Decision weights for experimental asset prices based on visual salience

Devdeepta Bose,[1] Henning Cordes,[2] Sven Nolte,[3]
Judith C. Schneider,[4] Colin F. Camerer[1] *

May 8, 2020

Abstract

Using a machine-learning algorithm that can predict visually salient portions of images, we construct decision weights based on salient parts of a stock price chart. We analyze these weights in three experimental studies that vary in the realism of the price path images and task complexity. We find that these decision weights are predictive of future investments. We conclude that visual salience captures attention paid to historical returns. Visual salience goes beyond overweighting returns at the tails of the historical distribution or with respect to their difference to a reference return as in the models of Barberis et al. (2016) and Bordalo et al. (2013). Moreover, we find that visual salience affects investment decisions independently from recency effects.

JEL classification: G4, G12
Keywords: Decision Weights, Eye-Tracking, Machine-Learning Algorithm, Visual Salience

---

*We thank Nicholas Barberis, Lawrence Jin, Alex Imas, Elise Payzan-LeNestour, and seminar and conference participants at Warwick Business School, Radboud University, University of Münster, Caltech, Southwestern Economic and Behavioral Economics 2019, Experimental Science Association 2019, Mid-Adlantic Meeting on Behavioral and Experimental Economics 2019, and the Boulder Summer Conference on Consumer Financial Decision Making 2019 for helpful comments and suggestions.

[1]California Institute of Technology, Corresponding author is devbose@caltech.edu
[2]University of Münster, Universitätsstr. 14-16., 48143 Münster, Germany
[3]Institute for Management Research, Radboud University, Nijmegen, The Netherlands
[4] Leuphana University Lüneburg, Universitätsallee 1, 21335 Lüneburg, Germany

1 Introduction and Motivation

Benartzi and Thaler (1995) in their seminal work on the equity premium puzzle propose that investors form expectations about future returns of an investment by using the historical distribution of annual stock market returns. Newer models suggest applying prospect theory to the historical return distribution (Barberis et al., 2016) or weighting historical returns by their salience compared to a benchmark return (Bordalo et al., 2013; Cosemans and Frehen, 2017). Both these types of models assign decision weights to individual returns from the past to form an expectation of future returns.

These classes of decision weights have psychological foundations. We add to this literature by relying on neuroscientifically motivated measures of attention for the decision weights. Attention is a scarce resource for investors during financial decision-making. Events that are likely to be attention-getting are thought to affect what investments people make (Barber and Odean, 2007; Hillert et al., 2014; Mohrschladt and Schneider, 2018). In decision weights as in Barberis et al. (2016) or Bordalo et al. (2013), the role of attention is accounted for by overweighting returns at the tails of the return distribution. A variety of evidence suggests that investors pay more attention to these extreme returns. These models are undoubtedly a portable approximation of important first-order salience effects. However, it is conceivable that the *visual* salience of price charts captures attentional effects toward certain types of returns features which are predictive of investment, beyond overweighting the tails.

A common representation of a stock's performance is a visualization of historical prices in a chart. We explore whether the decision weights investors assign to individual past returns are driven by the allocated attention. We propose that initial attention towards specific returns is determined by the visual salience of respective points on the visualized price path. An effective way to measure visual attention is by recording visual fixations and eye movements (saccades).

Early studies using measures of visual fixation in economics showed that what people looked at, and what they did not, could explain deviations of strategizing from equilibrium predictions (Camerer et al., 1993).

More recent research has extended this insight using a neuroscientifically-inspired algorithm (Cornia et al., 2018, Salience Affective Model (SAM)). SAM is one of the latest in a careful series of increasingly-improved algorithms which are trained to predict which points of an image are more visually salient than others. SAM is trained on actual eye-tracking data. Since price charts are a ubiquitous source of information for investors, we explore whether SAM can identify returns in these price chart images that are more (or less) visually salient, and assign higher (or lower) decision weights to these returns. Investors' valuations of historical returns can then be modeled by using a preference function consisting of the product of the salience-derived decision weights and a value function. We propose that visual salience enters the function by affecting the decision weights. This approach enables us to compare the visual salience model with the existing models of Barberis et al. (2016) and Bordalo et al. (2013).

We collect data from three experimental studies to evaluate the performance of the visual salience model. The three studies vary with respect to the degree of realism of the depicted price charts and the degree of control over the predictions of the different models. The purpose of Study I is to explore if the visual salience model has predictive power for investments based on empirical price charts that investors are exposed to in reality. We find that the model is, in contrast to the models of Barberis et al. (2016) and Bordalo et al. (2013), able to significantly predict invested amounts even without specifying a particular value function.

