

The Black Swan: Why Don't We Learn that We Don't Learn?

Nassim Nicholas Taleb¹

HIGHLAND FORUM 23, Las Vegas, November 2003

FIRST DRAFT, January 2004.

PLEASE DO NOT QUOTE WITHOUT EXPLICIT PERMISSION

¹ Empirica LLC and Courant Institute of Mathematical Sciences, New York University nnt@fooledbyrandomness.org. I thank John Stentbit, Richard O' Neill, Linton Wells, David Albers, Laurence Zuriff, and participants in Highland Forum 23 in Las Vegas for helpful comments. I am grateful to Didier Sornette for generous help, invaluable comments, and helpful discussions particularly with Lévy processes.

ABSTRACT

Does the structure of our mind cause us to underestimate the probability of the Black Swan classes of rare events? Are we made to *not learn rules* (rather than facts) from our past experiences of these events? There is evidence of a *scorn of the abstract* in our harm avoidance mechanism. This discussion presents the Black Swan problem and reviews the insights from a collection of disciplines in the behavioral and cognitive sciences (empirical psychology, neurobiology, evolutionary psychology) on the subject of the perception of harmful outcomes.

ABSTRACT

INTRODUCTION: A BRIEF INTRODUCTION TO THE BLACK SWAN

The three properties of the Black Swan

Causative properties and mislearning from the past

Black Swans and security

A brief presentation Hume's Problem of Induction

Risk, measurement, and uncertainty

Black Swans and "outliers"

Brief power law discussion

Outliers are prevalent in social variables

There are such things as positive Black Swans

*A dual environment representation: type 1 and type 2
environments.*

WHY DON'T WE LEARN META-RULES?

Scorn of the abstract

Behavior v/s cognition

HINTS FROM THE HEURISTICS AND BIASES RESEARCH APPROACH

The human brain's information limitations

Heuristics and Biases

Rationality

Availability and representativeness heuristics

Induction and small probabilities

"I knew it all along": the hindsight bias

Overconfidence in projects and off-model risks

Overconfidence with one's knowledge

THE PSYCHOLOGY AND NEUROBIOLOGY OF EMOTIONS

The affect heuristic and the "risk as feelings" theory

Biological explanations: the emotional brain

Fear and anatomical structure

PROBABILITY FOOLISHNESS AS MALADAPTATION: AN EVOLUTIONARY PERSPECTIVE

What is a module?

Ancestral Environment, Cognition and Preferences


CONCLUSION

REFERENCES

INTRODUCTION: A BRIEF INTRODUCTION TO THE BLACK SWAN

---FIGURE 1 Goes here --

FIGURE 1. *A sample large deviation with nothing in the past history of the process possibly indicating such possibility. Is there a mechanism of inferential myopia that causes us to ignore the fact that these happen routinely where they did not happen in the past?*



This section will present an intuitive non-technical summary of the ideas of the dynamics of the Black Swan process.

The three properties of the Black Swan

The Black Swan is defined here as a random event satisfying the following three properties: large impact, incomputable probabilities, and surprise effect. First, it carries upon its occurrence a disproportionately large impact. The impact being

extremely large, no matter how low the associated probability, the expected effect (the impact times its probability), if quantified, would be significant. Second, its incidence has a small but incomputable probability based on information available prior to its incidence². Third, a vicious property of a Black Swan is its surprise effect: at a given time of observation there is no convincing element pointing to an increased likelihood of the event. By some vicious dialectic, it is the surprise element that either causes the Black Swan or at least exacerbates its consequences. If there were anything convincing about the need for protection then agents would have taken preventive or protective actions. These would have either stopped it in their tracks or limited its impact. Consider that if the WTC attack of September 11 2001 were a plausible risk then planes would have protected New York City and airline pilots would have had locks on their doors. Accordingly its occurrence was linked to its implausibility and its harm to the surprise element; *the more inconceivable the more harmful* the event.

Figure 1 shows a graphical example of a Black Swan, using an unspecified quantitative variable (a physical variable, temperature, blood pressure, the number of citations of a scholar's work, an economic aggregate like Trade Balance or GDP, a price of a financial stock or commodity). Note that the Black Swan is not qualitatively different with economic entities from others in the social sciences (wars, sociological phenomena); economic variables presents the advantage of being quantifiable and lending themselves to testability.

Causative properties and mislearning from the past

These causative properties of the surprise element induce a condition that we cannot directly learn from the past since the

² There is no contradiction between "small" and incomputable in the sense that had the event been of a high rate of occurrence it would no longer cause a high rate of damage.

past has been already incorporated in the expectation of agents leading to a modification of their behavior –simply put, the next time it will cease to be a surprise and the damage will no longer be the same. It is unfortunate that, accordingly, we cannot develop convincing methods to infer the likelihood of a Black Swan from statistical-inductive methods (those based on the observation of the past) and the derivation of the likelihood on a future event based on the findings. Yet statistics is what we seem to resort to instinctively in the social sciences where controlled experiments are generally lacking. Scientific research on rare events, particularly in economics, has been hampered by the condition that after the actual occurrence of the event it becomes suddenly plausible, even obviously likely, by a mechanism called in the psychology literature the *hindsight bias* (a bit more on that later). Only successful methods “out of sample” can illuminate us, not those that fit stories backwards to match the events. While inductive-statistical methods (those relying on the examination of past data) can enlighten us on the properties a collection of physical events, such as the odds of a car crash, the death probability of a drunken retiree, or the success and failure of a cancer treatment protocol, events in the social, nonphysical world seem to elude such analyses.

Black Swans and security

It is difficult to motivate people in the prevention of negative Black Swans, a point this author argued in Taleb (2004). It is due to the lack of observable measures (and, as we will see, the beastly hindsight biases) and the absence of “long term” in the evaluation of performance. Furthermore what is “long” in “long term” is not easy to gauge with these complicated and nonstandard distributions. Prevention is not easily perceived, measured, or rewarded; it is generally a silent and thankless activity. Just consider that a costly measure is taken to stave off such an event. One can easily compute the costs while the results are hard to determine. How can one tell its effectiveness, whether the measure was successful or if it just coincided with no

particular accident? In conventional risks, those of a bell-shape (Gaussian) nature (events that happen more frequently), as we will see, a mere reduction in mortality can suffice to convince us of the effectiveness of the measures taken—not with Black Swans. Job performance assessments in these matters are not just tricky, but may be biased in favor of the observed “acts of heroism”. History books do not account for *heroic* preventive measures.

A brief presentation Hume's *Problem of Induction*

The metaphor Black Swan is historically attributed to the difficulty in epistemology called Hume's Problem of Induction³, of the complication that lie in deriving general rules from observed facts –and from those facts only (see the general discussion in Taleb and Pilpel, 2004). How many white swans does one need to observe before inferring that all swans are white and that there are no black swans? Hundreds? Thousands? The problem is that we do not know where to start –we lack a framework of analysis to know if our *ex ante* estimation is appropriate, which is key. The Black Swan is not just a hypothetical metaphor: until the discovery of Australia common belief held that all swans were white; such belief was shattered with the sighting of the first *cygnus atratus*.

The gravity of the problem lies in our lack of ability to estimate our error rate, except in a circular manner. To compute an error rate one needs a probability distribution; to be confident about the probability distribution one needs an error rate. Such circularity is not appreciated in the literature –indeed it leads to the condition that “induction is inductively justified” (it works because it has worked). More practically, is a large deviation a

³ *That there is nothing in any object, considered in itself, which can afford us a reason for drawing a conclusion beyond it; and, that even after the observation of the frequent or constant conjunction of objects, we have no reason to draw any inference concerning any object beyond those of which we have had experience.* Hume (1748).

