Cognitive, Affective, & Behavioral Neuroscience
2008, 8 (4), 348-354
doi:10.3758/CABN.8.4.348

Perspectives From Economics and Neuroeconomics

Understanding risk: A guide for the perplexed

Paul W. Glimcher
New York University, New York, New York

Over the course of the past decade, neurobiologists have become increasingly interested in concepts and models imported from economics. Terms such as "risk," "risk aversion," and "utility" have become commonplace in the neuroscientific literature as single-unit physiologists and human cognitive neuroscientists search for the biological correlates of economic theories of value and choice. Among neuroscientists, an incomplete understanding of these concepts has, however, led to a growing confusion that threatens to check the rapid advances in this area. Adding to the confusion, notions of risk have more recently been imported from finance, which employs quite different, although formally related, mathematical tools. Of course, the mixing of economic, financial, and neuroscientific traditions can only be beneficial in the long run, but truly understanding the conceptual machinery of each area is a prerequisite for obtaining that benefit. With that in mind, I present here an overview of economic and financial notions of risk and decision. The article begins with an overview of the classical economic approach to risk, as developed by Bernoulli. It then explains the important differences between the classical tradition and modern neoclassical economic approaches to these same concepts. Finally, I present a very brief overview of the financial tradition and its relation to the economic tradition. For novices, this should provide a reasonable introduction to concepts ranging from "risk aversion" to "risk premiums."

Over the course of the last decade, economic theory has begun to have a significant influence on the practice of neurobiological research. Neurobiological articles have begun to examine the neural representation of concepts such as "expected value," "utility," "expected utility," "risk taking," and "risk aversion." These are terms that have entered the neurobiological lexicon from economics specifically because they are precise mathematical concepts supported by well-developed proofs. Unfortunately, many neurobiological articles use these terms incorrectly, a trend that has led to a growing confusion about what we do and do not know, at a neurobiological level, about the mechanism of choice. To take one example: Many neurobiologists do not realize that within the framework of expected utility theory, it is impossible to measure both "risk aversion" and "utility" independently—because these terms refer to two approaches to quantifying the same underlying object. In a similar way, many neurobiologists do not realize that if they say: "My data set reveals risk-averse behavior," an economist colleague might hear them as insisting that the chooser they have studied strictly obeys expected utility theory (or a closely related theory) in all of his or her choices.

One response to the growing confusion this loose use of language has produced has been to argue that the actual mathematical meaning of words such as "utility" are of little relevance to neurobiologists, for whom these theories are simply coarse metaphors. Although this response is common today, it seems to me an odd choice. No computational neuroscientist would refer to a Kalman filter as a coarse metaphor ripe for redefinition, or suggest that the differences between the g_{Na} and g_{K} terms in the Hodgkin–Huxley equations (which represent features of the sodium and potassium channels, respectively) were of little significance. Instead of eschewing theory, neurobiologists must strive to be both clearer in their understanding of what these words mean and more precise in their language, if our field is to benefit from the well-developed theories of value that gave rise to these words we now try to employ.

Risk, Value, and Utility

Whereas all notions of risk, value, and utility have their roots in economics and finance, three separate but related traditions have used these words in slightly different ways. In the classical economic tradition, value, risk, and choice were related by simple but ad hoc mathematical functions now called "Bernoullis" by economists. In the neoclassical economic tradition, simple assumptions about the be-

P. W. Glimcher, glimcher@cns.nyu.edu

havior of human choosers were used to constrain a class of functions that could, in principle, relate value and risk to choice. In finance, yet another approach has been to relate value and a linear property called a "risk preference" to choice. Although all of these approaches are distinct, they are also heavily interrelated. What follows is a discussion of each of these sets of ideas, and a definition of the relationships between these approaches.

THE CLASSICAL ECONOMIC TRADITION

Expected Value

The classical economic tradition began with the work of Pascal (1670/1966), who sought to develop a normative technique for decision making. Pascal's idea was that decision makers should always act to maximize their average long-run payoff, whether that payoff would be in dollars or spiritual salvation. To understand how Pascal achieved this maximization, we begin by defining each of the available actions placed before a chooser as an "option." Pascal argued that a chooser should assess the desirability of each option by representing it as having a value (how much one stood to gain or lose) and a likelihood (the probability of that gain or loss). When an agent must choose between two or more options, he then simply multiplies the value of each option by the likelihood that option will pay off to yield what is called an expected value (EV):

$$\text{Expected Value} = \text{Probability} \times \text{Value}.$$

The chooser then simply selects the option having the highest EV.

