Belief Digitization in Economic Prediction

Samuel G. B. Johnson[1] & Faith Hill[2]
(sgbjohnson@gmail.com, fahill@vassar.edu)
[1]Department of Psychology, Yale University, New Haven, CT 06520 USA
[2]Department of Psychological Science, Vassar College, Poughkeepsie, NY 12604 USA

Abstract

Economic choices depend on our predictions of the future. Yet, at times predictions are not based on all relevant information, but instead on the single most likely possibility, which is treated as though certainly the case—that is, *digitally*. Two sets of studies test whether this digitization bias would occur in higher-stakes economic contexts. When making predictions about the future asset prices, participants ignored conditional probability information given relatively unlikely events and relied entirely on conditional probabilities given the more likely events. This effect was found for both financial aggregates and individual stocks, for binary predictions about the direction and continuous predictions about expected values, and even when the "unlikely" event explicitly had a probability as high as 30%; further, it was not moderated by investing experience. Implications for behavioral finance are discussed.

Keywords: Judgment & decision-making; probabilistic reasoning; explanatory reasoning; behavioral economics.

Introduction

Investors aim to buy low and sell high. Alas, this adage requires investors to predict the future—a feat known to be difficult for mortals (and even for economists).

People are famously biased in making predictions (Kahneman & Tversky, 1973), relying on a variety of useful but fallible heuristics. In economic contexts, a particularly worrisome bias would be *belief digitization*, as found in some other contexts (Johnson, Merchant, & Keil, 2015; Murphy & Ross, 1994). That is, when a reasoner is presented with data more consistent with one hypothesis than another, the reasoner acts as though the higher-probability hypothesis is *certainly* true when making predictions following from the hypothesis.

For example, in one study (Johnson et al., 2015), participants read about a pond that had ecological problems explainable either by an infestation of one type of snail (a simple explanation), or by an infestation of two types simultaneously (a complex explanation). The simple explanation was, reasonably, seen as more likely (about a 66% chance) than the complex explanation (about 34%). Yet, when using those explanations to make further predictions (e.g., about bacteria proliferation), people ignored this uncertainty. Manipulating the probability of bacteria proliferation given the simple explanation had a large effect on predictions about proliferation, but manipulating the probability given the complex explanation had no effect at all. People *digitized* the simple explanation, tacitly assigning 100% of their probabilistic weight to that possibility. Even though people often *explicitly* quantify uncertainty, this uncertainty does not propagate to subsequent computations but is instead rounded, in effect, to 0 or 1.

Such findings pose a challenge to probabilistic theories of cognition that treat humans as Bayesian thinkers who integrate across possibilities rationally (e.g., Anderson, 1991). Nonetheless, in many contexts, this strategy may a reasonably adaptive way to solve an otherwise intractable problem. The inference in this case (from ecological problems to snail infestation to the probability of bacteria proliferation) involves a fairly short chain of reasoning, yet people treated the first step in the inference as certain when making the second step. But we often rely on lengthy chains of reasoning, and propagating uncertainty through the entire chain may well be beyond our cognitive limits. If we must limit the complexity of these computations by prohibiting the consideration of multiple possibilities at each stage (e.g., thinking only about the consequences of the one-snail explanation or the two-snail explanation, but not integrating across both), then it is best to focus on the single *most* likely possibility. A person could do worse than this kind of belief digitization, even as it leads us astray relative to the optimal answer.

The current studies test whether such a digitization bias would influence judgments in economic contexts. In particular, digitization could affect predictions about future asset prices. Consider the impact of some piece of news, such as information about the government budget. Such information often has uncertain implications for future valuations, so rational investors would assign distributions over these possible futures and value assets according to their expected value. If voters elect a conservative populist (to take an example that is, of course, entirely hypothetical), this introduces uncertainty about the probability of fiscal stimulus. Perhaps there is a 70% probability of stimulus (with one set of implications for future valuations) and a 30% probability of fiscal austerity (with a different set of implications). Investors should rationally incorporate both possibilities into their valuations of the market, with a 70% weight on one possibility and a 30% weight on the other. Yet, if investors digitize, they would treat the likely event as *certain* when predicting the future value of the market. Rather than considering both possible futures, they would value assets assuming only the single most likely future.