Study II is designed to isolate the impact of the effect of temporal ordering of returns on predictability. Barberis et al. (2016) accomodate a recency effect by adding an exponential decay function to CPT which accounts for overweighting more recent returns (see also Bordalo et al., 2019). The Study II data suggest that there are at least two coexisting, independent effects:

a recency effect, and a visual attention and salience effect. Depending on the context, investors might either rely more on recent returns or on the most visually salient returns.

Study III considers the simplest possible depiction of price charts controlling for many factors which could influence investment decisions like image or task complexity. Moreover, in this study we have control over the environment and attention of the participants. We show that even in this simple and controlled setup, the visual salience model is able to predict future investments.

1.1 Connections to previous research

Our research is connected to four ongoing kinds of research in finance and economics. We summarize these briefly and elaborate on them (with detailed citations) in Section 1 of the Online Appendix.

First, several studies explore how the *framing or presentation* of historical returns seem to activate different biases and heuristics investors use, and and influence their decision making. We also show that presentation formats influence investment through their visual salience. Second, the large amount of information cheaply available to investors means that investors, and even institutional investors, cannot pay attention to everything. Many studies show that *heightened attention influence investment choices*. We do the same by predicting (and also measuring) visual attention in investment experiments. Third, behavioral models propose different ways in which past returns are weighted and integrated to form expectations about future returns, and drive investment. We *add a new model* of that type, using visual salience to weight past returns. Fourth, visual salience has also been used to explain choices in a couple of studies about consumer and strategic game choices. We use *similar measures to predict investment*, for the first time.

2 Theoretical Framework

Let X be a random variable with realizations $x_1, ..., x_m$ where $m \in M$ denotes either past realizations $k \in K$, or future realizations $s \in S$. M can be interpreted as the state space which describes the choice problem under risk. Since the entire state space M and possible future realizations $s \in S$ are unknown to the decision maker (DM), she relies on past realizations to help evaluate what is possible in the future. Barberis et al. (1998) argue that such an extrapolation is psychologically rooted in the "representativeness heuristic", a concept introduced by Kahneman and Tversky (1972, 1973). This heuristic leads the DM to think that the limited sample of past realizations is representative of the full distribution of possible outcomes, weighted by the likelihood that the DM associates with each past realization. In this paper's setting, investors use a combination of past returns and decision weights to evaluate the appeal of investing in the future. The investors evaluate X according to:

$$V(X) = \sum_{k \in K} \pi_k v(x_k), \tag{1}$$

where $v(\cdot)$ is a continuous, strictly increasing value function, and π_k represent the decision weights associated with historical realizations that occurred with objective probability p_k. In the following sections, we will explore how visual salience affects the decision weights that the investors assign to the past realizations, and how the model compares against the decision weights used in the canonical studies of Barberis et al. (2016) and Bordalo et al. (2013). We also explore how the results are impacted when value functions are replaced by correlations between returns and decision weights.[1]

[1] Consistent with most finance models, the analyses are based on returns. When confronted with price charts, investors might base their decisions on the visually accessible absolute changes in prices, rather than on returns. We explore this possibility further in Section 13 of the Online Appendix.

2.1 Decision Weights

Decision weights formalize how investors feel about probabilities (Wakker, 2010). In investment decision, it is common to use a two step procedure to determine how decision weights influence investment. First, a mathematical function that transforms objective probabilities into decision weights is developed. This weighting function is motivated by psychological insights or axioms (e.g. (Prelec, 1998) and has one or two parameters controllling its shape. In the second step, these parameters are statistically calibrated to fit observed behavior. This procedure underlies the decision weights in prospect theory (Tversky and Kahneman, 1992) or salience theory (Bordalo et al., 2012, 2013).

The visual salience decision weights do not follow this two-step procedure. Instead, objective probabilities are directly distorted by *predicted* attention weights. The only restriction is that the visual salience decision weights (attention weights $\times$ objective probabilities) have to add up to one. No further specific parametrization is needed, although a theory is needed to predict weights from visual attention. Recent research in neuroeconomics motivates the use of attention weights by showing that subjects in the laboratory attach more weight to attributes that they attend to more when making binary or multi-attribute choices, as well as choices among risky lotteries (Krajbich et al., 2010; Krajbich and Rangel, 2011; Mormann and Frydman, 2017; Mullett and Stewart, 2016).