“ten sigma” or is it that we have no clue about the probability of the events? The matter becomes of some significance when a single Black Swan can carry large-scale consequences. This leads to severe methodological problems in scientific knowledge.

The scientific methodologist Karl Popper used a similar argument to attack positivism (the prevailing belief early in the century that empirical and “scientific” research can strengthen our knowledge) and downgrade scientific knowledge by rebuilding it on a basis of skepticism. To him confirmation being impossible, we need to just work around the fact with a total reinterpretation of what the notions of corroboration and evidence mean. Now, the major practical concern of this author lies one step beyond: granted that one cannot posit certainty or general rules, we have absolutely nothing upon which to base a model for such uncertainty. We are uncertain about which model of uncertainty to use. Hence methods such as “Pascal’s wager” (avoid holding beliefs that can cause potential harm) seem the only plausible ones. *Skeptical empiricism* corresponds to rules of asymmetry in inference and the use of skepticism as a mode of operation. This paper’s conclusion will discuss the author’s advocacy of lessening the reliance on computable measures.

Note that the Problem of Induction is even more severe in the social sciences where one cannot observe the mechanism generating the process and conduct controlled experiments; what we witness are merely manifestations of it.

Risk, measurement, and uncertainty

This brings us to the distinction between “risk” and “uncertainty”, attributed to the economist Frank Knight¹ who introduced it in the 1920s, to very little follow-up in matters of risk management. The first category is measurable, frequently called “Knightian risk” but not the second, more properly called “Knightian uncertainty”. Note that a collection of other thinkers

have independently examined the topic, most explicitly John Maynard Keynes⁴. The distinction between the two notions of risk and uncertainty is not trivial; it leads to a gap in knowledge between the uncertainty treated by the literature and the one we witness as practitioners operating in the real world.

For instance the economic literature seems to have operated in a state of complete disconnect from the evidence from the data. We have had evidence with papers published since 1963 that economic variables and prices follow “wild uncertainty” (see Mandelbrot, 1963 and Mandelbrot, 1997)ⁱⁱ to no avail.

An illustration is provided with the Long Term Management fiasco where a failure of a fund that included among its founders two major economists with a Nobel Medal. The failure came obviously from a Black Swan but was attributed by the founders to a “ten sigma event”, that is, an event of such small occurrence that it would take place every several times the life of the universe --not to a Black Swan (for which he cannot emit such probabilistic statement).

There has been no follow through in the examination of incomputable risks as an active research went into the exclusive treatment of measurable risk –or considering that the risks they were discussing were measurable. Why? To be cynical, no academic would get rewards for writing a paper explaining that “nobody knows anything”, to popularize an expression attributed to concerning the distribution of returns in the movie business (De Vany, 2003). But the point I will make next is that one does not need an intractable distribution for the problem to occur, simply to be in a situation where computation is not possible for *lack of information* or *insufficient sample*.

⁴ Knight(1921)

⁴ Keynes (1937).

Black Swans and “outliers”

In statistical terms a Black Swan corresponds to the disproportionate contribution of a few observation to the total. It does not equate to the mechanisms described as power laws, but can certainly be caused by them.

Consider that the black swan can simply arise from a simple misunderstanding of the *law of large numbers*: we may live in a world of tractable randomness, *in theory*, but do not know, *in practice*, whether we have sufficient information to base our opinion. A simple underestimation of the total sample size necessary can cause a black swan – “cause” in the sense that an event that could have been expected if we knew the distribution should not be taken into account in the realm of the possible and can cause a surprise.

While much statistical complexity and pseudo-complexity went into the modeling of these events, a few points need to be made clear with the excessive application of the “law of large numbers”. In a simplified way, the law of large numbers indicates that the properties of a sample will converge to a well-known shape after a large number of observations. However the speed of convergence (or lack of it) is not known from the outset. The categories are as follows:

- a- tractable randomness where deviations are bounded; one can gauge a maximum and a minimum (a roulette table).
- b- tractable randomness but one that can, in theory, deliver large deviations (bell-shaped). These have a low

probability, and one that declines at very fast pace (exponentially)⁵.

c- tractable randomness (regular and occasional jumps Poisson-style) but observers are unable to gauge the sufficiency of the total sample size.

d-intractable randomness (Levi-stable, a.k.a. Pareto-Lévy, a.k.a. Pareto-Lévy-Mandelbrot distribution). No sample size will ever be sufficient as the properties will not be captured.

In the third case the law of large number can cause deviations but we will never know what magnitude these can take.

Outliers are considered by modelers in risk management *off-model risks*, those not captured by their models. A compelling example of on-model/off-model risks was discovered by the members of the Highland forum 23. The Casino's where an earlier version of this discussion was presented, the largest one in the world, presented us with their methods of risk management. Random variables affecting a casino's core business are Gaussian (physical variables) and can lend themselves to effective mathematical modeling; they can benefit from the laws of large numbers to diversify away their gambling risks (a large number of small bets, each with positive expected returns for the casino). Their risk management, aside from setting the gambling policies, is geared towards reducing the losses resulting from cheaters. The six largest risks losses incurred or narrowly avoided by the casino were completely outside their risk management model, ranking from a disgruntled contractor attempting to blow up the structure, to the near-loss of a gambling license owing to employee fraud and

⁵ To be more technical these distributions with rapidly declining densities bear the name compact support.

an irreplaceable performer in their major show being maimed by a tiger causing severe financial losses. A back-of the envelope calculation shows that the dollar-value of these off-model swamp the on-model risks by a factor of close to 1000 to 1. Now consider that the perception of risk and, accordingly, the emphasis and resources involved in risk management were in quite inverse proportion.

Brief power law discussion

Power laws is a broad name for a class of distributions with unstable properties that was introduced by Wilfredo Pareto in the late 19th Century without real follow through in the social sciences⁶. His observation of the distribution of wealth (even at the time) revealed that the divergence from the mean of the observed variable did not display the property promoted by Quételet (who brought the Gaussian curve into social science); such distribution shows most people huddling around the average and an exponentially smaller number of people having higher or lower income. Intuitively a power law distribution has the following property: if the power exponent were 2 and we applied it to the example of income, then there would be 4 times more people with a income higher than \$1 million than people with \$2 million. The effect is that there is a very small probability of having an event of an extremely large deviation. More generally, given a deviation x , the incidence of a deviation of a multiple of x will be that multiple to a given power exponent. The higher the exponent the lower the probability of a large deviation. Taking the previous example, if the power exponent were 3, then there would be 8 times more people with a income higher than \$1 million than people with one higher than \$2 million. Note that power laws with higher exponent offer milder properties. But this author views the representation as just

⁶ Such attitude is reminiscent of scholasticism, as we will see in the discussion of neoclassical interpretations of human behavior.

⁶ See Zajdenweber 2000 ; see Albouy, 2002.

pedagogical, if not simplistic: the processes are not easily subsumed into one single exponent⁷; many processes that bear more complexity than power laws are often represented into it⁸.

Benoit Mandelbrot provides a mathematical framework for the distinction (building on more sophisticated versions of the Pareto distributions, a generalization into Levy-stable now popular under the name Pareto-Lévy or Pareto-Lévy-Mandelbrot). These distributions are stable, i.e. unchanged, under addition or, more clearly, exhibit temporal invariance, i.e. they are similar at all scales. Monthly variations for prices are of the same distribution as hourly prices, except that variations are smaller.

