Taming Uncertainty

Ralph Hertwig, Timothy J. Pleskac, Thorsten Pachur,
and The Center for Adaptive Rationality

The MIT Press
Cambridge, Massachusetts
London, England

10 Experiences and Descriptions of Financial Uncertainty: Are They Equivalent?

Tomás Lejarraga, Jan K. Woike, and Ralph Hertwig

> No one can possibly have lived through the Great Depression without being scarred by it.... No "Depression baby" can ever be a yuppie. No amount of experience since the Depression can convince someone who has lived through it that the world is safe economically. One constantly waits for banks to close, for factories to shut down, for the pink slip of discharge.
>
> —Isaac Asimov, *I, Asimov: A Memoir*

10.1 Once Bitten, Twice Shy

In his memoir, science fiction writer Isaac Asimov described the hardships of growing up in Brooklyn during the Great Depression of the 1930s: the uncertainty, the fear, the pessimism. The effects of the crisis were devastating. Unemployment almost tripled in just three years, from 8.7% in 1930 to 25% in 1933. Half of the nation's banks failed due to customers defaulting on their obligations and withdrawing their savings in frantic bank runs. Many households lost everything. Countless families were forced to migrate to regions with better prospects, such as California, in an exodus unprecedented in the history of the United States. The Great Depression left an indelible mark on a generation of Americans that would become visible only decades later, as the economy recovered. Many, like Asimov, believe that growing up in economic hardship produces "Depression babies," who are pessimistic about the vagaries of the economy and unwilling to take economic risks. And their intuitions have recently gained scientific support. Using data spanning almost five decades from the Survey of Consumer Finances, Malmendier and Nagel (2011) found that people who have experienced economic shocks such as the Great Depression are indeed more averse to uncertainty. They are less likely to invest in the stock market and, if they do, they invest a lower proportion of their assets in stocks, favoring safer options. The more recent the experience of a shock, the more averse they are to uncertainty.

Financial concerns profoundly influence many important life decisions: the choice of career, place of residence, and even life partner are, to some extent, influenced by how people cope with financial uncertainty. How does experiencing a financial shock influence the perception of financial uncertainty? Is it possible to teach people about living with financial uncertainty before they are forced to learn from experience? In 2008, the world suffered the worst economic shock since the Great Depression. Again, major investment and savings banks went bankrupt and people around the globe lost their life savings. In the United States alone, households are estimated to have lost 11 trillion dollars, and 5.5 million people lost their jobs. Will this recession have a similar impact on millennials' risk taking as the Great Depression had on that of Depression babies? Although it is too early to draw definitive conclusions, financial analysts are already seeing signs that young adults who began investing in the last decade have much in common with Depression babies: "Millennials are really risk averse" (Russolillo, 2014).

Massive economic shocks are arguably exceptions to the normal functioning of the markets. Ignoring these rare events allows financial theorists to work with parsimonious and—some would argue—dangerously simplistic models of the markets (e.g., Taleb, 2007) and their actors. Traditional financial theory assumes that investors have stable risk attitudes and that they update their beliefs rationally (see also chapter 17). From this perspective, as long as wealth and income are kept constant, personally experiencing financial outcomes should be no different from learning about them from any other source of information, such as a newspaper. These assumptions, reasonable in stable markets, cannot accommodate the changes in risk attitudes that result from experiencing a financial shock. But experts are starting to believe that financial crises have become more frequent (Elliott & Milner, 2001), and acknowledging their impact is currently one of the key concerns for financial economists.

10.2 Does a Description of a Financial Shock "Burn" as Severely as the Experience of It?

Mark Twain's (1894/2004) fictional character Pudd'nhead Wilson recounts the story of a cat sitting on a hot stove, getting burned, and never sitting on a stove again. The story teaches an important lesson about coping with uncertainty: by avoiding all stoves, the cat never got burned again. But it forever forfeited the comfort of a warm stove on a cold winter morning. Much like Depression babies, who avoid the stock market after getting "burned," humans and other animals exhibit the "hot stove effect" (Denrell & March, 2001): after an unexpected negative experience, they avoid the source of the uncertainty altogether.

