Psychologists at the Gate: Review of Daniel Kahneman's Thinking, Fast and Slow

Citation

Shleifer, Andrei. 2012. Psychologists at the gate: review of Daniel Kahneman's Thinking, Fast and Slow. *Journal of Economic Literature* 50(4): 1080-1091.

Published Version

[dx.doi.org/10.1257/jel.50.4.1080](https://doi.org/10.1257/jel.50.4.1080)

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:10735580>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Psychologists at the Gate:

Review of Daniel Kahneman's *Thinking, Fast and Slow*

Andrei Shleifer¹

August 2012

The publication of Daniel Kahneman's book, *Thinking, Fast and Slow*, is a major intellectual event. The book summarizes, but also integrates, the research that Kahneman has done over the past 40 years, beginning with his path-breaking work with the late Amos Tversky. The broad theme of this research is that human beings are intuitive thinkers and that human intuition is imperfect, with the result that judgments and choices often deviate substantially from the predictions of normative statistical and economic models. This research has had a major impact on psychology, but also on such diverse areas of economics as public finance, labor economics, development, and finance. The broad field of behavioral economics – perhaps the most important conceptual innovation in economics over the last 30 years – might not have existed without Kahneman and Tversky's (hereafter, KT's) fundamental work. It certainly could not have existed in anything like its current form. The publication of Kahneman's book will bring to an even broader audience of economists some of the most innovative and fundamental ideas of 20th century social science.

In this review, I discuss some broad ideas and themes of the book. Although it would be relatively easy to carry on in the spirit of the first paragraph, constrained only by my limited vocabulary of adjectives, I will seek to accomplish a bit more. First, because the book mentions few economic

¹ Department of Economics, Harvard University. I have benefitted from generous comments of Nicholas Barberis, Pedro Bordalo, Thomas Cunningham, Nicola Gennaioli, Matthew Gentzkow, Owen Lamont, Sendhil Mullainathan, Josh Schwartzstein, Jesse Shapiro, Tomasz Strzalecki, Dmitry Taubinsky, Richard Thaler, and Robert Vishny. They are not, however, responsible for the views expressed in this review. I do not cite specific papers of Kahneman when the material is described in the book.

applications, I will describe some of the economic research that has been substantially influenced by this work. My feeling is that the most profound influence of KT's work on economics has been in finance, on what has now become the field of behavioral finance taught in dozens of undergraduate and graduate economics programs, as well as at business schools. I learned about KT's work in the 1980s as a graduate student, and it influenced my own work in behavioral finance enormously.

Second, I believe that while KT's work has opened many doors for economic research, some of the fundamental issues it raised remain work in progress. I will thus discuss what Kahneman's work suggests for decision theory, primarily as I see it through the lens of my recent work with Nicola Gennaioli and Pedro Bordalo (Gennaioli and Shleifer 2010, Bordalo, Gennaioli, and Shleifer 2012a,b,c).

Before turning to the book, let me briefly address the two common objections to the introduction of psychology into economics, which have been bandied around for as long as the field has existed. The first objection holds that, while psychological quirks may influence individual decisions at the boundary, the standard economic model describes first order aspects of human behavior adequately, and economists should focus on "first order things" rather than quirks. Contrary to this objection, Della Vigna (2010) summarizes a great deal of evidence of large and costly errors people make in important choices. Let me illustrate. First, individuals pay large multiples of actuarially fair value to buy insurance against small losses, as well as to reduce their deductibles (Sydnor 2010). In the standard model, such choices imply astronomical levels of risk aversion. Second, the standard economic view that persuasion is conveyance of information seems to run into a rather basic problem that advertising is typically emotional, associative, and misleading – yet nonetheless effective (Bertrand et al. 2010, Della Vigna and Gentzkow 2010, Mullainathan et al. 2008). Third, after half a century of teaching by financial economists that investors should pick low-cost index funds, only a minority do, while most select high cost actively managed funds that underperform those index funds. These kinds of behavior matter for

both prices and resource allocation. Explaining such behavior with the standard model is possible, but requires intellectual contortions that are definitely not “first order.”