To illustrate how CPT, salience theory, and the visual salience approach determine decision weights, consider the following example. Let x_t be the price for every period following a stationary AR(1) process with $x_{t+1} = \rho x_t + \varepsilon_t$ with $\rho = 0.6$. Three exemplary associated price paths for 15 time periods of historical realizations are depicted in the upper panel of Figure 1. All three paths are constructed from the identical 15 returns, only the temporal ordering of the returns differs. As a result, the visual appearance and the shapes of the price paths also differ. Associated returns are depicted in descending order below the price charts. Empirical

Figure 1: CPT and Salience

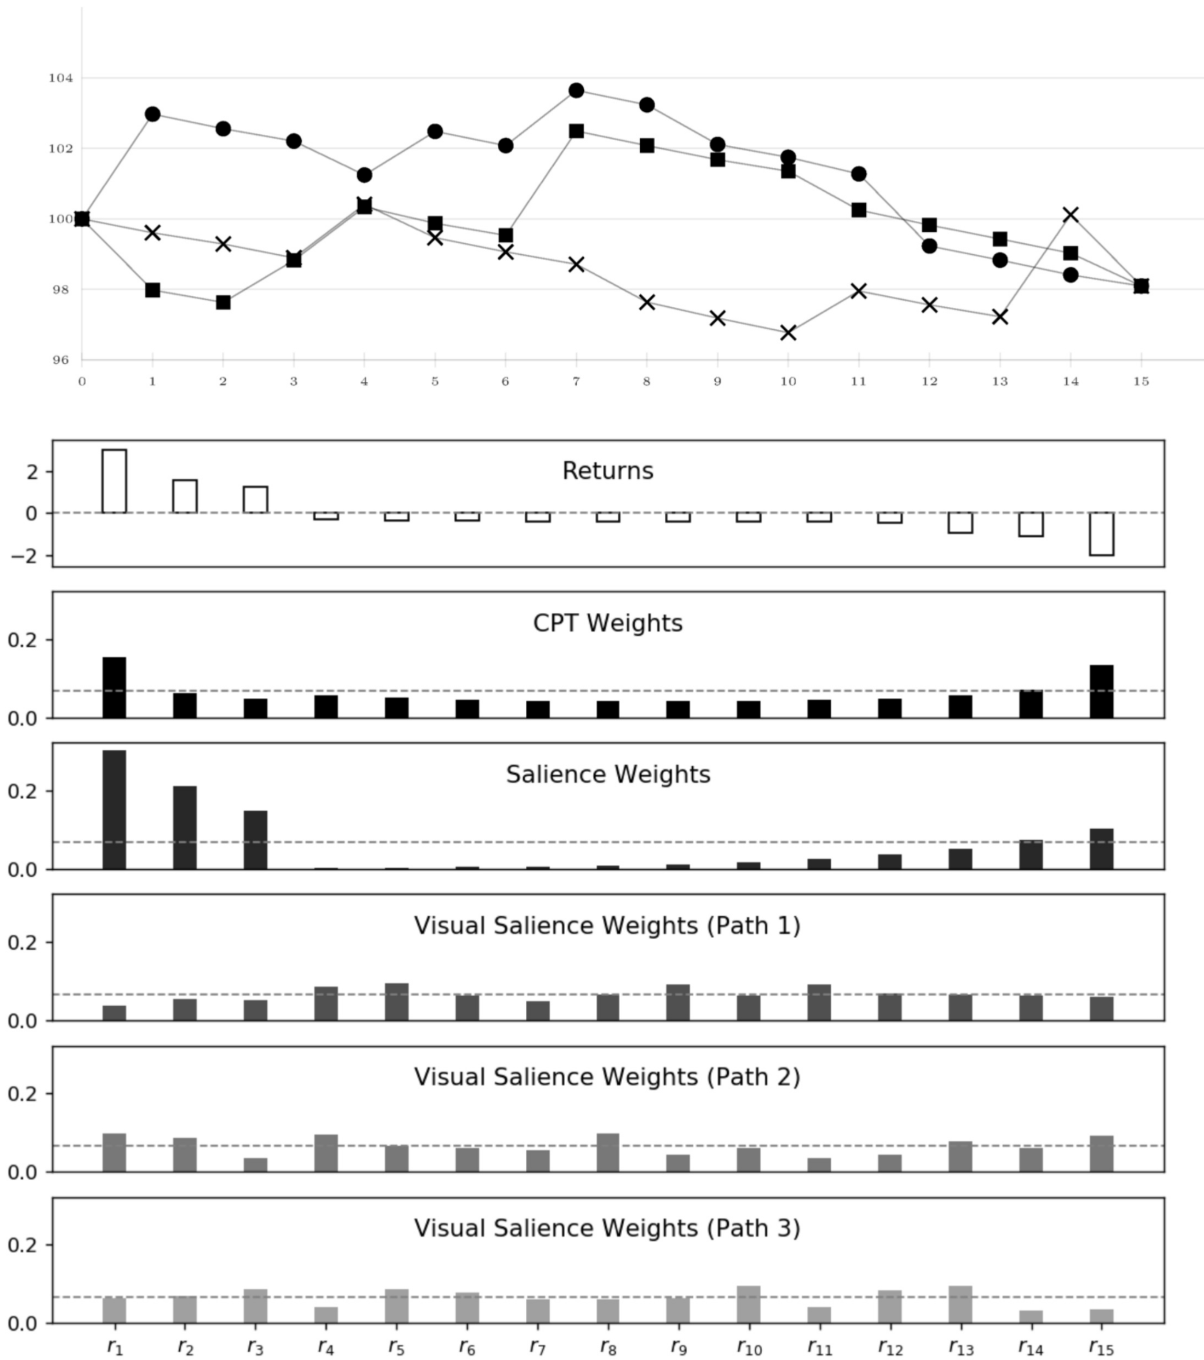

The figure shows three price paths of 15 periods historical realizations constructed from the returns in the graph below. The returns are the result of stationary process $x_{t+1} = 0.6x_t + \varepsilon_t$. The two graphs below present decision weights based on CPT and salience theory. The bottom three graphs are visual salience decision weights for path 1 (dots), 2 (crosses), and 3 (squares).