To Mandelbrot, random variables in the physical sciences are generally bounded by energy constraints and exhibit a well-behaved brand of uncertainty that lends itself to conventional analyses (such as convergence to the “law of large numbers”). They tend to exhibit the familiar bell-shapes can be called “Gaussian”. The distributions of height, birth weights, the dispersion of heat particles bear in the long run Gaussian properties. In the social sciences, variables have no such constraints; a price can take any value. This he calls “wild uncertainty”. One property of “wild” uncertainty is that the properties do not reveal themselves from past data. Broadly

⁷ Technical note: power laws are generally treated as an asymptotic behavior for a L-Stable (Levy stable) distribution with a specific parametrization (L-Stable do not have closed-form densities). Where “in the tails” they start matters a bit. There exist however some “pure power laws” like the Pareto distribution (one that is a power law with a constant parameter throughout) but it is rarely used in serious as it is hard to fit to the data.

⁸ The problem is that the estimation of the parameters is not easily done and sample size dependent. See Weron (2001) for a discussion of situations where a Levy-Stable process can tend to show a power law with an exponent of 3 or more when the sample size is 10^6 but a lower exponent at higher samples –giving more time to the process to deliver the unpredictable. In other words one can be fooled into reading from the sample that the distribution is a power law with more stability than it shows.

speaking, random variables that “converge” become bell-shaped, others do not and exhibit “fat tails”, i.e. hard to track.

Note again that “power laws” are not probability distributions per se, but correspond to the “tails” or what can be a large deviation. Some distributions become increasingly power laws or “switch” into it at some level called the “critical point” or “tipping point”.

Black Swan does not have to come from a “power law”⁹. As we mentioned, it only needs its probability to be incomputable. Some power laws have incomputable properties but the reverse is not true; many statistical distributions have computable properties but we may not have enough information to compute them.

“Variance of variance” and knowledge about knowledge. Note that in the recent booming literature on networks there is a growing discussion of class of random distributions called “scale free”ⁱⁱⁱ. This author’s approach (Taleb, 1997) is as follows. Consider that the variance is the magnitude of one’s ignorance around an expectation. An unknown variance reflects one’s ignorance of the rate of ignorance. Such layering of variances is *dubbed stochastic variance*, or more commonly stochastic volatility in the financial engineering literature.

Outliers are prevalent in social variables

Consequently, a remarkable feature of our social world^{iv} is that much of what we witness is attributable to the Black Swan classes of events. Consider that the variations in market prices (called “outliers” as their error lie outside the realm of forecasting models) and rare events constitute the bulk of economic growth.

⁹ For the universality of such acutely nonGaussian processes, see Sornette 2001. He sees these processes as pervasive in almost every walk of life.

If “outliers”, i.e. the off-model portion dominates the on-model one, mathematical modeling in the social sciences can be of smaller use than we think. The effect of network externalities is making these clusters of exceptional contributions more severe than the normal rest.

While a man on the street may be excused for not understanding such a point, scientists and researchers need to be put under more severe scrutiny. Now how come financial economists have labored under the assumption that such brand of uncertainty does not exist? Were they living on a different planet or did they just refuse to accept the fact that much of what is taking place is not reflected in their models? Clearly this is a classical element of misdirected research: the building of complicated mathematical models that are elaborate, but, in spite of their sophistication, have a track record similar to, say, astrology –leading to celebrated fiascos such as the one by two Nobel laureates in the Long Term Capital Management blowup of 1998¹⁰. It is obvious that the wild brand of uncertainty reflected in the saying “nobody knows anything”^v does not lend itself to computability. The mathematics simply elude us and *we have no way of knowing what we don't know*. But no “scientist” would be rewarded for not showing quantitative properties. This explains why we missed close to 120 years of research into these: the economist Wilfredo Pareto mentioned above showed in the late 19th century that the distribution of wealth had “fat tail” attributes of which we know nothing.

There are such things as positive Black Swans

It is key here that while we focus in this discussion on harmful effects (I use the designation Black Swan here as shorthand for negative Black Swan), their dynamics are not necessarily harmful; their impact can constitute both positive and negative

¹⁰ See Lowenstein (2000).

extreme deviations. The unexpected can be positive. These phenomena are popular in the common parlance as expressions such as “10% contribute to 90%”; their properties are getting more accentuated where “1% contribute to 99%” (see Frank, 1985, Frank and Cook, 1995 for a popular discussion). Consider that of around the hundreds of thousand of companies that have been publicly listed on United States stock markets, of which we can count 10,000 survivors, less than 100 represent half the current capitalization. The same dynamics applies to movies, book sales (to wit the Harry Potter craze), success of network nodes (the Google effect), popularity of actors, capitalization of companies a few years beyond startup (the Microsoft effect), social and political events (revolutions, coups, wars), financial crises. Recent events that have black swan properties include the World Trade Center attack of September 11, 2001, the stock market crashes of 1987 and 1989, and the technology bubble that started in the middle 1990s¹¹. Further back we have the world wars. The emergence of the United States as a world power would have been considered a lunacy if presented to a “reasonable” person in 1900; the ascent of Hitler laid beyond the ken of people at the time¹². The same applies to technological innovations –they do not correspond to what we expect from them. Yet we continue to extrapolate, never learning.

Note that for the purpose of clarity what we call risk avoidance here in this paper should be interpreted as equivalent of harm avoidance (as the two do not carry the same exact connotation in the economic literature where it is associated with a specific utility preferences).

¹¹ See the popular accounts by Watts (2003), Barabasi(2002) for the recent dynamic research in scale-free non Gaussian networks.

¹¹ De Vany (2003)

¹¹ Note that contrary to popular belief, the technology bubble had more of the properties of the Black Swan than its demise. This perceptual bias corresponds to the humans underestimate the randomness coming from positive events which give the impression of being unpredictable.

¹² I thank Linton Wells for a discussion of such “forward simulation” stripped of hindsight bias.

A dual environment representation: type 1 and type 2 environments.

We define a type 1 environment as one in which the contribution to randomness comes largely from the body (the Gaussian, for instance). A Type 2 environment is one where such contribution comes from a small number of events. As we will see in our discussion of evolution, much of physical randomness is of type 1 –and both our intuitions and the development of our tools (even the scientific ones) have been devoted to that. However we are accumulating enough evidence that the environment in which we live makes our social world increasingly of type 2.

WHY DON'T WE LEARN META-RULES?

Scorn of the abstract

Consider the following quizzical mechanism: we are not good at figuring out Black Swans, yet we don't know it, whether regular citizens or holders of a Nobel Memorial Prize in Economics (although anecdotal evidence indicates that the latter category is more prone to such mistake). How come we don't seem to figure out this apparently obvious and trivial point? How come we do not learn the meta-rule? One would expect that an event of the Black Swan variety, like, say, September 11 2001 would teach us that the incidence of such unpredictable events is higher than we expected before and would force us to adjust for them –in other words that we would learn general harm avoidance rules from the past such as the simple one that Black Swans do not resemble what we thought would be Black Swans. Indeed we learn from the past, but what we tend to learn are not general rules: what we learn is to avoid specific classes of events with the similarity of September 11. Our risk-avoidance mechanism is specific, too specific, and does not seem to accommodate abstract concepts.

We will see the exact mechanisms later with the “risk as feeling” theories; what I mean for now is that the human mechanisms do not seem to accommodate anything abstract like a rule, but requires something vivid and expressive. A special mention needs to be made as this scorn of the abstract is exacerbated by the press. For instance many journalists upon listening to the arguments above asked the author if he could be more “practical”.