The second objection holds that market forces eliminate the influence of psychological factors on prices and allocations. One version of this argument, made forcefully by Milton Friedman (1953) in the context of financial markets, holds that arbitrage bring prices and therefore resource allocation to efficient levels. Subsequent research has shown, however, that Friedman’s argument – while elegant – is theoretically (and practically) incorrect. Real-world arbitrage is costly and risky, and hence limited (see, e.g., Grossman and Miller 1988, DeLong et al 1990, Shleifer and Vishny 1997). Dozens of empirical studies confirm that, even in markets with relatively inexpensive arbitrage, identical, or nearly identical, securities trade at different prices. With costlier arbitrage, pricing is even less efficient.

A second version of the “forces of rationality” objection holds that participants in real markets are specialists invulnerable to psychological quirks. John List’s (2003) finding that professional baseball card traders do not exhibit the so-called endowment effect is supportive of this objection. The problem with taking this too far is that individuals make lots of critical decisions – how much to save, how to invest, what to buy – on their own, without experts. Even when people receive expert help, the incentives of experts are often to take advantage of psychological biases of their customers. Financial advisors direct savers to expensive, and often inappropriate, products, rather than telling them to invest in index funds (Chalmers and Reuter 2012, Gennaioli et al. 2012). Market forces often work to strengthen, rather than to eliminate, the influence of psychology.

I. System 1 and System 2.

Kahneman's book is organized around the metaphor of System 1 and System 2, adopted from Stanovich and West (2000). As the title of the book suggests, System 1 corresponds to thinking fast, and System 2 to thinking slow. Kahneman describes System 1 in many evocative ways: it is intuitive, automatic, unconscious, and effortless, it answers questions quickly through associations and resemblances, it is non-statistical, gullible, and heuristic. System 2 in contrast is what economists think of as thinking: it is conscious, slow, controlled, deliberate, effortful, statistical, suspicious and lazy (costly to use). Much of KT's research deals with system 1 and its consequences for decisions people make. For Kahneman, System 1 describes "normal" decision making. System 2, like the US Supreme Court, checks in only on occasion.

Kahneman does not suggest that people are incapable of System 2 thought and always follow their intuition. System 2 engages when circumstances require. Rather, many of our actual choices in life, including some important and consequential ones, are System 1 choices, and therefore are subject to substantial deviations from the predictions of the standard economic model. System 1 leads to brilliant inspirations, but also to systematic errors.

To illustrate, consider one of KT's most compelling questions/experiments:

An individual has been described by a neighbor as follows: "Steve is very shy and withdrawn, invariably helpful but with very little interest in people or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail." Is Steve more likely to be a librarian or a farmer?

Most people reply quickly that Steve is more likely to be a librarian than a farmer. This is surely because Steve resembles a librarian more than a farmer, and associative memory quickly creates a

picture of Steve in our minds that is very librarian-like. What we do not think of in answering the question is that there are 5 times as many farmers as librarians in the US, and that the ratio of male farmers to male librarians is even higher (this certainly did not occur to me when I first read the question many years ago, and does not even occur to me now as I reread it, unless I force myself to remember). The base rates simply do not come to mind and thus prevent an accurate computation and answer, namely that Steve is more likely to be a farmer. System 2 does not engage.

In another example (due to Shane Frederick), one group of respondents is asked (individually) to estimate the total number of murders in Detroit in a year. Another group is asked to estimate the total number of murders in Michigan in a year. Typically, the first group on average estimates a higher number of murders than the second. Again, System 1 thinking is in evidence. Detroit evokes a violent city, associated with many murders. Michigan evokes idyllic apple-growing farmland. Without System 2 engagement, the fact that Detroit is *in* Michigan does not come to mind for the second group of respondents, leading – across subjects – to a dramatic violation of basic logic.

Kahneman's other examples of System 1 thinking include adding $2 + 2$, completing the words "bread and ...," and driving a car on an empty road. Calling all these examples System 1 thinking captures the rapid, intuitive, automatic response, which usually gets the right answer, but sometimes – as with Steve and murders in Michigan – does not. Yet unfortunately things are not as clear as they look, once we apply our own System 2 thinking to System 1.