Figure 2: CPT and Salience

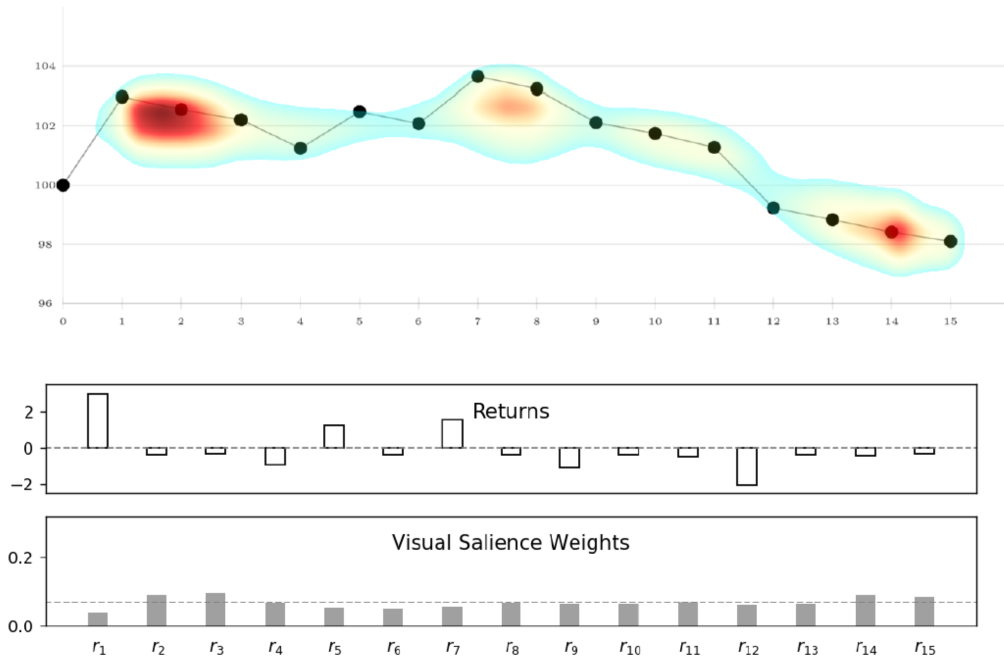

The figure shows the first price path from Figure 1 with an overlay of the heat map returned by SAM. The darker the overlay the more attention is allocated to that area of the price path. The returns in the order of occurrence and the resulting visual salience weights are depicted below.

frequencies for each of the 15 returns are $p_k = 1/15$, $k \in \{1, ..., 15\}$ and will be treated as objective probabilities.

Probability weighting as in CPT transforms the corresponding objective probabilities using the following functional. The historical returns are separated into n gains and m losses relative to a pre-specified point (e.g., the status quo), and sorted in ascending order. The objective probability of each return k is then transformed into a decision weight by applying: [2]

[2] We assume that $\pi^+(0) = \pi^-(0) = 0$ and $\pi^+(1) = \pi^-(1) = 1$.