Behavior v/s cognition

Indeed, this would not be puzzling if our mind could not handle abstract thought –it just seems that these just do not filter down to our behavior. Is it a surprise that humans endowed with intelligence and ability to solve abstract problems are unable to see obvious risk rules? First, behavior is not necessarily determined by cognition. Second, cognition itself is subjected to distortions and biases. There has been plenty of theories in several branches of the cognitive, neurobiological, and behavioral literature that may provide explanations of such anomalies in the mechanisms of risk avoidance. The rest of this essay will focus on possible explanations for this neglect of the abstract through three broad classes of research traditions in the cognitive sciences 1) the heuristics and biases approach (in empirical psychology), 2) recent neurobiology of emotions and risk avoidance, 3) evolutionary perspectives (mostly in evolutionary psychology and cognitive and computer science). Again this discussion should not be construed as an exhaustive presentation of these research traditions, but focusing of the limitations in our understanding of the risks of rare events.

Note that the point being made in this paper is not that there is a systematic underestimation of rare events, but rather that there is an underestimation of abstract ones.

HINTS FROM THE HEURISTICS AND BIASES RESEARCH APPROACH

The human brain's information limitations

Research by the artificial intelligence pioneer Herbert Simon in the 1950s¹³ shows that the information processing capabilities of the human brain are quite limited (a matter that surprised social science academics more than the regular persons as they were immersed in complicated models of rational choice). The central idea behind the research on “bounded rationality” is that an agent cannot have everything in his mind upon formulating choices and deriving beliefs; there is a need for shortcuts. There is a necessity for the human mind to “satisfice” (the melding of satisfy and suffice) and not endlessly compute optimal solutions. Just consider the costs and the time involved. Such shortcuts are called *heuristics*. Clearly these rules are shortcuts to avoid putting the mind's computational machinery at work in order to solve simple tasks. Psychology researchers Daniel Kahneman and Amos Tversky in the 1970s went beyond the belief held by these initial researchers. For them this need to use shortcuts induced some significant and meaningful biases –to the point of lack of internal consistency. Our brains cause us, among other things, to incur contradictions (one may believe a thing and its opposite); to estimate the probability of a subset B as higher than that of superset B that includes A; to violate rules of transitivity of preferences. For them these *heuristics* were not merely a simplification of rational models, but were different in methodology and category. They called them “quick and dirty” heuristics. There is a dirty part: these shortcuts came with side effects, these effects being the *biases*.

¹³ See Simon, 1987a and Simon 1987b for a review.

Heuristics and Biases

This started an empirical school of research called the “Heuristics and Biases Approach” that attempted to catalogue them – it is impressive because of its empiricism and the experimental aspect of the methods used (Tversky and Kahneman, 1974, Kahneman et al., 1982). Indeed the rigorous experimentation shows an unusual use of near theory-free scientific testing in the social sciences as these tests on biases have been run on different populations in different places. Consider the contrast with the methods used in economics (and much of social science) at the time: reliance on past statistics with all the problems attached to the interpretation of non-repeatable past historical data (and ex-post fitting of explanations and statistical problems called the biases of *data mining*) as compared to running an experiment and the verification of the results, just like physics. It is worth insisting here that the point made by these researchers is not that we make mistakes in probability (a frequent mistake made upon describing the contributions); it is that *these errors are systematic*.

Rationality

The heuristics researchers show series of disturbing departures from the rational model of Man developed in the neoclassical economics and the social sciences (indeed it is a specific definition of a rationality as complications appear when one considers the differences between individual and collective rationality). Indeed there prevailed a particular vision of humans – particularly that rationality in behavior retained by the literature because of its analytical convenience and handedness. Just as no academic project that focused in the “we don't know” in uncertainty could be an acceptable career course for a researcher, so did the assumptions normative rationality prevail as a platform of research and hamper research on risk assessment and control. Rational behavior implied the use of

axioms and the recourse to mathematical complications; “theorems” could be done thanks to the “uniqueness” of solutions –if everyone optimized and behaved in a coherent manner, it would be possible to embark on the project to model human behavior.

We next list the major heuristics. Note that some may partially overlap. The driving force behind them is the notion of quick and dirty shortcut.

Availability and representativeness heuristics

These closely related two heuristics have been studied for the longest and can be responsible for the bulk of the defects in our statistical machinery. They play quite a significant role in the denial of rare events.

The availability heuristic (Tversky and Kahneman 1973) makes people judge the importance and the relative probability of a given event as a function of the ease with which it comes to mind. The more easily available the example, the more likely it will appear to be. The mugging of a relative in Brooklyn or news of a plane crash can be used to estimate the probability of the corresponding risks.

The following experiment can be quite convincing. Imagine the following scenarios, and estimate their probability.

a) A massive flood some-where in America, in which more than 1,000 people die.

b) An earthquake in California, causing massive flooding in which more than 1,000 people die.

Respondents estimated the first event to be less likely than the second (Tversky and Kahneman 1983). An earthquake in

California, however, is a readily imaginable event which greatly increases the availability—hence the assessed probability—of the flood scenario.

People rely on their mental sampling affected by the ease of retrieval. The mental sampling of subjects is the cognitive equivalent of a statistician's method. It is not random and certainly not unbiased. Recency (how long ago did September 11 take place), salience (the news' ability to get our attention) and imaginability (how easy to visualize) are the main factors determining the assigned probability.

Closely related to the availability heuristic, the representativeness heuristic corresponds to humans' tendency to judge the probability that an event belongs to a category based on how representative it is to that category, not how likely it actually is (Kahneman and Tversky, 1972, Kahneman and Frederick, 2002). For instance it leads us to estimate the probability that a person belongs to a particular social group by assessing how similar the person's characteristics are to the "typical" group member's. A feminist-style philosophy student is deemed more likely to be a feminist bank teller than to be just a bank teller. This problem is known as the "Linda problem" (the feminist's name was Linda). Accordingly, people can estimate the risk of a plane crash from terrorism higher than the risk of a plane crash¹⁴.

¹⁴ There have been some research on whether, when agents are supplied with frequencies rather than probabilities some of these heuristics lose their bias. See Gigerenzer (1997), Gigerenzer and Hoppage (1995). Indeed there is a school that investigates "good heuristics" our innate ability to perform complicated calculations not too dissimilar to the bird's ability to perform complicated navigation; however the good side of heuristics is adapted to a more ancestral environment and fail us in a type 2 world.

Induction and small probabilities

Closer to our Black Swan point, these heuristics result in a bias called “belief in the law of small numbers” (Tversky and Kahneman, 1971; Rabin, 2000). This points to a quite severe flaw in human's native statistical inference machinery: people take the last few observations as descriptive of the general distribution (since these are readily available) and, accordingly, are very quick at making general rules. This author discussed (Taleb, 2004) how rare events are even more exacerbated by the effect: if an event deemed to happen every 5 years does not happen in a year, agents will be likely to believe that its incidence is greatly reduced (since the recent past is more easily retrieved in our memory).

Quite worrisome is the condition that people with statistical training seem to readily commit such mistake (the sample in Kahneman and Tversky, 1971, included authors of textbooks of statistical psychology called to make intuitive inferences). We will see further down that agents have a compressed, narrower distribution in their minds than warranted from the data. Adverse events and outliers will be underestimated (unless there is something vivid attached to them).