First, as Kahneman readily recognizes, the domains of System 1 and System 2 differ across people. For most (all?) readers of this review, computing 20×20 is a System 1 effortless task, largely because economists have both been selected to be good at it and have had lots of practice. But for many people who are not experts, this operation is effortful, or even impossible, and is surely the domain of System

2. In contrast, screwing in a light bulb is very system 2 for me – conscious, effortful and slow – but not so for most people, I gather. As people gain knowledge or expertise, the domains of the two systems change. In fact, the classification of decisions into products of System 1 and System 2 thinking seems to be even harder. Go back to murders in Detroit and in Michigan. The question surely evoked images of bombed-out Detroit and pastoral Michigan, but constructing the estimate also requires a substantial mental effort. Both systems seem to be in action.

Second, the challenge of going beyond the labels is that System 2 is not perfect, either. Many people would get 20 X 20 wrong, even if they think hard about it. The idea that conscious thought and computation are imperfect goes back at least to Herbert Simon and his concept of bounded rationality. Bounded rationality is clearly important for many problems (and in fact has been fruitfully explored by economists), but it is very different from Kahneman's System 1. Kahneman's brilliant insight – illustrated again and again throughout the book – is that people do not just get hard problems wrong, as bounded rationality would predict; they get utterly trivial problems wrong because they don't think of them in the right way. This is a very different notion than bounded rationality. Still, the challenge remains that when we see a decision error, it is not obvious whether to attribute it to System 1 thinking, System 2 failure, or a combination.

Third, the classification of thought into System 1 and System 2 raises tricky questions of the relationship between the two. Because System 1 includes unconscious attention, perception, and associative memory, much of the informational input that System 2 receives comes via System 1. Whether and how System 1 sends "up" the message if at all is a bit unclear. In other words, what prompts the engagement of System 2? What would actually trigger thinking about relative numbers of male librarians and farmers in the United States, or even whether Michigan includes Detroit? I am not sure that anything but a hint would normally do it. Perhaps System 2 is almost always at rest.

Furthermore, one function of System 2 appears to be to “check the answers” of System 1, but if information “sent up” is incomplete and distorted, how would System 2 know? To strain the legal analogy a bit further, appellate courts in the US must accept fact finding of trial courts as given, so many errors – as well as deliberate distortions – creep in precisely at the fact-finding trial stage, rather than in the appealable application of law to the facts. Kahneman writes that “the division of labor between System 1 and System 2 is highly efficient: it minimizes effort and optimizes performance (p. 25)”. I am not sure why he says so. If System 1 guides our insurance and investment choices described in the introduction, then System 2 seems rather disengaged even when the costs of disengagement are high.

To put these comments differently, each of System 1 and System 2 appears to be a collection of distinct mental processes. System 1 includes unconscious attention, perception, emotion, memory, automatic causal narratives, etc. I am worried that, once the biology of thought is worked out, what actually happens in our heads is unlikely to neatly map into fast and slow thinking. The classification is an incredibly insightful and helpful metaphor, but it is not a biological construct or an economic model. Turning metaphors into models remains a critical challenge.

II. Heuristics and Biases.

One of the two main bodies of KT’s work has come to be known as “Heuristics and Biases.” This research deals, broadly, with intuitive statistical prediction. The research finds that individuals use heuristics or rules of thumb to solve statistical problems, which often leads to biased estimates and predictions. KT have identified a range of now famous heuristics, which fall into two broad categories.

Some heuristics involve respondents answering questions for which they do not have much idea about the correct answer, and must retrieve a guess from their memory. The problem given to them is

not self-contained. As a consequence, respondents grasp at straws, and allow their answers to be influenced by objectively irrelevant frames. One example of this is the anchoring heuristic. A wheel of fortune, marked from 0 to 100, is rigged by experimenters to stop only at either 10 or 65. After a spin, students write down the number at which it stopped, and are then asked two questions: Is the percentage of African nations among UN members larger or smaller than the number you just wrote? What is your best guess of the percentage of African nations in the UN? For students who saw the wheel of fortune stop at 10, the average guess was 25%. For those who saw it stop at 65, the average guess was 45%. Similar experiments have been run with lengths of rivers, heights of mountains, and so on. The first question anchors the answer to the second. Kahneman interprets anchoring as an extreme example of System 1 thinking: planting a number in one's head renders it relevant to fast decisions.