Humans are practically alone in their ability to learn by means other than experience. Barely any aspect of modern life—from science, commerce, and the World Wide Web to fiction and poetry—would be conceivable without the human ability to produce and interpret symbolic descriptions (Schmandt-Besserat, 1996). Individuals are able to communicate and transmit their experiences to others through descriptions such as statistics, graphs, and texts. For example, later generations can learn about the Great Depression through tables of statistics, graphical representations of dramatic slumps in stock prices, or literary accounts that recreate the experiences of those who suffered. In John Steinbeck's (1939/2014) novel *The Grapes of Wrath*, a family of farmworkers migrates from Oklahoma to California—only to find that the conditions there are even worse than those they left behind. *The Grapes of Wrath* describes what Depression babies lived through in narrative form. But do descriptions of a crisis teach the same lessons about financial uncertainty as does actually experiencing a crisis?

To address this question, we (Lejarraga, Woike, & Hertwig, 2016) created a financial crisis in the laboratory and observed how people reacted to it. Clearly, unlike the Great Depression, our "experimental crisis" did not impact every dimension of investors' lives—but it potentially slashed the value of their portfolio by more than half. Our investors were given an experimental portfolio of €100 and asked to split this amount between a risky and a safe option across a number of monthly periods. The safe option was a cash deposit account offering a 0.25% rate of return each month (i.e., a 3% annual rate of return). The risky option was the Spanish stock market index, IBEX 35, offering the actual monthly rates of return obtained from

July 1999 to September 2013. Between 1999 and 2013, the IBEX 35 experienced two shocks—the first from 1999 to 2002 (resulting in a 57% drop in stock price); the second, from late 2007 to 2009 (resulting in a 52% drop in stock price). At the end of the experiment, investors earned an amount corresponding to the value of their respective portfolios.

For each monthly period, investors stipulated the proportion of their portfolio to be invested in stocks (i.e., the index fund); the rest was invested in a safe cash deposit account. The return on their investment was added to (or subtracted from) their current portfolio balance. Feedback on returns was given in a table (amounts earned from stocks and from the cash deposit account) and in three graphs. One graph showed the current and historical prices of the index fund during the experimental investment period. A second graph showed past and present rates of return on the index fund, the cash deposit account, and the portfolio. A third graph showed the development of the portfolio. The three graphs were updated after each investment.

The central feature of our investigation was that some investors "lived through a crisis," namely, they started making virtual investments before the 2000 dot-com crisis, whereas others learned about a crisis from descriptions. Our descriptions were not narratives like *The Grapes of Wrath*. Rather, investors learned about a past crisis by inspecting a graph—that is, they were informed in the same way as real-world investors are typically informed by banks, financial consultants, or online services. Another set of investors entered the investment game after the first crisis and without knowledge of it; here again, some investors learned from description and others from experience. Figure 10.1a describes the four conditions, overlaid on the evolution of the IBEX 35: solid lines indicate periods of investment and dotted lines indicate periods in which investors learned from descriptive sources (graphs and table). Box 10.1 gives the details of the experimental design. To experience being an investor in the market, please access interactive elements 10.1 and 10.2 (at https://taming-uncertainty.mpib-berlin.mpg.de/).

10.3 Do Descriptions and Experience Steer Investors toward the Same Investment Behavior?

We evaluated investors' behavior during a long evaluation window (72 investments, see figure 10.1). To this end, we defined the measure of risk taking R as the proportion of a person's assets invested in stocks (vs. in