We stop here and note that in exercises when probabilities are supplied to people during experiments, they overestimate them by a large factor which may on the surface contradict the message behind this essay (Kahneman and Tversky, 1979; Tversky and Kahneman, 1982). In addition, as we will see, the availability heuristic might make them overestimate them. Our case is that probabilities are not visible and agents in real life do not observe them –counter to research in economics and financial risk management where people are assumed to be able to “measure” probabilities and risk. In a relevant modern experiment (Barron and Erev, 2003), it was shown that agents do effectively underestimate small probabilities when calculating them themselves.

The availability heuristic and the belief in the law of small numbers are more severe with statistical processes where the adverse movements happen infrequently, the subject matter of this discussion. Why? Consider that most of the time the history of the process will reveal nothing taking place, leading to undue generalizations that All Swans Are White.

“I knew it all along”: the hindsight bias

Another problem with the perception of calamities can be associated with the hindsight bias. “This time is different” When asked to reflect on events after their occurrence, agents tend to overestimate how much they should have known at the time the event took place. We do not imagine the succession of chronological events in our minds as they actually happened but, literally, running backward. This is similar to classical psychological puzzles where the agent is supplied a picture of seemingly patternless dots that, after some visualization, reveal a dog; once they see the dog it is impossible to look at the graph without seeing the dog again. The effect is that agents overestimate how much they would have known had they not possessed the correct answer or prediction: events which are given an average probability of p percent before they are known to have occurred, are given, in hindsight, probabilities higher than p percent (Fischhoff 1982)—a phenomenon sometimes known as ‘I knew it all along. Both these biases can result from the availability heuristic. Events that took place are now readily and potently available in our minds.

As we discussed earlier, the hindsight bias can carry a severe effect in matters of security and can cause quite a great deal of unfairness in the evaluation of a given contribution. After the events of September 11 2001, it was obvious that things *should have been done* to counter them and that there were some “lapses”. The only way to do it accurately is to estimate the rate of lapses *compared* to those related to events that did not happen. Should an earthquake take place in San Francisco

(another case of incomputable probability, see Freedman and Stark, 2003) then it would be obvious *after the fact* that living there was not a reasonable choice and that residents of the area were the victims of lapses of judgment given what it was obvious that an earthquake would have to take place there.

Overconfidence in projects and off-model risks

Agents overestimate their skills owing to attribution bias. There is an ingrained asymmetry in the way events are perceived. Individuals ascribe their past failings to random events, but their successes to their skills. The consequence is that their projection¹⁵ of the space of eventualities will be rosy and they will underestimate the incidence of possible setbacks. There has been extensive investigation of the *planning fallacy*: why do people consistently carry rosy projections? Projects are rarely finished before the projected completion date and hardly ever below the projected budget. The literature frequently cites examples of excessive costs overruns and delays by the Sydney Opera House, Denver International Airport. Yet planners do not seem to learn from general history, not even their own.

One explanation of these failure, *focalism* can account for the mental elimination of off-model risks. People upon formulating a project eliminate factors lying outside the specifics of the project itself¹⁶.

People are unaware of their own track record and do not learn that their past projections were too optimistic and correct for it. There are always small Black Swans hampering projects; these do not seem to be taken into account. The “scenario analysis” style projections by economic agents hardly has allowance for these Black Swans.

¹⁵ See the review in Bueler, Griggin and Ross, 2002.

¹⁶ See Kahneman and Lovallo, 1993.

Overconfidence with one's knowledge

People over-estimate how much they actually know; studies have provided a quite extreme account of these. When they are p percent sure that they have answered a question correctly or predicted correctly, they are in fact right on average less than p percent of the time (e.g., Lichtenstein et al. 1982). Alpert and Raiffa (1982) studies have documented how agents underestimate the extreme values of a distribution in a surprising manner; violations are far more excessive than one would expect: events that are estimated to happen less than 2% of the time will take place up to 49%. There has been since a long literature on overconfidence (in the sense of agents discounting the probability of adverse events while engaging in a variety of projects), see Hilton (2003).

THE PSYCHOLOGY AND NEUROBIOLOGY OF EMOTIONS

The initial work by researchers in the Heuristics and Biases approach overlaps with that of modern neurobiology, particularly in the treatment of emotions. This section will merge the two approaches.

The affect heuristic and the "risk as feelings" theory

Affect can be loosely defined as "feeling" (we will get into the biological mechanisms later); cognition just as loosely defined as "thinking". The two are mechanisms that seem to either compete or complement each other in information processing. More specifically, affect is a more technical term used by psychologists to define emotion and desire (more broadly, to our desires and aversions, pleasures and pains); its handiness as a concept lies in that it can be positive or negative. "Cognition," in its traditional sense, refers to processes involved in the acquisition of knowledge (thinking, attention, memory). Zajonc (1980, 1984) introduces the mechanisms of competition between affect and

cognition as shown in the title of the influential paper that sparked a debate “*Feeling and thinking. Preferences need no inferences.*” Intuitively, one does not see just a tree, but a beautiful or ugly tree. Try to look at an object without it eliciting some reaction. Is it possible?

The idea put forth in Zajonc (1980, 1984) is that affective processing does not depend on controlled cognitive processing. That is, organisms are able to determine whether a stimulus is “good” or “bad” without engaging in intentional, goal-directed, conscious, or capacity demanding type of processing. Such automatic affective processing was believed to have an important impact on *subsequent* cognitive processing and behavior.

Psychologists later refined the problem by framing it as a dual system of processing by the brain. The system 1/ system 2 representation (Sloman, 1996; Sloman, 2002) gives the following clues. System 1 is self aware, effortful, reasoned, non-heuristic. By comparison, System 2 is where the heuristics (and their distortions) reside. It is opaque, effortless, automatic, and highly emotional.

This brings us to the affect heuristic; similar to the availability one, seems to make us gauge the likelihood of an event according to how strongly it affects our emotional well-being.

An intuitive presentation of the emotional computation of probability is well known in cancellation of airline trips after a plane crash. At the time of writing, people seem to worry more about the mad cow disease than general food poisoning. Probability assessments are determined by how emotionally powerful the outcome of a gamble. Slovic et al., 2002, write:

“The evaluability of a stimulus image is reflected in the precision of affecting feelings associated with that image. More precise affective impressions reflect more precise

meanings and carry more weight in impression formation, judgment, and decision making”.

Researchers figured out that in gambling, say with lottery tickets, the total stake matters far more than its probability. This leads to the pathology that ones images and feelings toward winning the lottery are likely to be similar whether the probability of winning is 1 in 10 million or 1 in 10,000.

A popularization of it with the “risk as feelings” theory. It seems that risk avoidance is not cognitive, mediated in the abstract part of the brain, but needs some specific visual association and emotional salience. Clearly there has to be processes in us that are automatic. The notions of opacity and automaticity of the process are key.

Research on the role of emotions on the brain has been popularized since the 1990s. It is to these discussions that we turn next.

Biological explanations: the emotional brain

Consider the well known mechanism of phobia where one's cognitive machinery accepts that overly worrying about a remote or nonexistent risk is irrational (say being tormented about contracting a deadly disease from shaking hands) yet one's emotional apparatus and behavior deliver an opposite message. One feels the emotions associated with risk but not cognition. This example of phobias illustrates a survival mechanism that went amok. People subjected to it are easily branded as “irrational” but now just think of the exact reverse situation: risks that we know do exist and are significant yet somehow do not seem to affect our behavior (see Berridge, 2003 for the difference between expected utility and decision utility). Smoking is a prime example. Such category of inverse phobias does not seem to have attracted much attention in the literature. There is no disease treatment for Black-Swan complacency.