The second category of heuristics is much closer to economics, and in fact has received a good deal of attention from economists. These heuristics describe statistical problems in which respondents receive all the information they need, but nonetheless do not use it correctly. Not all available information seems to come to the top of the mind, leading to errors. Examples of neglected decision-relevant information include base rates (even when they are explicitly stated), low probability but non-salient events, and chance. The finding that the causal and associative System 1 does not come up with chance as an explanation seems particularly important. Kahneman recalls a magnificent story of Israeli Air Force officers explaining to him that being tough with pilots worked miracles, because when pilots had a poor landing and got yelled at, their next landing was better, but when they had a great landing and got praised, their next landing was worse. To these officers, the role of chance and consequent mean reversion in landing quality did not come to mind as an explanation.

The best known problems along these lines describe the representativeness heuristic, of which the most tantalizing is Linda, here slightly abbreviated:

Linda is thirty-one years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

After seeing the description, the respondents are asked to rank in order of likelihood various scenarios: Linda is 1) an elementary school teacher, 2) active in the feminist movement, 3) a bank teller, 4) an insurance salesperson, and 5) a bank teller also active in the feminist movement. The remarkable finding is that (now generations of) respondents deem scenario 5) more likely than scenario 3), even though 5) is a special case of 3). The finding thus violates the most basic laws of probability theory. Not only do many students get the Linda problem wrong, but some object, sometimes passionately, after the correct answer is explained.

What's going on here? The description of Linda brings to mind, presumably from associative memory, a picture that does not look like a bank teller. Asked to judge the likelihood of scenarios, respondents automatically match that picture to each of these scenarios, and judge 5) to be more similar to Linda than 3). System 1 rather easily tells a story for scenario 5), in which Linda is true to her beliefs by being active in the feminist movement, yet must work as a bank teller to pay rent. Telling such a story for 3) that puts all the facts together is more strenuous because a stereotypical bank teller is not a college radical. The greater similarity of Linda to the feminist bank teller leads respondents to see that a more likely scenario than merely a bank teller.

Many studies have unsuccessfully tried to debunk Linda. It is certainly true that if you break Linda down for respondents (there are 100 Lindas, some are bank tellers, some are feminist bank tellers, which ones are there more of?) – if you engage their System 2 – you can get the right answer. But this, of course, misses the point, namely that, left to our own devices, we do not engage in such breakdowns. System 2 is asleep. In Linda, as in Steve the librarian and many other experiments, the full statistical

problem simply does not come to mind, and fast-thinking respondents – even when they do strain a bit – arrive at an incorrect answer.

There have been several attempts by economists to model such intuitive statistics (e.g. Mullainathan 2000, 2002, Rabin 2002, Rabin and Vayanos 2010, Schwartzstein 2010). In one effort that sought to stay close to Kahneman's System 1 reasoning, Gennaioli and I (2010) argued that individuals solve decision problems by representing them – automatically but incompletely -- in ways that focus on features that are statistically more associated with the object being assessed. In the Linda problem, the feminist bank teller is described comprehensively and hence represented as a feminist bank teller. A bank teller, in contrast, is not described comprehensively, and bank teller evokes the stereotype of a non-feminist because not being a feminist is relatively more associated with being a bank teller than being a feminist. The decision maker thus compares the likelihoods *not* of bank teller vs feminist bank teller, but rather of the stereotypical (representative) non-feminist bank teller vs feminist bank teller, and concludes that Linda the college radical is more likely to be the latter. This approach turns out to account for a substantial number of heuristics discussed in Kahneman's book. The key idea, though, is very much in the spirit of System 1 thinking, but made tractable using economic modeling, namely that to make judgments we *represent* the problem automatically via the functioning of attention, perception, and memory, and our decisions are subsequently distorted by such representation.

The representativeness heuristic had a substantial impact on behavioral finance, largely because it provides a natural account of *extrapolation* – the expectation by investors that trends will continue. The direct evidence on investor expectations of stock returns points to a strong extrapolative component (e.g., Vissing-Jorgensen 2003). Extrapolation has been used to understand price bubbles (Kindleberger 1978), but also the well-documented overvaluation and subsequent reversal of high performing growth stocks (De Bondt and Thaler 1985, Lakonishok et al 1994). Indeed, data for a variety of securities across

markets show that price trends continue over a period of several months (the so-called momentum), but that extreme performance reverts over longer periods (Cutler, Poterba, and Summers 1991). Even more dramatically, investors put money into well performing mutual funds, into stock funds and stock market-linked insurance products after the stock market has done well (Frazzini and Lamont 2008, Yagan 2012). Such phenomena have been described colorfully as investors “jumping on the bandwagon” believing that “the trend is your friend,” and failing to realize that “trees do not grow to the sky,” that “what goes up must come down,” etc.