Figure 3: CPT and Salience

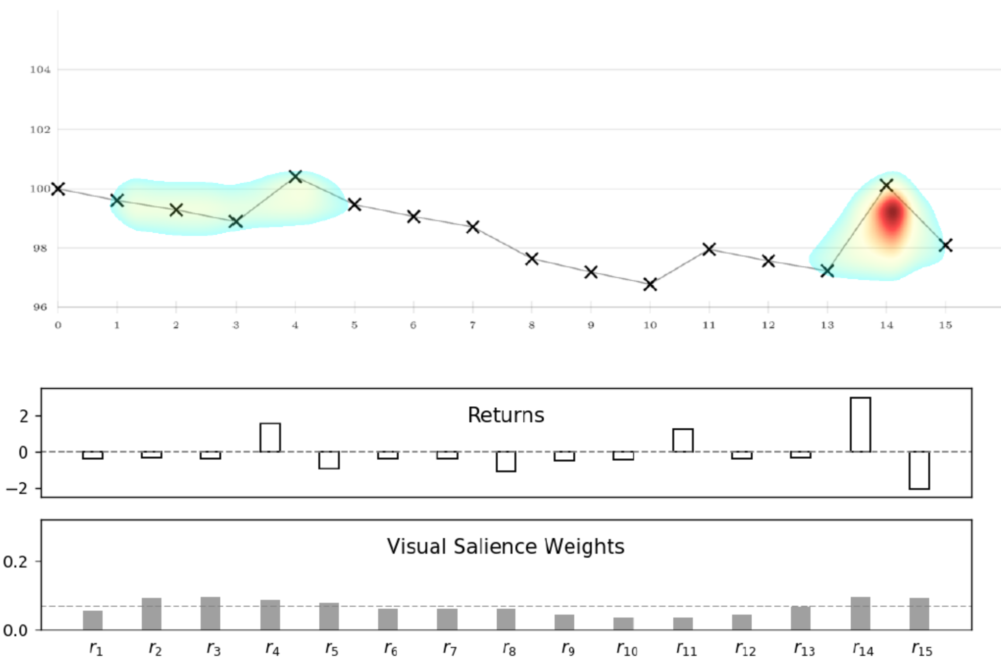

The figure shows the second price path from Figure 1 with an overlay of the heat map returned by SAM. The darker the overlay the more attention is allocated to that area of the price path. The returns in the order of occurrence and the resulting visual salience weights are depicted below.

$$\pi_k^{\text{CPT}} = \begin{cases} w^+(p_k + \ldots + p_n) - w^+(p_{k+1} + \ldots + p_n), \text{ for } 0 \leq k \leq n \\ w^-(p_{-m} + \ldots + p_k) - w^-(p_{-m} + \ldots + p_{k-1}), \text{ for } -m \leq k \leq -1 \end{cases} \quad (2)$$

$$\text{and } w^+(p) = \frac{p^{\delta_+}}{(p^{\delta_+} + (1-p)^{\delta_+})^{\frac{1}{\delta_+}}} \text{ and } w^-(p) = \frac{p^{\delta_-}}{(p^{\delta_-} + (1-p)^{\delta_-})^{\frac{1}{\delta_-}}} \quad \delta(0,1] \quad (3)$$

with n gains and m losses. Based on experimental evidence, probability weighting functions are predominantly inverse S-shaped, which means concave for small probabilities and convex for moderate and high probabilities (Abdellaoui, 2000). The original parametrization for this functional is $\delta^+ = 0.69, \delta^- = 0.61$ (Tversky and Kahneman, 1992), but the optimal parametrization depends on the underlying decision problem and elicitation method used (Wu and Gonzalez, 1996). Moreover, several alternative specifications of the probability weighting functional it-

