they have earned little, and have acquired no useful experience or references for their resumé. If we estimate the survival rate for different socio-economic variables (for example, the married status), then the variables have no significant effect on the survival rate, while the *Hit 15* affects it largely and significantly.

APPENDIX: RANDOM WALK WITH ENDOGENOUS BARRIERS

We denote the unknown parameter (for example, the state of nature) as $\theta \in \Theta$. The decision maker has an initial belief on the parameter, $\mu_0 \in \Delta(\Theta)$, and has to take an action $a \in A$. The utility she receives depends on the state of nature and the action taken, and is described by a function $u : \Theta \times A \rightarrow R$. She can, before she takes the action, observe a signal $x \in X$ with a probability that depends on the state of nature, denoted for example by $P_\theta \in \Delta(X)$. In classical statistical terminology, this is an experiment $P \equiv (P_\theta)_{\theta \in \Theta}$. For any given prior belief on the set Θ , this experiment induces a probability distribution on the set of signals:

$$P_\mu(x) = \int_{\Theta} \mu(d\theta) P_\theta(x).$$

The subject can observe independent replications of the signal as many times as she likes, stop, and then choose an action $a \in A$. The use of the experiment has a fixed cost c for every period in which it is used.

The information she has at time $t \in \{0, 1, \dots\}$ is the history of signals she has observed, an element

$(x_0, \dots, x_{t-1}) \in X^t$. The posterior belief at any time t is a random variable dependent on the history of signals she has observed, and is denoted by μ_t . Let $B(\mu, x)$ denote the posterior belief of a Bayesian decision maker with a prior belief μ after observing a signal x . We write $B(\mu, x, P)$ if we want to emphasize the dependence of the updating function on the experiment P .

The decision maker can make two separate choices in each period: first, whether to stop observing the signal, and second, if she decides to stop, which element of A to select. The action she chooses at the time in which she stops is optimal for her belief at that time. If her posterior is ν , her value at that time is equal to $v(\nu)$, the expected value conditional on the choice of the optimal action, namely:

$$v(\nu) = \max_{a \in A} E_{\nu} u(\cdot, a) \quad (4.7)$$

Conditional on stopping, the action in A is determined by the maximization problem we have just defined, and the value of stopping is given by v . We can therefore focus on the choice of when to stop.

A policy π is a sequence of mappings $(\pi_0, \dots, \pi_t, \dots)$, where each π_t maps the history of observations at time t , (x_0, \dots, x_{t-1}) into $\{0,1\}$, where 1 corresponds to *Stop*. The first component π_0 is defined on the empty history.

The initial belief μ and the policy π define a probability distribution over the set of infinite histories X^∞ , endowed with the measurable structure induced by the signal. We denote by $E_{\pi, \mu}$ the corresponding expectation. Also there is a stopping time T (a random variable) determined by the policy π , defined by

$$T \equiv \min_t \{\pi_t(x_0, \dots, x_{t-1}) = 1\}$$

The expected value at time zero with the optimal policy depends on the signal the subject has available, and is given by

$$V(\mu, P) = \max_{\pi} E_{\pi, \mu} \left[\delta^T v(\mu_T) - \sum_{t=1}^T c \delta^{t-1} \right]$$

where we adopt the convention that $\sum_{t=1}^0 c \delta^{t-1} = 0$.

Normally distributed signals An important example is the class of normally distributed experiments. Let Θ be a subset of the real line, indicating the expectation of a random variable.

An experiment P is defined as the observation of the random variable X distributed as $N(\theta, \sigma^2)$, where the variance σ^2 is known. An experiment Q , given by the observation of the variable Y distributed as

$N(\theta, \rho^2)$, is dominated by P if, and only if, $\rho > \sigma$. This is in turn equivalent to the existence of a normal random variable Z with zero expectation and variance equal to $\rho^2 - \sigma^2$ such that

$$Y = X + Z.$$

OPTIMAL POLICY

The operator M on the space of continuous functions on $\Delta(\Theta)$ with the sup norm is defined by

$$M(P, W)(\mu) \equiv \max\{v(\mu), -c + \delta E_{\mu, P} W(B(\mu, \cdot))\}$$

where the function v is defined in (4.7).