Heuristics provide a natural way of thinking about these phenomena, and can be incorporated into formal models of financial markets (see, e.g., Barberis et al. 1998). Specifically, when investors pour money into hot, well-performing assets, they may feel that these assets are similar to, or resemble, other assets that have kept going up. Many high tech stocks look like the next Google, or at least System 1 concludes that they do. Extrapolation is thus naturally related to representativeness, and supports the relevance of Kahneman’s work not just in the lab, but also in the field.

III. Prospect Theory.

Prospect Theory has been KT’s most influential contribution, and deservedly so. In a single paper, the authors proposed an alternative to standard theory of choice under risk that was at the same time quite radical and tractable, used the theory to account for a large number of outstanding experimental puzzles, and designed and implemented a collection of new experiments used to elucidate and test the theory. In retrospect, it is difficult to believe just how much that paper had accomplished, how new it was, and how profound its impact has been on behavioral economics.

Prospect Theory rests on four fundamental assumptions. First, risky choices are evaluated in terms of their gains and losses relative to a reference point, which is usually the status quo wealth. Second, individuals are loss averse, meaning extremely risk averse with respect to small bets around the reference point. Third, individuals are risk averse in the domain of gains, and risk loving in the domain of losses. And finally, in assessing lotteries, individuals convert objective probabilities into decision weights that over-weight low probability events and under-weight high probability ones.

The first assumption is probably the most radical one. It holds that rather than integrating all risky choices into final wealth states, as standard theory requires, individuals frame and evaluate risky bets narrowly in terms of their gains and losses relative to a reference point. In their 1979 paper, KT did not dwell on what the reference point is, but for the sake of simplicity took it to be the current wealth. In a 1981 *Science* paper, however, they went much further in presenting a very psychological view of the reference point: “The reference outcome is usually a state to which one has adapted; it is sometimes set by social norms and expectations; it sometimes corresponds to a level of aspiration, which may or may not be realistic (p. 456).” The reference point is thus left as a rather unspecified part of KT’s theory, their measure of “context” in which decisions are made. Koszegi and Rabin (2006) suggest that reference points should be rational expectations of future consumption, a proposal that brings in calculated thought. Pope and Schweizer (2011) find that goals serve as reference points in professional golf. Hart and Moore (2008) believe that contracts serve as reference points for future negotiations. A full elaboration of where reference points come from is still “under construction.”

The second assumption of Prospect Theory is loss aversion. It is inspired by a basic and intuitively appealing experiment in which people refuse to take bets that give them a 60% probability of winning a dollar and a 40% probability of losing a dollar, even though such a refusal implies an implausibly high level of risk aversion (Rabin 2000). Kahneman justifies this assumption by noting that,

biologically, losses might be processed in part in the amygdala in the same way as threats. KT modeled this assumption as a kink in the value function around the reference point. In fact, in its simplest version, Prospect Theory (without assumptions 3 and 4 described below) is occasionally presented graphically with a piecewise linear value function, with the slope of 1 above the origin and 2 below the origin (reference point), and a kink at the origin that captures loss aversion. Kahneman sees loss aversion as the most important contribution of Prospect Theory to behavioral economics, perhaps because it has been used to account for the endowment effect (the finding, both in the lab and in the field, that individuals have a much higher reservation price for an object they own than their willingness to pay for it when they do not own it).

The third assumption is that behavior is risk averse toward gains (as in standard theory) and risk seeking toward losses. It is motivated by experiments in which individuals choose a gamble with a 50% chance of losing \$1000 over a certainty of losing \$500. This assumption receives some though not total support (Thaler and Johnson 1990), and has not been central to Prospect Theory's development.