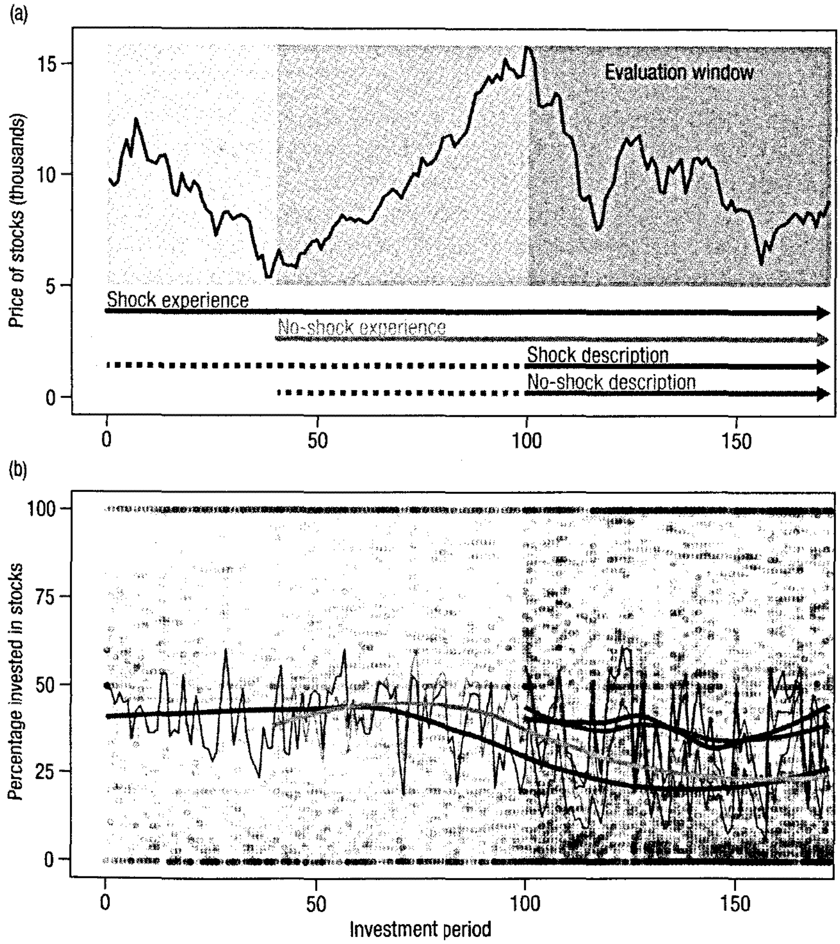

Figure 10.1

(a) Experimental conditions and price of stocks in thousands of euros (i.e., index fund) across 172 monthly periods. Solid arrow segments indicate periods of investment. Dotted arrow segments indicate periods of learning from descriptive sources. The four conditions were compared over the evaluation window from period 100 to period 172. (b) Percentage invested in stocks by condition. Dots indicate individuals' allocations. The thin lines show the mean percentages; the thicker lines show the data smoothed by local polynomial regression fitting. (Based on Lejarraga et al., 2016)

Box 10.1

Experimental market and investment decisions.

Two hundred investors (40% male, mean age 25 years, $SD=3.5$ years) were randomly assigned to one of four conditions that varied with respect to the length of historical data available and the mode of learning (see figure 10.1a for a summary of the four conditions). The 50 investors in the *shock experience* condition made investments in all 172 periods of the experiment (i.e., the 172 months between July 1999 and September 2013). These investors experienced a stock market that initially fell until about period 40. The 50 investors in the *no-shock experience* condition entered the market in period 40 and made decisions in 133 periods. These investors experienced a market that initially rose for around 60 periods; they were unaware of the previous downward trend.

The remaining 100 investors entered the market in period 100 and made 73 investment decisions. Of these investors, the 50 in the *shock description* condition were shown a graph plotting the price of the index fund since period 1. These investors thus learned from the graph what investors in the shock experience condition learned from experience (i.e., the development of the index fund across periods 1–99). The 50 investors in the *no-shock description* condition were shown a graph plotting the price of the index fund since period 40. Like their counterparts in the no-shock experience condition, they remained unaware of the market's initial downward trend, but they learned about the later upward trend—in this case, from the graph. Investors were not told in advance how many decisions they would make. They were paid according to the value of their portfolio at the end of the investment task.

Interactive element 10.1 allows you to experiment with a simplified version of the interface used by investors. Investors made their investment decisions by manipulating sliders in the upper right of the screen. Graphs and a table in the bottom panel provided information on the development of the stock index, the period-by-period return on each amount invested in stocks and in cash deposit, as well as the return on the whole portfolio. Investors first read the instructions on the computer screen and completed 10 periods as practice trials. They were informed that the return data in the practice trials were randomly generated. They were also given a printed booklet of instructions (including definitions of all concepts involved in the investment task) that they could consult at any time. The return data in the experimental trials reflected the actual returns on the IBEX 35, but the dates were shifted 25 years into the future to prevent investors from using historical knowledge to predict stock movements.