Although previous studies using non-financial stimuli are consistent with this possibility, it is not clear that digitization effects would generalize to these contexts. First, people are likelier to rely on multiple categories in category-based prediction tasks when the categories are

dangerous or threatening rather than emotionally neutral (Zhu & Murphy, 2013). If people adopt a more reflective, normative strategy under higher-stakes situations, perhaps they also do so when making economically relevant predictions. Second, and related, people are sometimes more rational when making decisions rather than logically equivalent inferences (Johnson, Zhang, & Keil, 2016). These two factors could lead people to integrate probabilities across potential futures.

Two sets of studies test whether people nonetheless make digitized predictions in economic contexts. Experiment 1 provides an initial test, asking participants to make probabilistic predictions about the direction of asset prices, given uncertain information. Experiment 2A tests whether digitization effects would occur only for binary predictions (i.e., will a price go up or not?) or would instead extend to predictions of expected value on a continuous scale. Finally, Experiment 2B tests a possible boundary condition by giving participants explicit posterior probabilities for the market's future direction. After examining these studies individually, we pool the data to examine whether expertise can combat digitization biases. In the General Discussion, we assess the implications of these findings for behavioral finance.

Experiment 1

Participants in Experiment 1 made predictions about the future prices of financial assets in light of information with uncertain implications. Experiment 1A looked at predictions about market aggregates (e.g., the S&P 500) and Experiment 1B looked at predictions about individual stocks (e.g., GE). Given that individual stocks seem to be priced more efficiently than the market as a whole (see Shiller, 2005), perhaps digitization mechanisms do not apply as robustly to predictions about individual stocks.

Participants predicted the probability of an increase in an asset price, denoted as $P(Z)$, based on information about two mutually exclusive possibilities, A and B. For instance, A might represent a stimulatory fiscal policy and B an austere fiscal policy. Participants were given information implying that $P(A) > P(B) > 0$, so that both A and B are possible even as A is likelier—the government may not have made a decision on its fiscal policy, but a stimulus is probable. In addition, participants were given information about the probability of Z conditional on A and B—$P(Z|A)$ and $P(Z|B)$. If people take both more and less likely possibilities into account, then they should rely on both $P(Z|A)$ and $P(Z|B)$ when predicting $P(Z)$. In contrast, if people digitize, relying only on the single most likely possibility, then only manipulations of $P(Z|A)$ should propagate to predictions of $P(Z)$.

Methods

We recruited 200 participants from Mechanical Turk, divided between Experiments 1A and 1B.

Participants each completed three items. For each item, participants read about an uncertain event, where one

possibility (A) seemed more likely than the other (B), given the available information. These likely and unlikely possibilities differed in their implications for future prices of financial assets. In the *high/low* condition, the more likely event A would have a high chance of leading to an increase in asset values (i.e., $P(Z|A)$ is high), whereas the less likely event B would have a low chance of leading to an increase (i.e., $P(Z|B)$ is low). One item in the *high/low* condition of Experiment 1A read:

Imagine that a foreign government is deciding what level of spending to adopt in the next fiscal year.

If they increase public spending, the value of the US stock market is *likely* to go up.

If they decrease public spending, the value of the US stock market is *unlikely* to go up.

Suppose that the leader of this government is concerned about the distribution of wealth in the country and is considering increasing public spending.

Participants reading this information should conclude that possibility A (public spending increase) was likelier than possibility B (public spending decrease). For instance, an investor might assign an 80% probability to possibility A and a 20% probability to possibility B.

Whereas $P(Z|A)$ was high and $P(Z|B)$ was low in the *high/low* condition, both $P(Z|A)$ and $P(Z|B)$ were low in *low/low* condition:

If they increase public spending, the value of the US stock market is *unlikely* to go up.