To take a specific example, consider a chooser faced with two options: (1) throwing a single six-sided die, which pays \$60 if a 4 appears and nothing if any other number appears, or (2) throwing two six-sided dice, which pays \$370 if two 4s appear and nothing if any other combination appears. Table 1 shows us how Pascal (1670/1966) would have us compute these expected values.

Table 1
Pascal's Method of Computing Expected Value

	Option 1	Option 2
Value	60	370
Probability	1/6	1/36
Expected value	10	10.28

According to Pascal (1670/1966), Option 2 is revealed by this operation and should therefore be chosen. There are two important things to note about Pascal's formulation. First, EV has no free parameters as a mode of choice—it is an entirely objective formulation. Second, it does not take into account the fact that having a wealth of zero (at least under some conditions) can be thought of as a special circumstance. This means, as shown in Table 2, that, by Pascal's formulation, a man with a net worth of \$0, if faced with the choice between a sure gain of \$5 and a 1/36th chance of \$370, should pick the option with an EV of \$10.28, even though this means he will almost certainly remain penniless.

Table 2
Further Example of Pascal's Method

	Option 3	Option 4
Value	5	370
Probability	1	1/36
Expected value	5	10.28

Bernoulli's Moral Value (Classical Expected Utility)

Historically, Pascal's (1670/1966) theory of expected value was immensely influential. It allowed insurance companies, for example, to determine the long-run average values of insurance policies (Huygens, 1657). The theory, however, was widely observed to do a poor job of predicting the choice behavior of actual human beings. In the choice presented in Table 2, for example, nearly all humans report that they would prefer Option 3 to Option 4, even though Pascal urges us to select Option 4. To explain this apparent hesitancy to accept risky options, the 18th-century mathematician Daniel Bernoulli suggested a modification to Pascal's approach (Bernoulli, 1738/1954). Bernoulli's specific suggestion was that humans did not in fact choose by multiplying objective value with probability to yield a decision variable (in this case, expected value) but rather made their choices by multiplying a "subjective value" with probability to yield a subjective decision variable (in this case, expected utility).

Bernoulli (1738/1954) thus proposed the following model: Begin with the current total objective wealth of the chooser. Take the log of this number to yield the "subjective wealth" of the chooser. To evaluate any option, one first assesses the "increment in subjective wealth" (known as the "marginal wealth") it would produce and then multiplies this increment by the probability of realizing that gain to yield the decision variable of interest. To make this process clearer, consider a chooser having an initial wealth of \$1 and facing the options presented in Table 3. Under these conditions, the chooser selects the sure win of \$5, because in subjective terms, that option is much more desirable.

The critical idea in Bernoulli's (1738/1954) formulation is that objective wealth is transformed into subjective units using a compressive function, and this, he argues, is what accounts for the unwillingness of people (especially poor people) to take risks in the way that Pascal (1670/1966) had imagined they should. Figure 1 shows this graphically. The lower axis plots dollars, and the ver-

Table 3
Bernoulli's Method of Computing Expected Utility

	Option 3	Option 4
Starting wealth	1	1
Log(starting wealth)	0	0
Ending wealth	6	371
Log(ending wealth)	0.78	2.57
Subjective wealth increase	0.78	2.57
Probability	1	1/36 (.028)
Decision variable	0.78	0.07