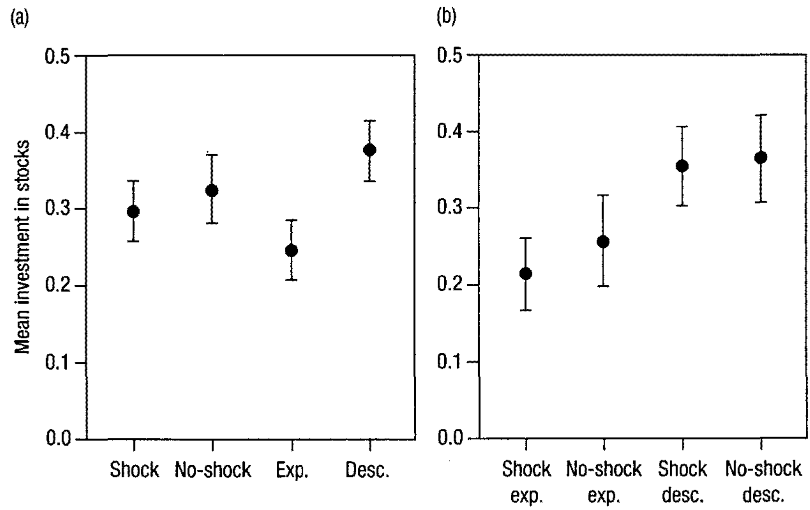

Figure 10.2
Mean investment in stocks (a) by mode of learning (experience, description) and by the experience of a shock or lack thereof; and (b) for the resulting four conditions (shock experience, no-shock experience, shock description, no-shock description). Error bars indicate the 95% confidence interval (CI). Means and confidence intervals were calculated by averaging risk taking for each individual across periods and computing the mean across individuals (i.e., means and CIs reflect independent observations).

the cash deposit account). Lower values of *R* indicate avoidance of financial uncertainty. Figure 10.1b shows the average trends for each condition; figure 10.2 shows the average *R* (collapsed across periods). Descriptions and experience were very clearly not equivalent forms of learning about a shock. Investors who experienced the shock (shock experience condition, R_{se}=22%) were much more uncertainty-avoidant than those who learned about it from a graph (shock description condition, R_{sd}=37%), taking 15% less risk. These results indicate that experiencing a market shock indeed makes people shy away from financial uncertainty in a way that exposure to a graphical description of that same shock does not. This result represents an experimental demonstration of the Depression baby effect (Malmendier & Nagel, 2011).

The asymmetry between learning from experience and learning from description goes beyond financial shocks. In fact, a description–experience gap emerged even without experience of a shock. Investors who experienced a shock-free market (no-shock experience, $R_{nse}=27\%$) were also more uncertainty-avoidant than those who learned about the same market from a graph (no-shock description, $R_{nsd}=38\%$), taking 11% less risk. Averaged across the shock and no-shock conditions, investors who learned from a graph took 13% more risk ($R_d=38\%$) than did investors who experienced the market firsthand ($R_e=25\%$). Moreover, those who learned about past market performance from a graph proved practically insensitive to the type of market history observed.

10.4 Do Recent Experiences Have More Influence on Uncertainty Avoidance than Less Recent Ones?

As figure 10.1b shows, the percentage invested in stocks in the four conditions (i.e., the thin lines) was highly volatile across periods. To determine the extent to which investors reacted to the most recent change in stock prices in each period, we calculated the individual-specific correlation between the change in price of the index fund (stock price_t/stock price_{t-1}) and the change in investment in the following period (proportion in stock_{t+1} − proportion in stock_t). This measure of reactivity is likely to underestimate the strength of the relationship, because the proportion invested in the index fund is bounded between 0 and 1, and an investor who is fully invested in stocks cannot increase the level of risk taking following an increase in stock prices. Across all investors, the mean correlation was 0.29, with little variation across conditions. Our analysis thus shows reactivity to recent changes in both description and experience, consistent with previous findings (Funk, Rapoport, & Jones, 1979; Gordon, Paradis, & Rorke, 1972; Kroll, Levy, & Rapoport, 1988; Rapoport, 1984) suggesting that people react to price changes in an attempt to capture the momentum of the market. Although the experimental design presented here compares the behavior of investors with different levels of wealth (investors who enter the market after the shock start with $100, whereas the endowments of those who experience it reflect the effects of the shock), a subsequent experiment (see study 3 in Lejarraga et al., 2016) showed that the effects described here persisted when wealth was set to be equal across conditions.