If they decrease public spending, the value of the US stock market is *unlikely* to go up.

Rationally, the probability of a price increase is lower in the *low/low* than the *high/low* condition, since possibility A has positive (indeed, high) probability of being correct. Thus, both rational and digitizing investors would distinguish between the *low/low* and *high/low* conditions.

A third condition, however, generates different predictions for these two groups of investors. In this *low/high* condition, $P(Z|A)$ is low and $P(Z|B)$ is high:

If they increase public spending, the value of the US stock market is *unlikely* to go up.

If they decrease public spending, the value of the US stock market is *likely* to go up.

In this *low/high* condition, a rational investor would judge the probability of a price increase likelier than in the *low/low* condition, since possibility B has positive probability (albeit lower than A). In contrast, if people digitize, tacitly assigning 0% weight to B, then the *low/high* and *low/low* conditions would not differ.

After reading each item, participants rated $P(A)$ and $P(B)$ (e.g., "Government intends to increase public spending" and "Government intends to decrease public spending") on a 0 to 100 scale. This measure was taken to ensure that people did not *explicitly* place a 0% weight on B, in which case rational prediction and digitization do

not diverge in their predictions. Further, explicitly quantifying uncertainty in the task produces a task demand to incorporate this uncertainty into predictions, working against our hypothesis.

Finally, on the same page, participants predicted P(Z) ("What do you think is the probability that the US stock market will go up?") on the same scale used above.

Experiments 1A and 1B differed only in the asset being judged. In Experiment 1A, the asset was the overall value of the US stock market and in Experiment 1B, the asset was the share price for stock in specific corporations.

The three probability conditions were counterbalanced with three different vignettes (one on fiscal policy, one on monetary policy, and one on regulatory policy) using a Latin square. The items were presented in a random order.

After the main task, participants completed 10 check questions and were excluded from analysis if they answered more than one-third incorrectly ($N = 19$). Another 14 participants were excluded because their total probability ratings for at least one item were not between 80% and 120%. However, including these two types of participants does not alter the pattern of results. Finally, 49 participants were excluded because they did not rate the A more likely than B for at least one of the items, since our predictions are predicated on participants' belief that $P(A) > P(B)$. (See Experiment 2B for a version that did not require the latter two categories of exclusions.)

Table 1: Results of Experiment 1

Condition		Predicted P(Z)	
P(Z|A)	P(Z|B)	Exp. 1A	Exp. 1B
Low	Low	28.8 (28.7)	30.1 (27.4)
High	Low	73.0 (17.7)	75.6 (12.8)
Low	High	30.3 (26.5)	32.3 (26.0)

Note. Entries are probabilistic predictions, expressed as percentages. SDs in parentheses.

Results and Discussion

As shown in Table 1, participants digitized in both Experiments 1A and 1B.

In both experiments, participants relied on P(Z|A) for in their predictions of future asset prices. The *high/low* and *low/low* conditions differed only in P(Z|A), and these conditions differed sharply in predictions [$t(62) = 10.38$, $p < .001$, $d = 1.85$ and $t(54) = 10.98$, $p < .001$, $d = 2.13$ for Experiments 1A and 1B]. Thus, people take account of high-probability possibilities when making predictions—consistent with both rational and digitizing strategies.

These two strategies differ, however, in their predictions about the *low/high* condition. This condition differs from the *low/low* condition only in P(Z|B). Thus, if people take account of less likely possibilities, they should differentiate between these two conditions, but if they digitize, these conditions should be rated similarly.

Supporting the latter possibility, there was no difference between these conditions for either experiment [$t(62) = 0.50$, $p = .62$, $d = 0.06$ and $t(54) = 0.47$, $p = .64$, $d = 0.08$, respectively]. Since we predicted null effects for these comparisons, we also computed Bayes Factors (Rouder et al., 2009; scale factor 1), which strongly favored the null hypothesis [$BF_{01} = 8.9$ and 8.5, respectively]. Further, based on participants' other judgments, we can calculate the normative mean difference between the *low/high* and *low/low* conditions (which would produce Ms = 37.6 and 40.3 for *low/high*, respectively). In both cases, the actual differences were less than these benchmarks [$t(62) = 2.41$, $p = .019$, $d = 0.30$ and $t(57) = 1.73$, $p = .089$, $d = 0.23$].