The fourth assumption of Prospect Theory is quite important. That is the assumption of an inverted S-shaped function converting objective probabilities into decision weights, that blows up low probabilities and shrinks high ones (but not certainty). The evidence used to justify this assumption is the excessive weights people attach to highly unlikely but extreme events: they pay too much for lottery tickets, overpay for flight insurance at the airport, or fret about accidents at nuclear power plants. KT use probability weighting heavily in their paper, adding several functional form assumptions (subcertainty, subadditivity) to explain various forms of the Allais paradox. In the book, Kahneman does not talk about these extra assumptions, but without them Prospect Theory explains less.

To me, the stable probability weighting function is problematic. Take low probability events. Some of the times, as in the cases of plane crashes or jackpot winnings, people put excessive weight on them, a phenomenon incorporated into Prospect Theory that Kahneman connects to the availability heuristic. Other times, as when investors buy AAA-rated mortgage backed securities, they neglect low probability events, a phenomenon sometimes described as black swans (Taleb 2007). Whether we are in the probability weighting function or the black swan world depends on the context: whether or not people recall and are focused on the low probability outcome.

More broadly, how people think about the problem influences probability weights and decisions. In one of KT's most famous examples, results from two potential treatments of a rare disease are described, alternatively, in terms of lives saved and lives lost. The actual outcomes – gains and losses of life -- are identical in the two descriptions. Yet respondents choose the “safer” treatment when description is in terms of lives saved, and the “riskier” treatment when description is in terms of lives lost. The framing or representation of the problem thus changes probability weights even when objective outcomes are identical. In another study, Rottenstreich and Hsee (2001) show that decision weights depend on how “affect-rich” the outcomes are, and not just on their probabilities. Bordalo, Gennaioli, and Shleifer (2012a) present a model in which attention is drawn to salient, or unusual, payoffs. In their model, unlike in Prospect Theory, individuals overweigh only low probability events that are associated with extreme, or salient, payoffs. The model explains all the same findings as Prospect Theory, but also several additional ones, including preference reversals (people sometimes prefer A to B, but are willing to pay more for B than for A when considering the two in isolation). Kahneman of course recognizes the centrality of context in shaping mental representation of problems when he talks about the WYSIATI principle (what you see is all there is).

Prospect Theory is an enormously useful model of choice because it accounts for so much evidence and because it is so simple. Yet it achieves its simplicity by setting to one side both in its treatment of reference points and its model of probability weights precisely the System 1 mechanisms that shape how a problem is represented in our minds. For a more complete framework, we need better models of System 1.

Prospect Theory has been widely used in economics, and many of the applications are described in Della Vigna (2010) and Barberis (2012). Finance is no exception. Benartzi and Thaler (1995) have argued, for example, that it can explain the well-known equity premium puzzle, the empirical observation that stocks on average earn substantially higher returns than bonds. Benartzi and Thaler observed that while stocks do extremely well in the long run, they can fall a lot in the short run. When investors have relatively short horizons, and also, in line with Prospect Theory, are loss averse, this risk of short term losses in stocks looms large, makes stocks unattractive, and therefore cheap, thus explaining the equity premium. More recently, Barberis and Huang (2008) argue that the probability weighting function of Prospect Theory has the further implication that investors are highly attracted to positive skewness in returns, since they place excessive weights on unlikely events. The evidence on overpricing of Initial Public Offerings and out of the money options is consistent with this prediction.

IV. What's Ahead?

In conclusion, let me briefly mention three directions in which I believe the ship launched by Kahneman and Tversky is headed, at least in economics. First, although I did not talk much about this in the review, Kahneman's book on several occasions discusses the implications of his work for policy. At the broadest level, how should economic policy deal with System 1 thinking? Should it respect

individual preferences as distinct from those dictated by the standard model or even by the laws of statistics? Should it try to debias people to get them to make better decisions?

I have avoided these questions in part because they are extremely tricky, at both philosophical and practical levels (Bernheim and Rangel 2009). But one theme that emerges from Kahneman's book strikes me as important and utterly convincing. Faced with bad choices by consumers, such as smoking or undersaving, economists as System 2 thinkers tend to focus on education as a remedy. Show people statistics on deaths from lung cancer, or graphs of consumption drops after retirement, or data on returns on stocks versus bonds, and they will do better. As we have come to realize, such education usually fails. Kahneman's book explains why: System 2 might not really engage until System 1 processes the message. If the message is ignored by System 1, it might never get anywhere. The implication, clearly understood by political consultants and Madison Avenue advertisers, is that effective education and persuasion must connect with System 1. Calling the estate tax "the death tax" may work better to galvanize its opponents than statistics on hard-working American farmers who may have to pay. Thaler and Sunstein's (2008) *Nudge* advocates policies that simplify decisions for people relying on System 1 in situations such as saving for retirement where even an educated System 2 might struggle.