Indeed we have been accumulating evidence for a long time to the effect that risk avoidance is not completely mediated in the rational part of the brain –hence the process does not result from reasoning but from other opaque factors that do not appear to affect our introspection. This brief discussion presents the main ideas on risk avoidance in the literature.

There has been a growth in the 1990s of work by neurobiologists on the role of emotions in decision-making. Damasio (1994) discusses the case of a patient who, having lost the ability to experience emotions, incurred severe degradation in his decision-making process. He was unable to make an appointment as he was lost between alternatives –Damasio's work provides a beautiful biological justification to Simon's "satisficing" theory. Cognition alone cannot help.

Fear and anatomical structure

Joseph Ledoux was among the earliest to study the localization of fear and the role of the limbic system that we share with other mammals in the risk avoidance process –to the point of overwhelming the cognitive system. Indeed to Joseph Ledoux the pathways from the emotional to the cognitive system are much wider than those from the cognitive to the emotional ones. What does it mean? That in many cases we experience an emotion, then fit an explanation to it. The connectivity of the amygdala¹⁷ with the neo-cortex is not symmetrical.

“The amygdala projects back to the neo-cortex in a much stronger sense than the neo-cortex projects to the amygdala. David Amaral has made this point from studies of primate brains. The implication is that the ability of the amygdala to control the cortex is greater

¹⁷ For the role of the amygdala, see Morris, Ohman and Dolan ,1998, Morris, Ohman and Dolan 1999, Ledoux, 1993, Ledoux, 2000.

than the ability of the cortex to control the amygdala. And this may explain why it's so hard for us to will away anxiety; emotions, once they're set into play, are very difficult to turn off. Hormones and other long-acting substances are released in the body during emotions. These return to the brain and tend to lock you into the state you're in at the time. Once you're in that state it's very difficult for the cortex to find a way of working its way down to the amygdala and shutting it off¹⁸.”

Patients with damage in their ventromedial frontal cortex seem to retain cognitive skills unimpaired, yet become erratic gamblers. Does it mean that underestimating the odds? Bechara, Damasio, Damasio, and Tranel (1994) show the degradation of the risk-avoidance behavior of patients with damage in their ventromedial frontal cortex, a part of the brain that links us to our emotions. While these patients have their cognition unimpaired, their decision making under uncertainty goes out of control.

The famous century-old Claparade's experiment enlightens us to the consciousness of the process. There is a famous case of a patient with severe amnesia who was unable to identify her doctor unless reintroduced to him every five minutes. One day he pricked her with a needle upon shaking hands with her. The following time she still did not recognize him but avoided shaking his hand (see Ledoux, 1998).

All thinkers now agree that emotions have an important role. They have been called “lubricants of reason” and play a major role in survival. The problem is that they can just as easily become misdirected. This links us to the next section on maladaptation.

¹⁸ Interview with John Brockman, www.edge.org

PROBABILITY FOOLISHNESS AS MALADAPTATION: AN EVOLUTIONARY PERSPECTIVE

Very little has been written on complicated probability distributions and evolution—evolutionary theory has only recently started to integrate more complicated probability models. Yet naively one can see biology as a series of power laws or complex systems owing to the dependence of organisms on each other's actions. If one were to put a scale of 0 to 10 in tractability of the probability distribution, with physics at 9 and economics at 1, biology would rank somewhere around 4. Consider that there are recursive mechanisms in evolution, the co-dependence between a given environment that prevents purely static analyses.

What is a module?

Steven Pinker's quote¹⁹ “our brains are made for fitness not for truth” summarizes the orientation of the research in this area. The paradigm used in evolutionary psychology is the “Swiss Army Knife” approach to the brain. In place of a central processing, here again our mind is viewed as a series of special-task “modules” that evolved to solve specific functions. Just as we saw with heuristics, modularity is the equivalent of a rapid and efficient system of processing, but with one difference: it is deemed to be biological.

A module is a specialized mental organ that has evolved to handle specific tasks and specific information. The principal properties of a module is encapsulation (we cannot interfere with its functioning), unconsciousness, and speed. Cognitive impenetrability: what is clear to us is that their efficiency does not let them lend themselves to much introspection. It is that property that is quite critical.

¹⁹ Pinker (1997).

Noam Chomsky and Jerry Fodor in the 1960s waged a complete intellectual war with the social science establishment about the efficiency of these modules—these modules need to be innate for the large part or, at least, have an innate biological structure. Chomsky showed it with linguistic acquisition: the mathematics of language acquisition are too complicated for them not to come from some innate organ in our brain. We can see some similarity between language acquisition and risk bearing. Causal thinking has been investigated in an evolutionary framework: we tend to make strong inferences about the physical habitat we live in (see Plotkin 1998). Experiments on infants clearly show, just as we do with language acquisition, they have an ability to infer causal structures *without knowing what a cause really is*. Infants develop language in a similar innate manner.

Ancestral Environment, Cognition and Preferences

The term ancestral “environment of evolutionary adaptedness” (EEA) was initially used to describe a habitat in equilibrium with human genes (Burnham, 2003²⁰, Tooby and Cosmides, 1989). The Pleistocene era, which lasted from around 2 million to close to 10,000 years ago, is generally considered the stable environment in which we developed the traits that we have today.

In his discussion of the *caveman economics* agenda, Burnham (2003) writes:

The caveman economics hypothesis claims that, in spite of human cultural transmission and behavioral flexibility, the ancestral environment selected genes that impose significant constraints on modern economic behavior. If human preferences evolved by natural selection, and the

²⁰ More modern variations, the “adaptedly relevant environment) ARE breaks down adaptations by traits, each corresponding to a different period

genetic mismatch hypothesis is true, then preferences were shaped in ancestral environments. Attempts to endogenize preferences should, therefore, model the ancestral and not modern world. This is the agenda for caveman economics.

Now consider the following properties of the Pleistocene:

- 1) Small groups, reduced information demands, no complications of sedentary life
- 2) Gaussian physical probabilities; the surviving population did not incur large-scale physical Black Swan

The sociobiologist (and pioneer of the field) E.O. Wilson (Wilson, 2002) looks beyond the Pleistocene at the Paleolithic. In a pessimistic discussion of the future of life in our planet, he writes:

The human brain evidently evolved to commit itself emotionally only to a small piece of geography, a limited band of kinsmen, and two or three generations into the future. To look neither far ahead nor far afield is elemental in a Darwinian sense. We are innately inclined to ignore any distant possibility not yet requiring examination. It is, people say, just good common sense. Why do they think in this shortsighted way?

The reason is simple: it is a hardwired part of our Paleolithic heritage. For hundreds of millennia, those who worked for short-term gain within a small circle of relatives and friends lived longer and left more offspring - even when their collective striving caused their chiefdoms and empires to crumble around them. The long view that might have saved their distant descendants required a vision and extended altruism instinctively difficult to marshal.

In an unrelated discussion on the development of brain size from sexual selection, Miller (2000) has a comment on our ability to capture rare events. He writes:

Evolution has no foresight. It lacks the long-term vision of drug company management. A species can't raise venture capital to pay its bills while its research team [...] Each species has to stay biologically profitable every generation, or else it goes extinct. Species always have cash-flow problems that prohibit speculative investments in their future. More to the point, every gene underlying every potential innovation has to yield higher evolutionary payoffs than competing genes, or it will disappear before the innovation evolves any further. This makes it hard to explain innovations.