10.5 How Do People Navigate Financial Uncertainty?

Uncertainty comes in many shades. Just one of the many ways in which uncertain situations vary is in the degree to which they defy predictability. Take a train's time of arrival at a station. Although the exact time of arrival is unknown, it generally falls within a certain range. Commuters who catch the train every day will have a good idea of its punctuality and should be able to plan ahead. There is some predictability in this sort of uncertainty. On the other hand, there are singular, rare events that cannot be quantified easily. Take, for example, the deadly hurricane-strength winds of Storm Friederike that ripped through northern Europe in January 2018, causing the German rail operator Deutsche Bahn to suspend all long-distance trains nationwide for the first time since 2007. Trains were not late—they simply never left the station. This kind of uncertainty is extremely hard to plan for; it is difficult even to have a sense of the class of events that could arise. There is no predictability in this sort of uncertainty.[1]

Financial markets involve both types of uncertainty, yet financial "storms" tend to be ignored in standard theories of portfolio choice. The prime example is Markowitz's (1952) mean-variance optimization model which, like other optimization models, assumes a world of uncertainty with predictable bounds. This is an important limitation, inasmuch as it can make simple rules better suited for making investment decisions than complicated calculations of mean variance. DeMiguel, Garlappi, and Uppal (2009) studied perhaps the simplest heuristic for dealing with financial uncertainty. The $1/N$ heuristic, also known as "naive diversification," consists of dividing one's budget equally across all investment options (Benartzi & Thaler, 2001). They took a normative approach, examining how well this heuristic performed against several complex instantiations of the mean-variance model and other asset allocation models of optimization. The simple and complex allocation strategies were evaluated, among other

1. Makridakis, Hogarth, and Gaba (2009) have referred to these two types of uncertainty as "subway" and "coconut" uncertainty, respectively. The former relates to uncertain but predictable events such as the arrival time of a subway train at the station; the latter refers to singular, unpredictable events, such as a coconut falling on one's head. Similarly, in *Fooled by Randomness*, Taleb (2007) refers to rare events of extreme impact as "black swans" and emphasizes their unpredictability.

benchmarks, in terms of their performance on the Sharpe ratio, a measure of return adjusted for risk, across six datasets. For example, in the S&P 500 dataset, the simple $1/N$ heuristic had the highest Sharpe ratio across all strategies, substantially higher than the mean-variance model. In fact, the mean-variance model did not significantly outperform the $1/N$ heuristic in any of the datasets. Other research has corroborated this finding (Jacobs, Müller, & Weber, 2014; Tu & Zhou, 2011).

Whether heuristics represent good investment strategies is one question; whether people actually use them is another. Although there has been some work on the first question, little is known about the second. Our experimental design provides some answers. We first looked closely at individual investment paths; figure 10.3 gives an idea of the dramatic individual variation we found. We next defined nine simple investment strategies and examined their prevalence in our experiment (see section 10.5.1 and table 10.1

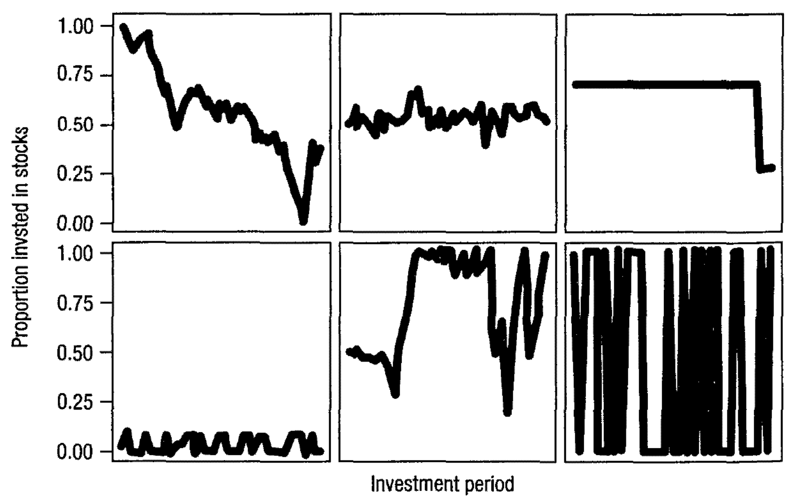

Figure 10.3
Proportion invested in stocks by six individuals in each trial of the evaluation window. These investment paths illustrate the individual variation in the data collected, and helped us to infer the strategies that people used. For example, the investor in the upper right panel was unreactive to price changes in most periods. The investor in the lower left panel took low to no risk throughout, and showed little to no reaction to the market. The investor in the lower right panel did not diversify their portfolio, as the proportion invested in stocks fluctuated from 0 to 1, possibly in response to price changes.