This operator is a contraction on that space, because it satisfies the conditions of Blackwell's theorem. Hence the value function V exists, and is the solution of the functional equation:

$$V(\mu, P) = M(P, V)(\mu), \text{ for every } \mu.$$

The value function equation describes implicitly the optimal policy, which is a function $\hat{\Pi}$ of the current belief. As in our simple example, the policy is to stop at those beliefs in which the value function V is equal to the value of stopping immediately, v . Formally we define the stopping time region $S(P) \subseteq \Delta(\Theta)$ as

$$S(P) \equiv \{\mu : v(\mu) = V(\mu, P)\}$$

The optimal policy is stationary: the function π_t depends on the history of signals only through the summary given by the current belief. This optimal policy is described by the function $\hat{\Pi}$

$$\hat{\Pi}(\mu) = 1 \text{ if and only if } \mu \in S(P).$$

VALUE AND QUALITY OF SIGNALS

Consider now two experiments of different quality, P and Q say. Let \succeq denote the partial order (as defined by Blackwell, 1951; Targerson, 1991) over experiments. We now show that if the experiment is more informative, then the set of beliefs at which the decision maker continues to observe the signal is larger than it is for the worse signal.

Theorem 4

1. The operator M is monotonic in the order \succeq , that is, for every function $W : \Delta(\Theta) \rightarrow \mathbb{R}$, if $P \succeq Q$, then $M(P, W) \geq M(Q, W)$, and therefore $V(\cdot, P) \geq V(\cdot, Q)$
2. The optimal stopping time region S is monotonically decreasing, namely if $P \succeq Q$ then $S(P) \subseteq S(Q)$.

In terms of our main application, decision under risk and uncertainty, the conclusion is that with a richer information (risk) the barrier where the random walk stops is farther than it is with the more poor information (ambiguity). As a consequence, the updating process may take longer in risk than in ambiguous choices.

Quality of signals and response time We now present formally the argument presented informally in our analysis of the simple example. Recall first that:

1. An experiment is called totally un-informative, denoted by P^u , if

$$\text{for all } \theta^1, \theta^2 \in \Theta, P_{\theta^1}^u = P_{\theta^2}^u$$

2. An experiment is called totally informative, denoted by P^u , if

$$\text{for all } \theta^1, \theta^2 \in \Theta, P_{\theta^1}^u \perp P_{\theta^2}^u,$$

that is the two measures are mutually singular.

We now have:

Lemma 5

1. If the experiment P is totally informative, then at the optimal policy the stopping time $T \leq 1$, $(\pi, \mu) - a.s.$;
2. If the experiment P is totally un-informative, then at the optimal policy the stopping time $T = 0$, $(\pi, \mu) - a.s.$;

As in the analysis of our simple example, note that the two conclusions are independent of the discount factor δ and the cost c .

We now turn to the analysis of the response times when the experiments have intermediates, namely for experiments P such that $P^i \succeq P \succeq P^u$.

Define the function

$$U(\mu) \equiv \int_{\Theta} \max_{a \in A} u(\theta, a) d\mu(\theta)$$

This is the value at the belief μ of a decision maker who is going to be completely and freely informed about the state before she chooses the action. Information is always useful for her (for every belief that is different from complete certainty about a state) if the value of the optimal choice at μ is smaller than the expected value when complete information will be provided:

Assumption 6

Information is always useful, namely

$$\text{For every } \mu \in \Delta(\Theta), \text{ if } \mu \notin \{\delta_{\theta} : \theta \in \Theta\}, \text{ then } U(\mu) > v(\mu). \tag{4.8}$$

Assumption 7

An experiment P is intermediate, that is:

1. For every finite number n of independent observations, and initial belief in the relative interior of $\Delta(\Theta)$,

$$B(\mu, P^n)$$

is in the relative interior of $\Delta(\Theta)$.

2. As the number of independent observations tends to infinity, the product experiment converges to the totally informative experiment.

Theorem 8

If information is always useful (assumption 9.3) and the experiment is intermediate (assumption 9.4), then

$$\lim_{c \downarrow 0, \delta \uparrow 1} T = \infty, (\pi, \mu_0) - a.s.$$

where π is the optimal policy.