Together, these results show that people fail to account for low-probability possibilities when making economic predictions. This was true both when predicting the overall level of financial aggregates as well as the value of stock shares in individual companies.

That said, one may raise some concerns about these results. Perhaps of most concern, the information given in the problem could plausibly have implied near-certainty in its predictions (e.g., "the leader of this government is concerned about the distribution of wealth in the country and is considering increasing public spending"). To assess this possibility, we looked at participants' explicit judgments about P(A) and P(B). Unlike their implicit judgments, which assigned essentially a 100% probability to A, participants assigned more reasonable probabilities when asked explicitly (83% and 82% in Experiments 1A and 1B, respectively). Nonetheless, we address this concern head-on in Experiment 2B.

Experiment 2

Experiment 2 examines two possible boundary conditions on belief digitization in economic contexts.

First, Experiment 1 asked for predictions about the probability of binary events (increases or decreases in value). The direction of future gains or losses is likely to be the dominant factor in real investing decisions, but the extent of these predicted gains or losses is also important. In some cases, people are better at reasoning about continuous rather than binary events (e.g., in covariation-based causal reasoning; Alloy & Tabachnik, 1984). Experiment 2 therefore tests whether digitization effects extend to continuous judgments of expected value.

Second, participants in Experiment 1 arrived at estimates of P(A) and P(B) on the basis of other, ambiguous information, as has been the case in most prior work finding digitization effects (Johnson et al., 2015; Murphy & Ross, 1994). Would such effects occur even when the problem explicitly quantifies the uncertainty? Experiment 2B addresses this question by assigning a 30% probability to the less likely event. This further rules out the concern that participants may have rationally ignored a low probability. This also addresses the concern that participants in Experiment 1 may have *actually* assigned extremely low explicit probabilities to the

unlikely events and *reported* biased explicit judgments due to task demands. In that case, it would not be their implicit judgments that are biased (for interesting reasons) but their explicit judgments (for deflationary reasons).

Methods

We recruited 200 participants from Mechanical Turk, divided between Experiments 2A and 2B.

The procedure of Experiment 2A was identical to Experiment 1A, except that the dependent measure was a continuous price, on either the NASDAQ, DJIA, or S&P 500, instead of the probability of a directional change. Participants were given approximately the current value of one of these indices (e.g., "Suppose the current value of the United States stock market, as indexed by the S&P 500, is $2,000") and then asked to predict the future value of that index ("Please estimate what you think the value of the S&P 500 will be 3 months from today") on a scale ranging from 10% lower than its current value (e.g., $1800) to 10% higher than its current value (e.g., $2200).

Experiment 2B was identical, except explicit probabilities were given for *A* and *B* ("Analysts say there is a 70% chance that this foreign government will increase public spending, and a 30% chance that it will decrease public spending") and thus participants were not asked to rate the probabilities of these events.

After the main task, participants completed 10 check questions and were excluded from analysis if they answered more than one-third incorrectly ($N = 14$). Another 7 participants from Experiment 2A were excluded because their total probability ratings for at least one item were not between 80% and 120%. Finally, 31 participants from Experiment 2A were excluded because they did not rate P(*A*) higher than P(*B*) for at least one of the items. Since Experiment 2B explicitly provided these probabilities, participants were not excluded for this reason. Analyses including all participants found similar results for both experiments.

Table 2: Results of Experiment 2

Condition		Predicted Change	
P(*Z*|*A*)	P(*Z*|*B*)	Exp. 2A	Exp. 2B
Low	Low	–0.21%	0.33%
		(2.96%)	(3.44%)
High	Low	2.86%	3.57%
		(3.50%)	(2.89%)
Low	High	–0.32%	0.37%
		(3.18%)	(3.54%)

Note. Entries are predicted changes in stock market value. SDs in parentheses.