Table 10.1
Definition of investment strategies according to four criteria.

	Mean	*SD*	Trend	Correlation
*Strategies **unreactive** to stock price fluctuations*				
Naive diversification (1/*N*)	[0.4, 0.6]	<0.1	[0.1, −0.1]	ns
Constant risky	>0.8	<0.1	[0.1, −0.1]	ns
Constant safe	<0.2	<0.1	[0.1, −0.1]	ns
Nondiversified		>0.4		ns
*Strategies **reactive** to stock price fluctuations*				
Momentum				
Nondiversified		>0.4		s (+)
Diversified	[0.2, 0.8]	<0.4		s (+)
Risky	>0.8			s (+)
Safe	<0.2			s (+)
Contrarian				s (−)

Note: ns denotes nonsignificant correlations (with $N=73$, $r<.31$), s (+) denotes significant positive correlations, and s (−) denotes significant negative correlations.

for a definition of each strategy). Some of these strategies have been studied previously (e.g., naive diversification, momentum, and contrarian strategies); others we inferred from the investment paths observed in our data. We classified the strategies as either reactive or unreactive to changes in stock prices.

10.5.1 Strategies Unreactive to Fluctuations in Stock Price

Naive diversification (1/*N* heuristic). Investors using the naive diversification strategy divide their budget evenly among the *N* options available (Benartzi & Thaler, 2001). This strategy does not depend on the attractiveness of the options; it is not susceptible to the fluctuations of the market. In our setup, naive diversification means consistently investing 50% of one's budget in stocks.

Constant-target strategy. Investors using the constant-target strategy select a target level of risk taking in the first investment period (in our experiment, measured in terms of a specific proportion invested in the index fund, e.g., 30% or 80%) and maintain this target throughout. We distinguish two types of the constant-target strategy, depending on the level of risk taken: constant risky (>80%) and constant safe (<20%).

Nondiversified strategy. Investors using this strategy put all of their eggs in one basket—in our experiment, either the "risk" basket (100% index fund) or the "safe" basket (100% cash deposit account).

10.5.2 Strategies Reactive to Fluctuations in Stock Price

Momentum strategies. Investors using momentum strategies continuously adjust their allocations as a function of market changes. Specifically, in period t they increase allocations in the options that increased in price in period $t-1$ (Grinblatt, Titman, & Wermers, 1995). In our setup, momentum strategies imply increasing the investment in stocks after an increase in stock prices, and decreasing the investment after a decrease in price. We distinguish four types of momentum strategies: Using *momentum-nondiversified* strategies means putting the total budget in stocks following a rise in prices and moving the total budget out of stocks following a drop. *Momentum-diversified* strategies are more moderate, and follow stock fluctuations in a proportional manner. *Momentum-risky* and *momentum-safe* strategies respond to changes in stock prices, but adopt different levels of risk: momentum risky entails that more than 80% of the portfolio is invested in the index fund across all periods, and momentum safe denotes that less than 20% of the portfolio is invested in the index fund.

Contrarian strategies. Investors using contrarian strategies reduce their investment in stocks in period t after a price increase in period $t-1$, and increase their investment in stocks after a price drop (Gregory, Harris, & Michou, 2001).

Participants were classified as using a strategy only if all conditions in table 10.1 were met. About 31% of investors were left unclassified. As reported in table 10.2, a substantial proportion of investors (61%) were identified as using momentum strategies—that is, they increased the proportion invested in stocks after a price rise and reduced it after a drop. Two types of momentum strategies were predominant: investors who tracked the stock price while taking generally low risks (27%, momentum safe), and investors who followed stock fluctuations in a proportional manner (27%, momentum diversified). A minority of investors seemed to be unreactive to changes in the stock price, with 2% using naive diversification (roughly a 50/50 split) and 4% using constant-safe strategies. Our classification suggests that the investment strategies used appear not to depend on the mode of learning. Importantly, however, the mode of learning can shift

Table 10.2
Number of investors classified to each investment strategy.