References

Anscombe, A. and Aumann, R. (1963). A definition of subjective probability. *Ann. Math. Stat.* 34, 199–205.

Benjamin, D., Brown, S., and Shapiro, J. (2007). Who is “behavioral”? Discussion Paper.

Blackwell, D. (1951). Comparison of experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley, CA: University of California, 93–102.

Burks, S., Carpenter, J., Goette, L. et al. (1943). Using behavior economic experiments at a large motor carrier: the context and design of the truckers and turnover project. IZA DP No 2789.

Burks, S., Carpenter, J., Goette, L., and Rustichini, A. (2007). Cognitive abilities explain economic preferences, strategic behavior and job performance. Working Paper, University of Minnesota.

Cartwright, D. (1941a). The relation of the decision time to the categories of response. *Am. J. Psychol.* 54, 174–196.

Cartwright, D. (1941b). The decision time in relation to the differentiation of the phenomenal field. *Psychological Rev.* 48, 425–442.

Cartwright, D. and Festinger, L. (1943). A quantitative theory of decision. *Psychological Rev.* 50, 595–621.

Davidson, D. and Marschak, J. (1959). Experimental tests of stochastic decision theory. In: C.W. Churchman (ed.), *Measurement Definitions and Theories*. New York, NY: John Wiley & Sons, pp. 233–269.

Debreu, G. (1958). Stochastic choice and cardinal utility. *Econometrica* 26, 440–444.

Dickhaut, J., McCabe, K., Nagode, J. et al. (2003). The impact of the certainty context on the process of choice. *Proc. Nat. Acad. Sci.* 100, 3536–3541.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2007). Are risk aversion and impatience related to cognitive abilities. IZA DP 2735.

Gul, F. and Pesendorfer, W. (2003). Random expected utility. *Econometrica* 74, 121–146.

Krantz, D., Luce, R., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, vol. I, *Additive and Polynomial Representations*. New York, NY: Academic Press.

- Köbberling, V. (2006). Strength of preferences and cardinal utility. *Economic Theory* 27, 375–391.
- Maccheroni, F., Marinacci, M., and Rustichini, A. (2007). Preference based decision process. Working Paper, University of Minnesota.
- McFadden, D. and Richter, M. (1991). Revealed stochastic preferences. In: J.S. Chipman, D. McFadden, and M.K. Richter (eds), *Preferences, Uncertainty and Optimality*. Boulder, CO: Westview Press, pp. 161–186.
- Moyer, R. and Landauer, T. (1967). Time required for judgements of numerical inequality. *Nature* 215, 1519–1520.
- Padoa-Schioppa, C. and Assad, J. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223.
- Peters, E., Vastfjall, D., Slovic, P. et al. (2006). Numeracy and decision making. *Psychological Sci.* 17, 407–413.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Rev.* 85, 59–108.
- Rustichini, A., Dickhaut, J., Ghirardato, P. et al. (2005). A brain imaging study of the choice procedure. *Games Econ. Behav.* 52, 257–282.
- Savage, L.J. (1954). *The Foundation of Statistics*. New York, NY: Dover Publications Inc.
- Schall, J.D. (2001). Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* 2, 33–42.
- Shadlen, M.N. and Newsome, W.T. (1996). Motion perception: seeing and deciding. *Proc. Natl Acad. Sci. USA* 93, 628–633.
- Shadlen, M.N. and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (Area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.
- Shapley, L. (1975). *Cardinal utility comparisons from intensity comparisons*. Report R-1683-PR, The Rand Corporation, Santa Monica, CA.
- Sigman, M. and Dehaene, S. (2005). Parsing a cognitive task: a characterization of mind’s bottleneck. *PLoS Biol.* 3, e37.
- Smith, P. (2000). Stochastic dynamic models of response time and accuracy: a foundational primer. *J. Math. Psychol.* 44, 408–463.
- Torgersen, E. (1991). *Comparison of Statistical Experiments*. Cambridge: Cambridge University Press.
- von Neumann, J. and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*, 2nd edn. Princeton, NJ: Princeton University Press.