Beyond the changing thinking on economic policy, Kahneman's work will continue to exert a growing influence on our discipline. A critical reason for this is the rapidly improving quality of economic data, from the field, from experiments, and from field experiments. Confronted with the realities of directly observed human behavior – financial choices made by investors, technology selection by farmers, insurance choices by the elderly – economists have come to psychology for explanations, especially to the work described in Kahneman's book. Rapidly expanding data on individual choices is the behavioral economist's best friend.

But it seems to me that some of the most important advances in the near future both need to come, and will come, in theory. Economics, perhaps like any other discipline, advances through changes in standard models: witness the enormous influence of Prospect Theory itself. In contrast, we do not have a standard model of heuristics and biases, and as I argued, Prospect Theory is still a work in progress. Fortunately, the broad ideas discussed in Kahneman's book, and in particular his emphasis on the centrality of System 1 thinking, provide some critical clues about the features of the models to come.

In particular, the main lesson I learned from the book is that we represent problems in our minds, quickly and automatically, before we solve them. Such representation is governed by System 1 thinking, including involuntary attention drawn to particular features of the environment, focus on these features, and recall from memory of data associated with these perceptions. Perhaps the fundamental feature of System 1 is that what our attention is drawn to, what we focus on, and what we recall, is not always what is most necessary or needed for optimal decision making. Some critical information is ignored; other -- less relevant -- information receives undue attention because it stands out. In this respect, the difference from the models of bounded rationality, in which information is optimally perceived, stored, and retrieved is critical. System 1 is automatic and reactive, not optimizing.

As a consequence, when we make a judgment or choice, we do that on the basis of incomplete and selected data assembled via a System 1-like mechanism. Even if the decisions are optimal at this point given what we have in mind, they might not be optimal given the information potentially available to us both from the outside world and from memory. By governing what we are thinking about, System 1 shapes what we conclude, even when we are thinking hard.

Kahneman's book, and his lifetime work with Tversky, had and will continue to have enormous impact on psychology, applied economics, and policy making. Theoretical work on Kahneman and

Tversky's ideas has generally modeled particular heuristics and choices under risk separately, without seeking common elements. A potentially large benefit of Kahneman's book is to suggest a broader theme, namely that highly selective perception and memory shape what comes to mind, before we make decisions and choices. Nearly all the phenomena the book talks about share this common thread. In this way, Kahneman points toward critical ingredients of a more general theory of intuitive thinking, still an elusive, but perhaps achievable, goal.

References

- Barberis, Nicholas. 2012. "Thirty Years of Prospect Theory in Economics." *Journal of Economic Perspectives*, forthcoming.
- Barberis, Nicolas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Pricing." *American Economic Review* 98(5): 2066-100.
- Barberis, Nicholas, Andrei Shleifer, and Robert W. Vishny. 1998. "A Model of Investor Sentiment." *Journal of Financial Economics* 49(3): 307-43.
- Benartzi, Shlomo, and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics* 110(1): 73-92.
- Bernheim, B. Douglas, and Antonio Rangel. 2009. "Beyond Revealed Preference: Choice Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124(1): 51-104.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2010. "What's Advertising Content Worth? Evidence from Consumer Credit Marketing Field Experiment." *Quarterly Journal of Economics* 125(1): 263-306.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012a. "Salience Theory of Choice Under Risk." *Quarterly Journal of Economics*, forthcoming.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012b. "Salience and the Experimental Tests of the Endowment Effect." *American Economic Review Papers and Proceedings* 102(3):47-52.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012c. "Salience and Consumer Choice." National Bureau of Economic Research Working Paper 17947.
- Chalmers, John, and Jonathan Reuter. 2012. "What is the Impact of Financial Advisors on Retirement Portfolio Choices and Outcomes?" National Bureau of Economic Research Working Paper 18158.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers. 1991. "Speculative Dynamics." *Review of Economic Studies* 58(3): 529-46.
- Della Vigna, Stefano. 2010. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47(2): 315-72.