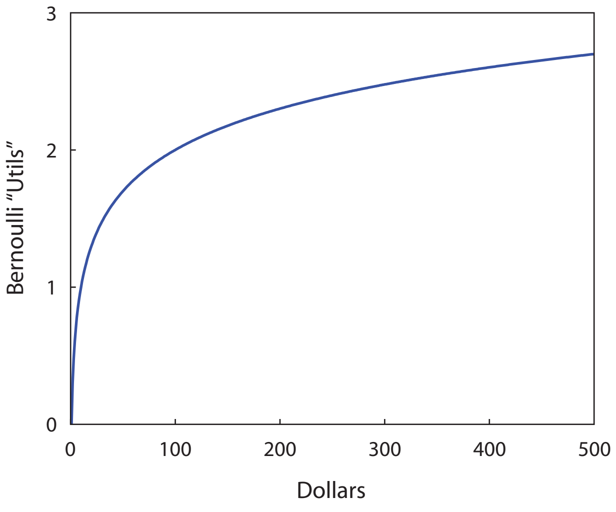

Figure 1. A logarithmic utility curve of the kind envisioned by Bernoulli (1738/1954).

tical axis plots subjective values. The logarithmic curve, which Bernoulli selected on totally ad hoc grounds, relates these two quantities.

What is amazing about Bernoulli's (1738/1954) formulation, and what has made it so influential, is that it yields a subjective unwillingness to take risks by compressively transforming value. To make this very important point clear, consider the two options shown in Figure 2. Option 5 is a guaranteed $100, which is worth 2 in subjective units. If we contrast that with Option 6, where the probability has been halved and the objective value doubled to $200, note that the subjective value has been increased only by about a third. In other words, the log transformation means that to compensate for halving the likelihood, we need to *more* than double the objective value if we want

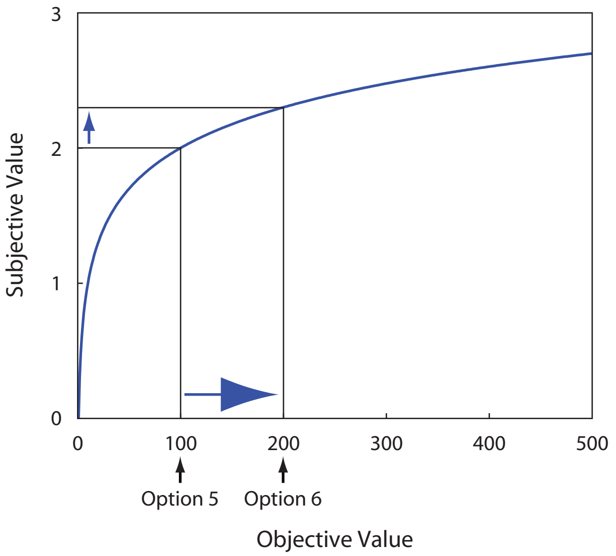

Figure 2. A demonstration of why halving probability and doubling objective value lead to a reduction in expected utility.

to keep overall desirability constant. In humans making choices among risky options, we do observe a preference for lower risks, but the decision maker avoids risks in the Bernoulli formulation *simply because of the logarithmic transformation of value*. Put most explicitly, in this model, risk aversion comes from the subjective representation of value and from nowhere else.

Bernoulli (1738/1954) referred to the logarithmically transformed objective value as a "moral value." In the classical economic tradition, this "moral value" came to be known as a "utility," and the product of utility and probability has come to be known as "expected utility." It should be immediately clear, however, that Bernoulli's adoption of the log transform was entirely ad hoc. He could just as easily have said Utility = (Objective Value)$^{.6}$, and this alternative formulation would have yielded broadly similar results.[1] For this reason, modern economists refer to any function that maps objective monetary value to subjective value as a "Bernoulli."

Bernoulli's (1738/1954) model has both good and bad features that need always to be borne in mind. Its best feature is that it accounts for the observation that human choosers are averse to risks, and it accomplishes this in a very subtle and elegant way. A bad feature is that it does this on very ad hoc grounds, and his selection of the logarithmic function as a way to describe all existing human choosers seems not particularly well supported by our available data (Holt & Laury, 2005; Wu & Gonzalez, 1999). Still, these general notions of EV and expected utility form the core of many analyses of choice and served as motivating ideas in the neoclassical economic revolution of the 20th century.

THE NEOCLASSICAL ECONOMIC TRADITION

During the classical period, it was widely acknowledged that different people placed different subjective values on different goods, that people had idiosyncratic *preferences* (Smith, 1776/1976). It was also clear from Bernoulli's (1738/1954) work that if (1) people made choices on the basis of these hidden preferences, and (2) the hidden functions that related the amount of a good to the strength of the preference (the utility) had certain forms, then one would predict an aversion in choosers to risky options. What troubled the economists of the early 1900s, however, was that this chain of reasoning was both ad hoc and based firmly on an object that by definition could never be observed directly: utility.