Figure 4: CPT and Salience

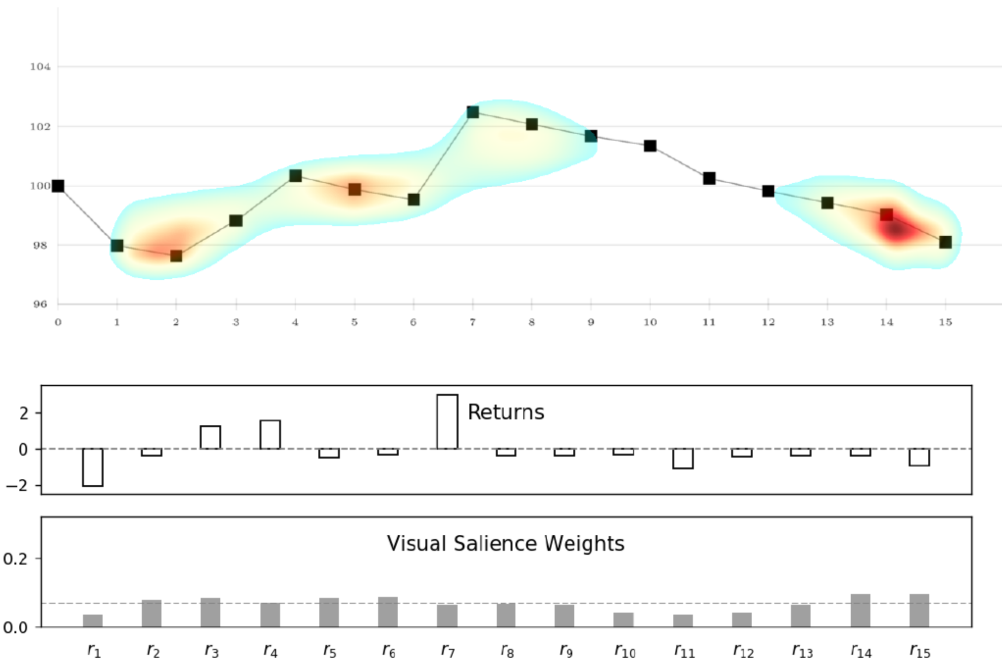

The figure shows the third price path from Figure 1 with an overlay of the heat map returned by SAM. The darker the overlay the more attention is allocated to that area of the price path. The returns in the order of occurrence and the resulting visual salience weights are depicted below.

self are discussed in the literature (Goldstein and Einhorn, 1987; Prelec, 1998; Diecidue et al., 2009). While the exact decision weights change depending on the parametrization or weighting function, the general idea remains: extreme returns are overweighted. This can be seen from the third graph in Figure 1. Moderate returns are underweighted while extreme returns are overweighted. One characteristic that is important to note is that CPT decision weights do not depend on the temporal order of returns.

In salience theory (Bordalo et al., 2012, 2013), decision weights are derived from salience weights. Salience overweights returns that have a large magnitude of difference from specified reference value, and underweights returns that are close to this reference value. This idea of salience is operationalized by assigning each return a salience rank, depending on its difference from the reference value.

$$\pi_k^{Sal} = p_k \times h_k \text{ with the salience weight } h_k = \frac{\nu^{\kappa_k}}{\sum_i \nu^{\kappa_i} p_i} \quad \nu \in (0, 1] \tag{4}$$

where κ_k denotes the salience rank of x_k, which is measured by:

$$\kappa_k = \sigma(x_k, \bar{x}_k) = \frac{|x_k - \bar{x}_k|}{|x_k| + |\bar{x}_k| + \theta} \text{ where } \bar{x}_k \text{ is a reference value}$$

These decision weights formalize how investors attend to differences rather than absolute values (Kahneman, 2003). They have the properties of *ordering* and *diminishing sensitivity*: salience increases in the distance from the reference value, and salience decreases as the reference value gets farther from zero. Often the reference value is the aggregated stock market return or a similar benchmark (Bordalo et al., 2013; Cosemans and Frehen, 2017). For a single price path, such market benchmarks are not very feasible as the investor lacks information about the market. Instead, the price path can be treated as resembling a lottery with its historical returns as potential outcomes (Barberis et al., 2016). In this setting, a possible reference value in the spirit of Bordalo et al. (2012) can be the average of all historical returns. This implies that returns that are dramatically different from this reference value are salient in the minds of investors, and are thus overweighted when aggregating historical returns into future expectations.

For the analyses, we use the parameter values suggested in Bordalo et al. (2012) of $\nu = 0.7, \theta = 0.1$.[3] An illustration is in the lower graph of Figure 1. The reference value is close to the moderate realized returns; that is why those returns receive the lowest salience weight. The reference value makes salience theory context-dependent, and thus flexible and applicable to many economic decisions. At the same time, eliciting the reference value is difficult, and

[3]The smaller the ν, the more weight is put on the extreme returns (farther from the reference value).

there exists no standard procedure. As a result, predictions can be sensitive to the specification of reference points which changes salience weights. Also, just like CPT weights, they do not depend on the temporal order of returns.[4]

For decision weights based on visual salience, we rely on eyetracking measures of subjects' eye movements, as well as general predictions of rapid visual attention. For prediction, we apply a neural network model called the Salience Attentive Model (SAM) developed in computer vision. This model uses direct measures of attention measured through eye-tracking of a large sample of average people freely gazing at thousands of images. The measured fixations are used to "train" an algorithm which creates a many-layered convolutional neural network, that takes images as an input and outputs predictions of where the fixations will occur. Overfitting is prevented by evaluating the trained model in a separate "test" set of eyetracking data.