CONCLUSION

Recall our discussion of the type1/type 2 environments; it seems that evolution prepared us to a type 1 world –and perhaps too well at that. Just as we do not seem to tolerate a modern lifestyle (in the physical equivalent of the type 2 world) owing to our inability to process sugars, so we do not seem to be adapted to process the spate of Black Swan risks that we face; our risk avoidance machinery is not adapted to it.

There is very little discussion in the scientific literature on cures and palliatives for the severe behavioral problems when facing severe random outcomes. Clearly we do not understand the Black Swan. Will we ever do so? Will education help? To use the same analogy to nutrition: complacency towards this class of

risks does not seem to be cured by mere awareness or cognitive action.

The only hint: We know that, owing to the availability heuristic the media distorts the risk perception of individuals, but it is unclear how to correct these biases. We need more. What? Regulation? Force the media to educate the public about statistical significance? How about warning labels?

This calls for further discussion of measures to counter such risk-foolishness –these stand outside the specialty of this author.

REFERENCES

1. Albouy, François-Xavier, 2002, *Le temps des catastrophes*, Paris: Descartes & Cie
2. Arrow, Kenneth., 1987, "Economic Theory and the Postulate of Rationality," in Eatwell, J., Milgate, M., & Newman, P. (eds.), 1987, *The New Palgrave: A Dictionary of Economics*, vol. 2, 69-74, London: Macmillan
3. Arthur, Brian W., 1994, *Increasing Returns and Path Dependence in the Economy*, Ann Arbor: University of Michigan Press
4. Barabási, Albert-László, 2002, *Linked: The New Science of Networks*, Boston: Perseus Publishing
5. Barron, G., and Erev, I. ,2003, "Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions", *Journal of Behavioral Decision Making*, 16, 215-233.
6. Bechara, A., Damasio, A. R., Damasio, H. and Anderson, S. W., 1994. "Insensitivity to future consequences following damage to human prefrontal cortex." *Cognition*, 50:1-3, pp. 7-15.
7. Berridge, Kent C., 2003, *Irrational Pursuits: Hyper-Incentives from a Visceral Brain*, in Brocas & Carillo
8. Bouvier, Alban (ed.), 1999, *Pareto aujourd'hui*, Paris: Presses Universitaires de France
9. Brocas, I. & Carillo J., (eds.), 2003, *The Psychology of Economic Decisions: Vol 1: Rationality and Well-Being*, Oxford: Oxford University Press
10. Brock, W. A. & De Lima, P. J. F., 1995, "Nonlinear Time Series, Complexity Theory, and Finance", University of Wisconsin Madison – Working Papers 9523
11. Buehler, R., Griffin, D. & Ross, M., 2002, "Inside the Planning Fallacy: The Causes and Consequences of Optimistic Time Predictions" in Gilovich, Griffin & Kahneman

12. Camerer, C., Loewenstein, G. & Prelec, D., 2003, "Neuroeconomics: How neuroscience can inform economics", Caltech Working Paper.
13. Conlan, Roberta (ed.), 1999, *States of Mind: New Discoveries About How Our Brains Make Us Who We Are*, New York: Wiley
14. Damasio, Antonio, 1994, *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Avon Books
15. Damasio, Antonio, 2000, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, New York: Harvest Books
16. De Vany, Arthur, 2003, *Hollywood Economics: Chaos in the Movie Industry*, London: Routledge
17. Finucane, M. L., Alhakami, A., Slovic, P. & Johnson, S. M., 2000, "The affect heuristic in judgments of risks and benefits", *Journal of Behavioral Decision Making*, 13, 1-17.
18. Fischhoff, Baruch, 1982, "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight", in Kahneman, Slovic & Tversky
19. Fodor, Jerry A., 1983. *The Modularity of Mind: An Essay on Faculty Psychology*, Bradford Books, MIT Press, Cambridge, Massachusetts.
20. Frank, R. H. & Cook, P. J., 1995, *The Winner-Take-All Society: Why the Few at the Top Get So Much More Than the Rest of Us*, New York: Free Press
21. Frank, Robert H., 1985, *Choosing the Right Pond: Human Behavior and the Quest for Status*, Oxford: Oxford University Press
22. Freedman, D. A. & P. B. Stark, 2003, What Is The Chance Of An Earthquake? Department of Statistics University of California Berkeley, CA 94720-3860 Technical Report 611. September 2001; Revised January 2003.
23. Gehring, W. J. & Willoughby, A. R., 2002, "The Medial Frontal Cortex and the Rapid Processing of Monetary Gains and Losses", *Science*, 295, March

24. Gigerenzer G., Todd, P. M. & ABC Research Group, 2000, *Simple Heuristics That Make Us Smart*, Oxford: Oxford University Press
25. Gigerenzer, G., Czerlinski, J. & Martignon, L., 2002, "How Good are Fast and Frugal Heuristics?", in Gilovich, Griffin & Kahneman
26. Gigerenzer, G. 1997. Ecological intelligence: An adaptation for frequencies. *Psychologische Beiträge* 39: 107–125.
27. Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review*, 102, 684-704
28. Gilbert, D., Pinel, E., Wilson, T. D., Blumberg, S. & Wheatley, T., 2002, "Durability Bias in Affective Forecasting" in Gilovich, Griffin & Kahneman
29. Gillies, Donald, 2000, *Philosophical Theories of Probability*, London: Routledge
30. Gilovich, T., Griffin, D. & Kahneman, D. (eds.), 2002, *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge: Cambridge University Press
31. Gilovich, T., Vallone, R. P. & Tversky, A., 1985, "The hot hand in basketball: On the misperception of random sequences", *Cognitive Psychology*, 17, 295-314
32. Glimcher, Paul, 2002, *Decisions, Uncertainty and the Brain: The Science of Neuroeconomics*, Cambridge: MIT Press.
33. Granger
34. Hilton, Denis, 2003, "Psychology and the Financial Markets: Applications to understanding and remedying irrational decision-making" in Brocas and Carillo
35. Hsee, C. K., & Rottenstreich Y.R., 2004, "Music, Pandas and Muggers: On the Affective Psychology of Value". Forthcoming, *Journal of Experimental Psychology*.
36. Hume, David, 1999, (1748), *An Enquiry Concerning Human Understanding*, Oxford: Oxford University Press
37. Johansen A. & Sornette, D., 1999, Stock market crashes are outliers, *European Physical Journal B* 1, 141-143 (1998)

38. Johansen A. & Sornette, D., 2001, Large Stock Market Price Drawdowns Are Outliers, *Journal of Risk* 4(2), 69-110, Winter 2001/02
39. Kahneman D. & Tversky, A. (eds.), 2000, *Choices, Values, and Frames*, Cambridge: Cambridge University Press
40. Kahneman D. & Tversky, A., 1972, "Subjective probability: A judgment of representativeness", *Cognitive Psychology*, 3, 430-454
41. Kahneman D. & Tversky, A., 1973, "On the psychology of prediction", *Psychological Review*, 80: 237-51
42. Kahneman, D. & Lovallo, D., 1993, "Timid choices and bold forecasts: A cognitive perspective on risk-taking", *Management Science*, 39, 17-31
43. Kahneman, D. & Tversky, A., 1979, "Prospect Theory: An analysis of decision under risk", *Econometrica*, 47, 263-291
44. Kahneman, D. & Tversky, A., 1982b, "On the study of statistical intuitions", *Cognition*, 11: 123-141
45. Kahneman, D. & Tversky, A., 1996, "On the reality of cognitive illusions", *Psychological Review*, 103, 582-591
46. Kahneman, D., Diener, E. & Schwarz, N. (eds), 1999, *Well-being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation
47. Kahneman, D., Knetsch, J. L. & Thaler, R. H., 1991, "Anomalies: The endowment effect, loss aversion, and status quo bias", in Kahneman and Tversky (2000)
48. Kahneman, D., Knetsch, J. L. & Thaler, R. H., 1986, "Rational Choice and the Framing of Decisions," *Journal of Business*, Vol. 59 (4) pp. 251-78.
49. Kahneman, D., Slovic, P. & Tversky, A. (eds.), 1982, *Judgment under Uncertainty : Heuristics and Biases*, Cambridge: Cambridge University Press
50. Kahneman, D. & Frederick, S., 2002, "Representativeness Revisited: Attribute Substitution in Intuitive Judgment", in Gilovich, Griffin & Kahneman
51. Keynes, John Maynard (1937), *The General Theory*. In *Quarterly Journal of Economics*, Vol. LI, 209-233.