To try to overcome this limitation, a group of economists, probably most notably Paul Samuelson (1938), wondered whether economics could be rebuilt in a way that avoided both ad hoc modeling and a reliance on "unobservable" objects. To that end, Samuelson introduced a new and more rigorous modeling technique now known as "revealed preference." Samuelson's strategy was to define on a priori grounds a very small number of assumptions about the behavior of choosers and then to ask mathematically whether these very limited assumptions, or "axioms," allowed one to make any predictions at all about the choice behavior of individuals who obeyed these axioms.

To understand this approach to studying choice behavior, consider a set of axioms widely known as the generalized axiom of revealed preferences, or GARP (Houthakker, 1950). GARP assumes (states as a testable mathematical axiom) that if a chooser reveals that he prefers apples to oranges, and oranges to peaches, then apples are "indirectly" revealed by his behavior as preferred to peaches (and similarly for longer chains of indirect revelation). This is the core falsifiable statement in GARP. It need not be true that a human chooser have this feature, but what Houthakker asked was: If a human did obey this rule, then what else might we be able to prove we already know about him? Taking this approach, what Houthakker proved was that if the axiom of GARP holds for binary choices among pairs of objects, then a small set of choices can be used to make predictions about the relative desirability of pairs of objects that have never been directly compared by the consumer. Consider a situation in which a consumer chooses an apple over an orange and then an orange over a peach. If the assumption of GARP is correct, then this consumer must not consistently choose a peach over an apple even if this is a behavior we have never observed before. Further, Houthakker went on to show, and this is the important part, that making such a statement *is formally equivalent to saying* that there is a utility function of some kind (or, more formally, a "preference function") that describes the values people place on the objects of their choices. The critical idea here is that observing that a chooser obeys the axioms of GARP is the same as observing that a chooser has a utility function of some kind. Let me hasten to add, however, that GARP made no particular predictions about risk. Choosing among risky gambles like the ones Bernoulli (1738/1954) had explored was not a feature of this particular formulation.

In their famous book, von Neumann and Morgenstern (1944) extended this approach, rooted in Samuelson's (1938) insight, into the realm of risk. They developed a larger assumption set (four specific axioms) that allowed them to make specific predictions about how a chooser who obeyed their axioms would behave when he or she faced risky options.[2] (For an overview of axiomatics in general and von Neumann and Morgenstern's [1944] contribution, in particular, see Mas-Collel, Whinston, & Green, 1995; Rustichini, 2009). What von Neumann and Morgenstern showed in their "theory of expected utility" was that any chooser who obeyed (in their choices) these four axioms behaved *exactly as if* they had a continuous monotonic utility function. That is to say, the general logic Bernoulli's (1738/1954) intuition had suggested was correct *if* (and this is a big if) the person being studied obeyed the axioms of "completeness," "transitivity," "continuity" (the Archimedian axiom), and "independence." What this means is that if one tests each of these axioms in a given chooser and he obeys all four axioms, then one knows that the chooser has a monotone utility function. One does not, however, know whether that function is compressive/concave (whether an individual is risk averse). To determine that, an additional test is required: A chooser must be shown to (on average) prefer a lower level of risk when choosing between two options of equal EV. Such a chooser can then be said, in the expected utility tradition, to be "risk averse."

In the neoclassical tradition, then, risk itself has a highly specific meaning (Rothschild & Stiglitz, 1970). One option can be defined as riskier than another if the probability distribution of the first option can be seen as less certain than the other.[3] Alternatively, if a rational chooser strictly prefers one of two options having equal EV, then the preferred option can be defined as less risky than the other option. This is the meaning of risk in the neoclassical framework.

It is incredibly important to point out here that simply showing that a chooser prefers a lower level of risk when choosing between two options of equal EV is *not* sufficient evidence to conclude that he or she is risk averse, as some neurobiologists have (probably accidentally) claimed. A chooser must both show this feature and demonstrate rationality in his or her choice patterns to be called risk averse.