For our application, we input images of price charts into the algorithm. SAM outputs a greyscale image assigning higher values (close to 1) to pixels that contain features that are more visually salient and lower values (close to 0) to pixels that contain features that are less visually salient. We aggregate the visual salience values for all the pixels associated with each price, and average the values across consecutive prices to obtain a decision weight for each return.

Visual salience probability weighting is defined as:

$$\pi_k^{\text{VS}} = p_k \times l_k, \ \sum_k \pi_k = 1, \ \text{with the visual salience weight } l_k = \frac{\gamma_k}{\sum_i \gamma_i p_i} \tag{5}$$

$$\text{and } \gamma_k = \frac{SAM(P_{k-1}) + SAM(P_k)}{2}$$

where $SAM(P_k)$ denotes the visual salience of price P_k as predicted by SAM.

Figures 2 - 4 illustrate the decision weights and the outcome of SAM for the three price charts used in Figure 1. Obviously, decision weights are now dependent on the order of return

[4]In Section 2, Figure S1 of the Online Appendix, we vary the parameters the CPT weighting function, and the reference value for salience, and analyze the impact on the decision weights.

realizations. The decision weights assign higher probabilities to areas of the price paths where SAM predicts higher attention. Thus, these decision weights do not depend on the extremity of the return but on the attention it receives, i.e., it is not necessarily the highest return which receives the highest weight but is dependent on the occurrence and attention allocation. While in Figure 2 the highest return gets assigned a small weight, in Figure 3 the highest return gets assigned a relatively high weight.

Different decision weights from the three theories are compared by calculating distance between them and their correlation. Three metrics are used the sum of absolute distances from empirical frequencies ($\sum_{k \in K} |\pi_k - p_k|$), the sum of absolute distances between decision weights ($\sum_{k \in K} |\pi_k^i - \pi_k^j|$), and the correlation between decision weights ($\frac{E[(\pi_k^i - \bar{\pi_k^i})(\pi_k^j - \bar{\pi_k^j})]}{\sigma_{\pi_k^i} \sigma_{\pi_k^j}}$).

Each metric is calculated for each path. Table 1 show the three distance metrics for the price paths used in the Figure 1 stylized example. The absolute distances are exactly the same for the three paths for CPT and salience theory (left table). The visual salience return weights can be different because they based on visual features, which are likely to depend on the temporal ordering of returns (e.g. whether high and low returns are adjacent in time and general sharp visual contrast or unusual line orientations). The visual salience weight absolute distances are also generally smaller in magnitude.

Correlations between Salience theory and CPT decision weights are large (r=.68) because they both assign larger decision weights to more extreme returns. Correlations from path to path range from slightly negative to slightly positive, enabling a good separation between visual salience and other theories across paths. The fact that visual salience decision weights have a wider distribution and are uncorrelated to the other two theories' decision weights is a general property of the price paths used in Study 1 (see Online Appendix Section 3).

To establish a proof-of-concept that visual salience decision weights can carry information which influences investment, let us consider an example based on the data collected in Nolte and

Table 1: Distance and correlation between decision weights from three different theories

Sum of absolute distance from empirical frequencies	Path1	Path2	Path3
CPT	0.36	0.36	0.36
Sal	1.01	1.01	1.01
VS	0.20	0.29	0.26

Sum of absolute distance between theories	Path1	Path2	Path3
CPT - Sal	0.78	0.78	0.78
CPT - VS	0.49	0.38	0.53
VS - Sal	1.14	0.93	1.09

Correlations between theories	Path1	Path2	Path3
CPT - Sal	0.68	0.68	0.68
CPT - VS	-0.47	0.48	-0.39
VS - Sal	-0.67	0.33	-0.04

The theories are Barberis et al. (2016) (CPT), Bordalo et al. (2013) (Sal), as well as the visual salience theory (VS) applied to the three sample paths in Figure 1. (Left to right) The sum of absolute distances from empirical frequencies ($\frac{1}{15}$); the sums of absolute distances between the different theories' decision weights; the degree of correlation between the different theories' decision weights.

Schneider (2018). In this study investors are primed with four price path shapes with the same returns but different temporal ordering, and then asked for their willingness to invest. In Figure S5 of the Online Appendix (Section 4), we depict the two price paths with the highest and lowest visual salience value. The two price paths are designed such that the return realizations, i.e., the expected value is the same for both. If visual salience weights indeed have forecasting ability, the price path which results in the highest (lowest) expected value based on salience decision weights (left graph, right respectively) should also show the highest (lowest) invested amount. Indeed, for the two paths the average invested amount is equal to 65.48% (left graph) and to 54.76% (right graph). This difference of more than 10 percentage points is both economically and statistically significant ($p < 0.001$).