52. Keynes, John Maynard ,1937, "The General Theory". In Quarterly Journal of Economics, Vol. LI, 209-233.
53. Knight, Frank,1921 (1965), *Risk, Uncertainty and Profit*, Harper Torchbook Edition, New York: Harper and Row.
54. Knight, Frank,1921 (1965), *Risk, Uncertainty and Profit*, Harper Torchbook Edition, New York: Harper and Row.
55. LeDoux, J. E. (1993). Emotional memory systems in the brain, Behavioural Brain Research, 58, 69-79
56. LeDoux, J. E. (2000). Emotion circuits in the brain. Annual Review of Neuroscience, 23, 155-184.
57. Ledoux, Joseph, 1998, *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, New York: Simon & Schuster
58. Ledoux, Joseph, 2002, *Synaptic Self: How Our Brains Become Who We Are*, New York: Viking
59. Lichtenstein, S., Fischhoff, B. & Phillips, L., 1977, *Calibration of Probabilities: The State of the Art*, in Kahneman, Slovic & Tversky (1982)
60. Lichtenstein, S., Fischhoff, B., & Phillips L. D. (1982). Calibration of probabilities: The state of the art in 1980. In D. Kahneman, P. Slovic,&A. Tversky (Eds.), *Judgment under uncertainty: Heuristics andbiases*. New York: Cambridge University Press.
61. Loewenstein, G. F., Weber, E. U., Hsee, C. K. & Welch, E. S., 2001, "Risk as feelings", *Psychological Bulletin*, 127, 267-286.
62. Lowenstein, Roger, 2000, *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, New York: Random House
63. Mandelbrot , Benoit B., 1997, *Fractals and Scaling in Finance*, New York: Springer-Verlag.
64. Miller, Geoffrey F., 2000, *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*, New York: Doubleday

65. Morris, J.S., Ohman, A. & Dolan, R.J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, 393 (6684), 467-470.
66. Morris, J.S., Ohman, A. & Dolan, R.J. (1999). A subcortical pathway to the right amygdala mediating “unseen” fear. *Proceedings of the National Academy of the Sciences*, 96, 1680-1685.
67. Pinker, Steven, 1997, *How the Mind Works*, New York: W.W. Norton
68. Plotkin, Henry, 1998, *Evolution in Mind: An Introduction to Evolutionary Psychology*, Cambridge: Harvard University Press
69. Rabin, Mathew, 2000, “Inference by Believers in the Law of Small Numbers”, Economics Department, University of California, Berkeley, Working Paper E00-282, <http://repositories.cdlib.org/iber/econ/E00-282>
70. Simon, Herbert A., 1987a, “Bounded rationality”, in Eatwell, J., Milgate, M. & Newman, P. (eds.), 1987, *The New Palgrave: A Dictionary of Economics*. London: Macmillan
71. Simon, Herbert A., 1987b, “Behavioral economics”, in Eatwell, J., Milgate, M. & Newman, P. (eds.), 1987, *The New Palgrave: A Dictionary of Economics*. London: Macmillan
72. Sloman, Steven A., 2002, “Two Systems of Reasoning” in Gilovich, Griffin & Kahneman
73. Sloman, Steven A., 1996, “The empirical case for two systems of reasoning”, *Psychological Bulletin*, 119, 3-22. [See link for reasons for suggested change.]
74. Slovic, P., Finucane, M., Peters, E. & MacGregor, D. G., 2002, “The Affect Heuristic”, in Gilovich, Griffin & Kahneman
75. Slovic, P., Finucane, M., Peters, E. & MacGregor, D. G., 2003a, “Rational Actors or Rational Fools? Implications of the Affect Heuristic for Behavioral Economics”, working paper, www.decisionresearch.com

76. Slovic, P., Finucane, M., Peters, E. & MacGregor, D. G., 2003b, "Risk as Analysis, Risk As Feelings: Some Thoughts About Affect, Reason, Risk, and Rationality", Paper presented at the Annual Meeting of the Society for Risk Analysis, New Orleans, Louisiana, December 10, 2002
77. Slovic, Paul, 1987, "Perception of Risk", *Science*, 236, 280-285
78. Slovic, Paul, 2000, *The Perception of Risk*, London: Earthscan Publications
79. Sornette, Didier, 2002, Predictability of catastrophic events: material rupture, earthquakes, turbulence, financial crashes and human birth, Proceedings of the National Academy of Sciences USA
80. Sornette, Didier, 2003, *Why Stock Markets Crash: Critical Events in Complex Financial Systems*, Princeton: Princeton University Press
81. Sornette, Didier, 2004 (2000), *Critical Phenomena in Natural Sciences : Chaos, Fractals, Selforganization, and Disorder: Concepts and Tools*, Heidelberg: Springer Verlag
82. Stanovich, K. & West, R., 2000, "Individual Differences in Reasoning: Implications for the Rationality Debate", *Behavioral and Brain Sciences*, 23, 645-665
83. Taleb, Nassim Nicholas, 1997, *Dynamic Hedging: Managing Vanilla and Exotic Options*, New York: Wiley.
84. Taleb, Nassim Nicholas, 2004 (2001) 2nd ed., *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*, New York & London: Thomson Texere.
85. Taleb, Nassim Nicholas, 2004, "Bleed or Blowup? Why Do We Prefer Asymmetric Payoffs ?", forthcoming, *Journal of Behavioral Finance*, 5
86. Tversky, A. & Kahneman, D., 1973, "Availability: A heuristic for judging frequency and probability", *Cognitive Psychology*, 5, 207-232
87. Tversky, A. & Kahneman, D., 1982, Evidential Impact of Base-Rates, in *Kahneman, Slovic & Tversky*, 153-

- 160
Voit, Johannes, 2001, *The Statistical Mechanics of Financial Markets*, Heidelberg: Springer
88. Tversky, A. & Kahneman, D., 1992, "Advances in Prospect Theory: Cumulative Representation of Uncertainty", *Journal of Risk and Uncertainty* 5, 297–323
89. Tversky, A. and Kahneman, D., 1971, "Belief in the Law of Small Numbers", *Psychology Bulletin*, Aug. 76(2), 105-10
90. Watts, Duncan, 2003, *Six Degrees: The Science of a Connected Age*, New York: W.W Norton
91. Wilson, Edward O., 2002, *The Future of Life*, New York: Knopf
92. Zajdenweber, Daniel, 2000, *L'économie des extrêmes*, Paris: Flammarion
93. Zajonc, R.B. ,1980, "Feeling and thinking. Preferences need no inferences". *American Psychologist* ,35, 151–175.
94. Zajonc, R.B.,1984, "On the primacy of affect". *American Psychologist* , 39, 117–123.114