Thus, risk aversion has a fairly complete and technical meaning for a modern economist. It means that (1) the chooser obeys a set of axioms that allow us to make global predictions about the pattern of choice behavior and (2) a small set of measurements have been made on choices between risky alternatives showing a particular class of preferences within the constraints of the axiomatic theory. Importantly, within the specific framework of expected utility theory, the two statements made above are formally equivalent to saying (1) the chooser has a continuous monotonic (never falling) utility function and (2) that function is strictly concave (curves in the same direction as Bernoulli's [1738/1954] function curved).

A central point of modern economic theory is that these two sets of statements are fully interchangeable. There is no way one could independently test the hypothesis that a set of choices obeys the axioms of expected utility theory and that the choices are based on a monotonic utility function. These are saying exactly the same thing. In a similar way, one cannot ask whether a chooser who obeys expected utility theory *both* is risk averse *and* has a concave utility function. If they are risk averse, they must have a concave utility function; if they have a concave utility function, they must be risk averse. This is what these phrases mean.

One final point worth stressing is that simply observing that a chooser obeys the axioms of, for example, expected utility theory places very few constraints on the actual shape of his or her utility function—it only constrains the utility function to be monotone and continuous. Other measurements can then further constrain the set of possible utility functions. (For an example of how this is done in practice, see Holt & Laury, 2005.) Choices among risky gambles can be used to determine whether the function must be concave or convex. More subtle choice sets can determine whether the curvature of the function has a constant first derivative or a constant ratio of the second and first derivatives. More measurements place more constraints in a hierarchical fashion that proceeds from the axioms themselves.

THE FINANCIAL TRADITION

Approaches to risk and decision in financial theory begin from a very different conceptual point. For about a half-century, financial theorists have been interested both in how asset prices are determined in existing financial

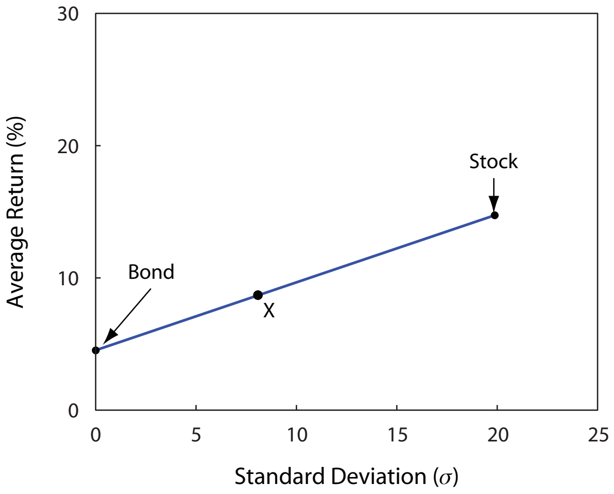

Figure 3. The asset allocation problem in finance.

markets and in determining how to design efficient portfolios of investments in the real world for real investors. To understand risk aversion in the financial tradition, it is therefore helpful to begin thinking about investments. Here, we begin with stock markets and government bonds. Let us define two assets. The first is a risk-free investment that yields an annual rate of return of 5%. One can think of this as something like a United States Treasury Bill. The second is, for the purposes of this example, a single stock. Let us simplify by asserting that this stock has a known average annual rate of return of 15%, with an annual standard deviation (σ) of 20%. Figure 3 shows the options for an investor deciding how to divide her money between these two assets. At the far left of the line, she invests all her assets in the bond. At the midpoint of the line, she invests half her money in each. At the far right, she invests all of her money in the stock. What is observed in financial situations is that human choosers place themselves somewhere on this line by picking a proportion of their assets and placing them in the stock and placing the remainder in the bond. Those more tolerant of risks will naturally place a larger proportion in the stock; those more averse to risk will place them in the bond.[4] What a chooser selects therefore reveals the degree of risk aversion he or she possesses, from a financial perspective. Point X in the figure gives an example of this kind of choice.