Results and Discussion

As shown in Table 2, participants once again digitized.

In Experiment 2A, participants predicted a significantly higher change in asset price in the *high*/*low* condition than

in the *low*/*low* condition [$t(56) = 6.59$, $p < .001$, $d = 0.95$]. Thus, participants did consider the likely event when making predictions. However, participants again ignored the less-likely event *B*, since they did not use P(*Z*\|*B*). Predicted changes did not differ across the *low*/*high* and *low*/*low* conditions [$t(56) = –0.23$, $p = .82$, $d = –0.04$, $BF_{01} = 9.4$]. Further, as in Experiment 1, the difference in predicted changes between the *low*/*high* and *low*/*low* conditions was marginally lower than it normatively should have been (for a *low*/*high* mean of 0.59%), based on the other judgments [$t(56) = 1.93$, $p = .059$, $d = 0.26$]. Thus, digitization occurs even for predictions made on a continuous scale rather than probabilities of binary events.

Experiment 2B provided explicit probabilities of P(*A*) and P(*B*), ensuring that the "unlikely" event *B* had a rather serious chance of occurring (30%). Nonetheless, the results are similar to Experiment 2A. While participants again differentiated between the *high*/*low* and *low*/*low* conditions [$t(90) = 7.06$, $p < .001$, $d = 1.02$], they did not differentiate between the *low*/*high* and *low*/*low* conditions [$t(90) = 0.08$, $p = .93$, $d = 0.01$, $BF_{01} = 12.0$]. Further, the difference between conditions was dramatically lower than it normatively should have been (for a *low*/*high* mean of 1.75%) [$t(90) = 3.41$, $p < .001$, $d = 0.36$]. Thus, people are willing to ignore even a 30% probability of an event's occurrence when predicting assets' future value.

One possible objection is that participants may have been giving an appropriate answer, depending on their interpretation of the question. That is, whereas participants' judgments of probabilities in Experiment 1 normatively should accommodate the possibility of lower-probability events (as is provable from the laws of probability), predictions of future value may be reports of the most likely single value, rather than the expected value. In fact, the single most likely value of the market does depend greatly on P(*Z*\|*A*), given that *A* is the single most likely event, but to a much lesser degree on P(*Z*\|*B*).

However, there are two reasons to doubt this interpretation. First, although the maximum-probability and expected value interpretations of the question are both reasonable, participants would have to uniformly adopt the maximum-probability interpretation to produce our results. That is, if half of participants took the maximum-probability interpretation and therefore did not use P(*Z*\|*B*) in their predictions, the other half of participants were still making a mistake in failing to use P(*Z*\|*B*).

Second, even though ignoring P(*Z*\|*B*) is appropriate in estimating the maximum-probability value of the price, people tend to probability-match rather than maximize in tasks of this sort. For example, suppose there is one button that has a 70% chance of giving a positive payoff and another button that has a 30% chance of giving the payoff. If you are supposed to predict which button will produce the payoff on a given trial, the rational thing to do would be to choose the 70% button every time. In fact, people will predict the 30% button a significant fraction (roughly 30%) of the time. The only way to reconcile this

result with the current task is to assume that participants have tacitly assumed that the 30% probability event has a 0% chance of occurrence and can thus be safely ignored.

Overall, Experiment 2 helps to address alternative interpretations of Experiment 1, and shows that people do not need to arrive at estimates of event probabilities themselves in order to digitize them. Together, these two experiments demonstrate that digitization effects may be a pervasive force in investors' judgments of future value.