	N	Percentage
Strategies ***unreactive*** *to stock price fluctuations*	**16**	**8**
Naive diversification (1/*N*)	4	2
Constant risky	1	1
Constant safe	8	4
Nondiversified	3	2
Strategies ***reactive*** *to stock price fluctuations*	**122**	**61**
Momentum		
Nondiversified	13	7
Diversified	54	27
Risky	2	1
Safe	53	27
Contrarian	0	0
Unclassified	**62**	**31**
Total	200	100

users of a strategy toward more or less risk seeking. For instance, a large proportion of investors relied on momentum strategies, whether they learned from experience or from description. However, more investors employed a momentum-safe strategy (with a lower level of risk) in the experience condition than in the description condition, again indicating that investors who learn from experience prefer less exposure to risk. The levels of risk taking following a shock also differed: more investors adopted safe strategies (constant-safe strategy and momentum strategy) in the shock than in the no-shock condition.

10.6 Can Experience Be Harnessed to Help People Navigate Financial Uncertainty?

Our results show that there is no substitute for experience. But if descriptions are not an adequate tool for warning investors about the possibility of future crises, one possibility is to "create" experiences. Kaufmann, Weber, and Haisley (2013) allowed experimental investors to "try out" as many allocations as they wanted before deciding on a final allocation. Investors were provided with immediate feedback on the risk–return profile of each simulation (try the simulation tool yourself by accessing interactive element 10.2). Kaufmann et al. found that people felt more knowledgeable and better informed when they were allowed to experience the allocations they were considering. As a result, they took more financial risk and their investments were closer to what would be expected on the basis of an optimal investment model. On the face of it, these results seem to contradict our findings—which indicate that experience leads to less risk taking. But Kaufmann et al. had a different goal. Their intention was to devise a way to make investors aware of the various possible outcomes of their investments. Simulating experience of a wide range of outcomes allows investors to evaluate their investments prospectively and at no cost. This approach seems to boost investors' knowledge and confidence in their choices, resulting in more risk taking. In our study, investors did not have the luxury of simulated experience of potential investments; they had to experience the full consequences of their investments to learn—retrospectively—about the market. Bradbury, Hens, and Zeisberger (2014) also found that simulated experience helps people to understand financial uncertainty better, to improve their decisions, and to avoid regret.

The power of simulated experience is enormous (see Hertwig, Hogarth, & Lejarraga, 2018, for a conceptual discussion of experience). Professionals learning to deal with complex situations also use simulations when real experience is costly and dangerous. Training aircraft pilots, for example, relies heavily on simulation. Some universities use simulated hospitals to teach medical students about the surgical procedures to be followed in operations or to teach them how to break bad news. Mistakes made in the context of piloting airplanes or performing heart surgery can clearly be costly. But so can mistakes made by those investing under financial uncertainty. Whereas simulations for pilots and surgeons help them to acquire motor and procedural skills, the crucial factor in the investment case is the simulation of uncertainty: investors need to experience the range of possible outcomes and the potential frequencies of their occurrence. When events are frequent, they need to experience that frequency; when events are rare, their scarcity must be felt. Financial simulations can also recreate collapsing markets for those "Boom babies" who fail to acknowledge the risks lurking in the fog of financial uncertainty.

Research is beginning to reveal some of the many practical implications of simulating financial uncertainty. The European Council Directive 2004/39/EC on markets in financial instruments (European Parliament and European Council, 2004) requires banks and other financial institutions to assess the degree of uncertainty that their clients are willing to accept. Our experimental results suggest that the willingness to be exposed to financial uncertainty differs dramatically depending on whether investors have learned about the uncertainty from description or from experience. Moreover, it is highly dependent on the outcomes experienced (e.g., booms or busts). Financial institutions currently use Likert scales to assess their customers' risk attitudes using a simple description-based approach. The benefits of providing customers with simulated experience of the financial markets and only then assessing their risk attitude could be enormous for them and society. Isaac Asimov described the crippling effects of traumatic experience—it is time to explore its empowering effects.