To understand this behavior in more economic terms, let us recast the line of Figure 3, which describes the options available to our chooser, as a set of choices between conventional economic lotteries. To make the math easy, let us further assume that the investor has \$100 in hand. On the far left is a certain gain of \$5. On the far right, the stock describes a set of possible outcomes, with a Gaussian distribution centered at \$15, shown in Figure 4. The stock thus offers all of the values on the lower axis at all of the probabilities on the left axis. The EV of the set of lotteries offered by the stock here is \$15. So one can think of a risk-neutral chooser as someone who would pick to invest entirely in the stock and an infinitely risk-averse chooser as one who would pick to invest only in the bond. Of course, the line that connects these two extreme options thus describes a family of "lottery sets" that are all offered to the chooser simultaneously, and the point that the chooser selects suggests a degree of risk aversion.

Now, consider a situation in which the stock has the same standard deviation but yields a higher rate of return. The lottery it presents (and thus the intermediate points connecting it with the bond) have higher values for the same standard deviation. Under these conditions, our chooser would naturally be expected to shift her choice to the right, because she is in essence being paid more for taking on these risks. In a similar way, if the risky investment yielded lower gains (relative to the risk-free investment), her choice would shift to the left. Theoretical work by Markowitz (1991) and Sharpe (1964) established that for a particular kind of efficient chooser, the way this point moves as the value of the stock changes can be computed analytically. Once you know the chooser's degree of risk aversion (and make some assumptions about the shape of the chooser's utility function) from a single such choice, you can fully specify the efficient portfolio allocation for any line of this kind, at any slope, for this individual.

The key to understanding this fact is to reconceptualize risk aversion as a "risk premium." The idea here is that the infinitely risk-averse subject, the one who buys only the bond, will not take any risk at all, no matter how much more it pays. The chooser at Point X, however, was willing to accept an increase in the standard deviation of her

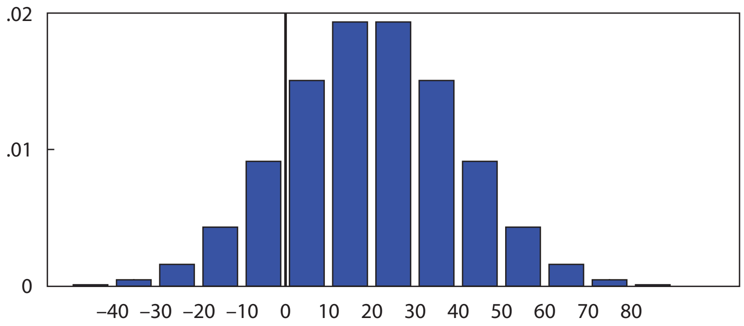

Figure 4. A lottery-style representation of the possible gains and losses associated with a stock.

payoff of 10% as long as she was compensated with an increase in the expected value of $5. To put that in more traditional financial terms: In order to accept an increase in the percent variance (which is equal to the square of the standard deviation) of .01, she had to be compensated by an increase in mean payoff of 4.8%. If she had been less risk averse, she would have required less mean compensation to accept the additional variance. Of course, this explains why less risk-averse investors move to the right on the line in Figure 3; the compensation offered at a standard deviation of 10 is more than enough for them, so they move their choice point farther to the right than Point X.

To return to our canonical investor, she has a degree of risk aversion (in the neoclassical economic sense) that caused her to pick the lottery with an expected value of $10 as the most preferred (from among all of the stock–bond combinations she could have chosen). In financial theory, this means she placed a risk premium of 4.8% on a percent variance of .01. What Markowitz (1991) and Sharpe (1964) showed was that this is enough information for us to predict all of her variance-associated behavior *if* (and again, this is a big if) she is an efficient investor with a particular kind of utility function. In fact, what they showed was that, for a chooser of this type, the risk premium as a function of percent variance should be a straight line (Markowitz, 1991; Sharpe, 1964). Figure 5 shows this in a very simplified graphical format. The darkest (blue) line plots the risk premium required by this chooser for any degree of variance. The line connects the premium she would require to accept a variance of .01—in this case, 4.7%—with the premium she would require to accept a sure thing—of necessity, a risk premium of $0. The straight line that connects these two points thus captures all possible risk premiums, *given the assumption that her risk premium can be described as a straight line in this space*.