- Della Vigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics* 2010(2): 643-69.
- De Bondt, Werner F. M., and Richard H. Thaler. 1985. "Does the Stock Market Overreact?" *Journal of Finance* 40(3): 793-805.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy* 98(4): 703-38.
- Frazzini, Andrea, and Owen A. Lamont. 2008. "Dumb Money: Mutual Fund Flows and the Cross-section of Returns." *Journal of Financial Economics* 88(2): 299-322.
- Friedman, Milton. 1953. "The Case for Flexible Exchange Rates." In *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Gennaioli, Nicola, and Andrei Shleifer. 2010. "What Comes to Mind." *Quarterly Journal of Economics* 125(4):1399-1433.
- Gennaioli, Nicola, Andrei Shleifer, and Robert W. Vishny. 2012. "Money Doctors." National Bureau of Economic Research Working Paper 18174.
- Grossman, Sanford J., and Merton H. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43: 617-33.
- Hart, Oliver, and John Moore. 2008. "Contracts as Reference Points." *Quarterly Journal of Economics* 123 (1): 1-48.
- Kindleberger, Charles P. 1978. *Manias, Panics, and Crashes: A History of Financial Crises*. New York: Basic Books.
- Koszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121(4): 1133-65.
- Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny. 1994. "Contrarian Investment, Extrapolation, and Risk." *Journal of Finance* 49(5): 1541-78.
- List, John A. 2003. "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics* 118(1): 41-71.
- Mullainathan, Sendhil. 2000. "Thinking through Categories." Unpublished.

- Mullainathan, Sendhil. 2002. "A Memory-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 117(3): 735-74.
- Mullainathan, Sendhil, Andrei Shleifer, and Joshua Schwartzstein. 2008. "Coarse Thinking and Persuasion." *Quarterly Journal of Economics* 123(2): 577-619.
- Pope, Devin G., and Maurice E. Schweitzer. 2011. "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes." *American Economic Review* 101: 129-57.
- Rabin, Matthew. 2000. "Risk Aversion and Expected Utility: A Calibration Theorem." *Econometrica* 68: 1281-92.
- Rabin, Matthew. 2002. "Inference by Believers in the Law of Small Numbers." *Quarterly Journal of Economics* 117(3): 775-816.
- Rabin, Matthew, and Dimitri Vayanos. 2010. "The Gambler's and Hot-Hand Fallacies: Theory and Applications." *Review of Economic Studies* 77(2): 730-78.
- Rottenstreich, Yuval, and Christopher K. Hsee. 2001. "Money, Kisses, and Electric Shocks: On the Affective Psychology of Risk." *Psychological Science* 12(3): 185-90.
- Schwartzstein, Joshua. 2010. "Selective Attention and Learning." <http://www.dartmouth.edu/~jschwartzstein/papers/sa.pdf>.
- Shleifer, Andrei, and Robert W. Vishny. 1997. "The Limits of Arbitrage." *Journal of Finance* 52(1): 35-55.
- Stanovich, Keith, and Richard West. 2000. "Individual Differences in Reasoning: Implications for the Rationality Debate." *Behavioral and Brain Sciences* 23(5): 645-65.
- Sydnor, Justin. 2010. "(Over)insuring Modest Risks." *American Economic Journal: Applied Economics* 2(4): 177-99.
- Taleb, Nassim Nicholas. 2007. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.
- Thaler, Richard H., and Eric J. Johnson. 1990. "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice." *Management Science* 36(6): 643-60.
- Thaler, Richard H., and Cass R. Sunstein. 2008. *Nudge*. New Haven: Yale University Press.

Tversky, Amos, and Daniel Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice." *Science* 211: 453-458.

Vissing-Jorgensen, Annette. 2003. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions." *NBER Macroeconomics Annual* 18(1): 139-94.

Yagan, Danny. 2012. "Why Do Individual Investors Chase Stock Market Returns? Evidence of Belief in Long-Run Market Momentum." http://www.people.fas.harvard.edu/~yagan/papers/Chasing_Returns.pdf.