Expertise Effects

Amateur investors are often referred to as "noise traders" in financial models (Shleifer & Summers, 1990) and the behavior of these models depends greatly on these traders' beliefs and choices (Shleifer, 2000). Although professional investors may use different strategies from amateurs (but see Tuckett, 2011), the behavior of amateurs contributes to market dynamics and is therefore important to characterize. Given that our participants are laypeople, but some have investing experience whereas others do not (about half of Mechanical Turk participants own financial assets and about half have taken at least one finance course; Johnson & Tuckett, 2017), would we see expertise effects within this sample?

Participants in both studies were asked to rate their investing experience and knowledge. If people who have more domain expertise are likelier to consider low-probability events in making predictions, then the effect of $P(Z|B)$—converted to a z-score to aggregate data across studies—ought to be larger for individuals with more experience and knowledge. This was not the case, either for self-reported experience [$r(264) = .02, p = .72$] or for knowledge [$r(264) = -.02, p = .70$].

This result, although preliminary, suggests that domain expertise may not be sufficient to overcome digitization effects even in a context like financial prediction that has obvious real-world implications. This does not necessarily undermine the argument often advanced by economists that highly incentivized individuals can avoid such biases, nor the possibility that in market contexts corrective forces can emerge if a subset of investors exploit the suboptimal behavior of others. Nonetheless, this result does suggest that quite extensive expertise—outside the range of experience of our sample—is necessary for such mechanisms to apply. Digitization appears to be a robust cognitive bias at the individual level, and is therefore likely to cause suboptimal performance from investors at a variety of skill levels unless explicitly checked.

General Discussion

Economic choices, such as investment allocations, depend on our predictions about the future. Rational predictions require us to integrate over multiple uncertain possibilities; failing to do so leads to overconfident predictions that are too near to 0% or 100%. Yet, participants in our studies consistently failed to account for lower-probability possibilities in making predictions.

Digitization is broadly consistent with *conviction narrative theory* (e.g., Tuckett, 2011), the idea that decisions in highly uncertain environments are made by constructing a narrative to explain the past, projecting this narrative into the future, and using emotional reactions to the projected narratives to guide choices. For example, amateur investors use company performance news to guide predictions and choices even once the market has had time to "price in" that information, particularly if the news concerns the future rather than the past (Johnson & Tuckett, 2017). This follows from narrative thinking, since narratives are emotionally valenced and temporally oriented. Another important feature of narrative thinking is that it is linear—it concerns a single sequence of events rather than a web of possibilities. The current work shows that people indeed focus on a single narrative to explain the past and project the future, rather than integrating across multiple possible narratives.

In addition to this theoretical contribution, these results have two kinds of practical implications. First, these biases may persist at the market level, leading to mispricing. A previous study examined explanatory biases in the context of Wall Street Journal headlines (Johnson, 2016). For instance, one headline read "ECB Move Crushes Hopeful Markets." There had recently been a downturn in European markets because the European Central Bank (ECB) had chosen to follow a less inflationary monetary policy than markets had expected. Had investors been "counting on" monetary expansion, tacitly assigning it a 100% probability? Or had the market priced in this uncertainty already (as mainstream financial theory suggests; e.g., Malkiel & Fama, 1970)?

Investors made an uncertain diagnosis (the meaning of the ECB chair's statements) and a prediction based on that diagnosis (the implications for monetary policy). Normatively, uncertainty about the interpretation of ECB statements should propagate to any predictions based on such inferences. If the market digitizes at an aggregate level, however, this could have led the market to react strongly to disconfirmed expectations: If the expectations are formed based on uncertain information treated as certain, the market would be overconfident. This could lead prices to be either too high or too low—and indeed to oscillate between those extremes. New information may cause an investor to rationally move from predicting, say, a 70% probability to a 30% probability of some event. If these probabilities are treated as 100% and 0%, respectively, this will lead to a much larger shift in asset valuation than is justified by fundamentals.