To relate this back to the neoclassical tradition, let us redescribe Figure 5. The lower axis of Figure 5 plots the variance of all Gaussian distributed lottery sets that one could possibly offer to the chooser. The dark (blue) line tells about this chooser's indifference point when comparing these lottery sets with certain gains. If one offers this subject a lottery set with a variance of .01 (a standard deviation of 10%) versus a certain gain worth 4% less than the EV of the lottery set, she will prefer the certain gain. If one offers her the same lottery versus a certain gain worth 6% less than the EV of the lottery set, she will prefer the lottery set. The slope of an individual's risk premium line is thus a measure of aversion to risk. Since the risk premium line is straight and has an intercept at 0, we can represent it, and thus capture the preferences of the chooser (given our assumption of linearity in this space), with a single number: the slope of that line.

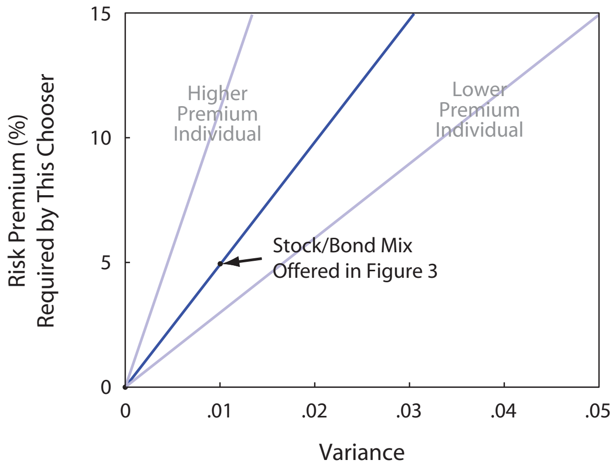

Figure 5. A representation of risk aversion in the financial domain.

Relating the Economic and Financial Traditions

What should be immediately obvious is that the objects of choice differ significantly in these two traditions. The financial tradition has focused on representing stocks and stock-like objects that are well described and tractably analyzed by mean and variance.[5] In contrast, the economic tradition has focused on goods and lotteries. The differences between these classes of objects are one of the things that differentiate these two approaches most clearly.

A larger difference when it comes to understanding risk, however, is the philosophically different way economics and finance approach the causes of risk-averse behavior. At an almost algorithmic level, expected utility theory (as stated by von Neumann & Morgenstern, 1944) argues that humans understand and are largely objective in their assessment of probabilities. In utility theory, the subjective value of a good grows more slowly than objective value; this is the curvature of the utility function. When combined lawfully with probabilities, this yields an aversion to risk. In contrast, finance theory argues that choosers have an objective representation of both value and probability, but that they require a subjectively defined inducement to accept risks; this is captured in the slope of the risk premium line.

These philosophical differences lead quite naturally to terminological differences that can be quite confusing. In the economic tradition, saying that a subject is "risk averse" means that he has a concave utility function. In the finance tradition, saying that someone is "risk averse" means that he has a nonzero slope to his risk premium line. These are similar, although not completely identical, concepts. In the economic tradition, when one speaks of risk, one refers to lotteries (or events) with probabilistic outcomes. The risk associated with a lottery that yields $15 with a probability of .5 is greater than the risk associated with a sure payoff of $7.50. In finance, risk is typically used to refer to variance or standard deviation. The risk associated with the stock in the example above is .04. Of course, one can translate between these two notions of risk, and under some limited circumstances, they can be shown to be equivalent. One can compute the variance of the single lottery described above. It has an EV of $7.50 and a variance of 28. One can represent a stock as a lottery set, as I have done in Figure 4. But it is important to keep in mind that there are subtle differences here that must be considered, and under many conditions, these statements are not equivalent.

As a final note, it is probably also important to point out that these traditions are also related mathematically. As we have seen, a chooser who obeys GARP can be described as having a utility function. A chooser who obeys expected utility theory can be described as having a continuous monotone utility function. In a similar way, a chooser who always shows a straight risk premium line can be described even more restrictively as having a continuous monotone utility function drawn from the class of quadratic equations. In this sense, financial theory specifies a more restrictive description of the decision maker than does expected utility theory.

Which Model for Neuroeconomics?

A healthy debate is going on today over which of these two models will prove to be more useful in studies of the brain. One advantage of the utility representation is that it is very closely related to classical psychophysics. We know that both the psychological and neural representations of qualities such as brightness or sweetness are concave and monotonic with regard to objective properties such as number of photons or sugar concentration (see, e.g., Stevens & Stevens, 1961). This may suggest that utility-like representations are abundant in the brain and that concave representations of this type are an essential property of most neural encoders (Glimcher, 2009). One advantage of the financial representation is that it allows for the explicit representation variance and even perhaps for the explicit representation of higher order features of complex distributions. In fact, recent neurobiological studies have begun to suggest that higher order features of complex distributions may be represented in the cortex (Preuschoff, Bossaerts, & Quartz, 2006).

In any case, it is critical that neurobiologists studying risk and valuation be clear about the points that they make. Notions of risk, risk aversion, utility, and variance are well-defined and important concepts that must be treated with care if scholars from other disciplines are to understand the measurements neurobiologists make.

AUTHOR NOTE

The author thanks Kenway Louie and Daniel Burghart for helpful discussions. Correspondence concerning this article should be addressed to P. W. Glimcher, Center for Neuroeconomics, New York University, 4 Washington Place, 809, New York, NY 10003 (e-mail: glimcher@cns.nyu.edu).

REFERENCES

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, **22**, 23-36. (Original work published 1738)

Glimcher, P. W. (2009). Choice: Towards a standard back-pocket model. In P. W. Glimcher, C. F. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 503-522). New York: Academic Press.

Holt, C. A., & Laury, S. K. (2005). Risk aversion and incentives: New data without order effects. *American Economic Review*, **95**, 902-912.

Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica*, **17**, 159-174.

Huygens, C. (1657). De ratiociniis in aleæ ludo. In F. van Schooten and L. Batava's *Exercitationum mathematicarum*. Reprinted (1920) in *Oeuvres complètes de Christiaan Huygens* (Dutch, with French translation).

Markowitz, H. M. (1991). Foundations of portfolio theory. *Journal of Finance*, **26**, 469-477.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. Oxford: Oxford University Press.

Pascal, B. (1966). *Pensées* (A. J. Krailsheimer, Trans.). New York: Penguin. (Original work published 1670)

Preuschoff, K., Bossaerts, P., & Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*, **51**, 381-390.

Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, **2**, 225-243.

Rustichini, A. (2009). Neuroeconomics: Formal models of decision-making and cognitive neuroscience. In P. W. Glimcher, C. F. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 33-46). New York: Academic Press.

Samuelson, P. A. (1938). A note on the pure theory of consumer behavior. *Economica*, **1**, 61-71.

Sharpe, W. F. (1964). Capital asset prices—A theory of market equilibrium under conditions of risk. *Journal of Finance*, **19**, 425-442.

Smith, A. (1976). *An inquiry into the nature and causes of the wealth of nations.* Chicago: University of Chicago Press. (Original work published 1776)

Stevens, J. C., & Stevens, S. S. (1961). Scaling problems in psychophysics. *Acta Psychologica*, **19**, 192-193.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Wu, G., & Gonzalez, R. (1999). Nonlinear decision weights in choice under uncertainty. *Management Science*, **45**, 74-85.

NOTES

1. In fairness, the logarithmic function does have many elegant features that make it desirable for a mathematician and hugely useful in finance.

2. Risky options are here defined as choices that lead to outcomes that are not certain or, more precisely, outcomes that are described by nondegenerate probability distributions.

3. More formally, this is the claim that Y is riskier than X if the distribution of outcomes associated with Option Y can be described as the distribution of outcomes associated with X plus some random variable. It turns out that this is equivalent to saying that if an expected utility-compliant chooser with a concave utility function prefers X to Y, and X and Y have the same expected value, then Y can be defined as riskier than X. Interestingly, neither of these notions of risk are equivalent to the statement "Y has greater variance than X," a point developed by Rothschild and Stiglitz (1970).

4. In fact, highly risk-tolerant individuals can borrow money to place themselves even farther out along this line, but that is a detail beyond our discussion here.

5. However, in fairness, there is now significant evidence that real-world fluctuations in stock prices are not well described as Gaussian, under some conditions.

(Manuscript received June 9, 2008;
revision accepted for publication July 12, 2008.)