That said, such an interpretation of these experimental results is controversial, as are many efforts in behavioral finance to generalize from individual behavior to market-level behavior (Shleifer, 2000). A common rejoinder from a neoclassical approach is that behavioral biases can often be neutralized in market contexts. Markets create incentives for accuracy, which are often lacking in behavioral experiments. Markets allow for specialization

so that investors can learn over time to correct their biases (though our expertise analysis suggests that such learning is non-trivial). And perhaps most importantly, self-correcting market-level phenomena may emerge. If some, potentially small, subset of investors comes to understand the biases of other investors, they can trade against that bias and capitalize on others' irrationality. Because of these mechanisms, market prices may be less likely to be seriously afflicted by digitization biases than are individual investors' decisions. However, given that stock markets appear to be more volatile than is justified by fundamentals (de Bondt & Thaler, 1985; Shiller, 1981), digitization of hypotheses could be a partial explanation of this excess volatility. Nonetheless, this issue will not be adjudicated by lab experiments alone.

Second, however, these biases are troubling not only because of potential market inefficiencies they might cause. Even if financial markets do have self-correcting forces that lead experienced investors to profit from the errors of novice investors, the losses of these novices are still cause for concern. Digitized predictions of asset prices can lead to several errors in the investing strategies of amateur "noise traders." First, if one has a high valuation of an asset relative to the market, one may overpay for that asset. For instance, if one is purchasing a house and has an unreasonably high valuation of that house, the buyer may not adequately negotiate the price. Second, extreme asset valuations may potentially lead to suboptimal patterns of diversification. A very bullish assessment of the tech industry accompanied by a very bearish assessment of the financial sector may lead one to prioritize the former over the latter, when a diversified investor would spread her exposure over all sectors. Third, if one's valuations are oscillating faster than the market's valuations, this may lead investors to overtrade, which leads to portfolio value loss due to transaction costs. Finally, overconfident predictions about asset prices may lead investors to inadequately hedge: If the cost of insurance is high relative to the perception of the risk being insured, there is less incentive to insure. This may lead some investors to be overexposed to unexpected downturns in the market—why hedge against something that is deemed, at some level, to be impossible?

Nassim Taleb (2010) warns of "black swans"—"unknown unknowns" of high impact that we discount on the basis of their low probability. Our participants exemplified this problem, and indeed took it one step further: An event with a 30% chance is not exactly on the tail of a distribution. Investors would do well to consider all the swans—both black and white.

References

Alloy, L.B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, *91*, 112–149.

Anderson, J.R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.

De Bondt, W. F. M., & Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance*, *40*, 793–805.

Johnson, S.G.B. (2016). Explaining December 4, 2015: Cognitive science ripped from the headlines. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Merchant, T., & Keil, F.C. (2015). Predictions from uncertain beliefs. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Johnson, S.G.B., & Tuckett, D. (2017). *Narrative decision-making in investment choices: How investors use news about company performance*. In preparation.

Johnson, S.G.B., Zhang, M., & Keil, F.C. (2016). Decision-making and biases in causal-explanatory reasoning. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus, and Giroux.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237–251.

Malkiel, B.G., & Fama, E.F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, *25*, 383–417.

Murphy, G.L., & Ross, B.H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, *27*, 148–193.

Rouder, J.N., Speckman, P.L., Sun, D., Morey, R.D., & Iverson, G. (2009). Bayesian *t*-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *The American Economic Review*, *71*, 421–436.

Shiller, R. J. (2005). *Irrational exuberance* (2nd Ed.). Princeton, NJ: Princeton University Press.

Shleifer, A. (2000). *Inefficient markets: An introduction to behavioral finance*. Oxford, UK: Oxford University Press.

Shleifer, A., & Summers, L. H. (1990). The noise trader approach to finance. *Journal of Economic Perspectives*, *4*, 19–33.

Taleb, N. N. (2010). *The black swan: The impact of the highly improbable* (2nd Ed.). New York, NY: Random House.

Tuckett, D. (2011). *Minding the markets: An emotional finance view of financial instability*. London, UK: Palgrave Macmillan.

Zhu, J., & Murphy, G. L. (2013). Influence of emotionally charged information on category-based induction. *PLoS ONE*, *8*, e54286.