2.2 Value Functions and Correlation Measure

To set up the full model as outlined in Equation 1, we still need to define the value function. We rely on the CPT value function which incorporates three different aspects that are psychologically motivated: reference dependence, loss aversion, and diminishing sensitivity. The functional form we use is:

$$v(x_s) = \begin{cases} x_s^{\alpha} & \text{for } x_s \geq 0 \\ -\lambda(-x_s)^{\alpha} & \text{for } x_s < 0 \end{cases} \tag{6}$$

with $\lambda > 1$ and $\alpha < 1$.

All x_s are expressed as gains or losses relative to a reference point. Thus, the DM evaluates relative returns rather than the absolute value of returns when making her decision. How reference points are formed remains an open question. For now, we encode positive returns as gains and negative returns as losses.[5]

Loss aversion is the idea that people are more sensitive to losses than equally sized gains. This is captured in the CPT value function by λ, the loss aversion parameter where $\lambda > 1$ implies that losses loom larger than gains.

Finally, the parameter α measures the diminishing marginal sensitivity that the DM has towards large losses or large gains in the value function. As the basic parameter setup, we will stick to the common choice of $\alpha = 0.88$ and $\lambda = 2.25$.

Both the decision weights and the value functions affect the evaluation of the historical price path. Since the visual salience measure only enters the functional through the decision weights, we propose an additional measure based on correlations that serves as a proxy for the direct influence of decision weights on a price path's attractiveness independent of specifics of the value function, like curvature or loss aversion.[6] Note that we can rewrite Equation 1 as:

$$V(X) \propto E(\pi_k v(x_k))$$

$$V(X) \propto E(\pi_k)E(v(x_k)) + Cov(v(x_k), \pi_k)$$

[5]For price paths, Baucells et al. (2011) and Nolte and Schneider (2018) have proposed that the reference point is some weighted combination of the initial, final, average, highest, and lowest price.

[6]When the number of decision weights is much larger than the number of possible returns, we artificially depress any covariational/correlational structure. In Study III, we explore risky assets with only two possible returns. In this case, to have a meaningful setup weights should collapse to the number of outcomes. This only holds for the model of Bordalo et al. (2013), but not for the specification of CPT proposed by Barberis et al. (2016) or for the visual salience weights. Thus, we will not use the correlational measures for the analyses in Study III.

To capture the direct influence of different decision weights on $V(X)$, we pay attention directly to the covariance term. We use the correlational measure between returns and decision weights $Corr(x_k, \pi_k)$ as a proxy for the covariance term.[7] The intuition behind this measure is as follows: if higher (lower) returns receive a higher decision weight, the correlation between returns and decision weights is high (low). As a consequence, this correlation should predict attractiveness of a price path regardless of the underlying value function. In our analyses, we investigate correlations between decision weights and returns as well as CPT value functions to determine the performance of different decision weights. Their ranking remains identical for either type of analysis.

3 Saliency Attentive Model (SAM)

Our hypothesis is that when investors look at a price chart, their gaze is drawn towards salient regions. But what exactly is salient? There is a long, impressively-accumulating tradition of computational neuroscientists using facts about actual brain processes, to emulate human salience in the form of models that predict likely human gaze locations for a wide range of natural images. The earliest models were Itti and Koch (2000). They used evidence about coarse to fine encoding in layers of visual cortex to construct an analogous computational process, inputting natural images and outputting predicted salience.

We use one such model created by the AImage Lab at the University of Modena and Reggio Emilia called the Salience Attentive Model (SAM) (Cornia et al. (2018)). SAM is a saliency prediction neural network model consisting of three main blocks: a feature extraction Convoluted Neural Network (CNN) that identifies features of input images that coincide with a high incidence of eye fixations, a Long Short Term Memory (LSTM) attention based feature encoding network that refines and balances the weights assigned to different high and low level

[7]For linear utility without loss aversion, the two measures have similar effects on $V(X)$.

feature maps when making predictions for out-of-sample images, and a prior learning network which accounts for the tendency of humans to fixate on the center region of an image.

3.1 Off-the-shelf SAM

SAM has been trained and validated on a combination of four image datasets: SALICON (Jiang et al., 2015), MIT300 (Judd et al., 2012), MIT1003 (Judd et al., 2009), and CAT2000 (Borji and Itti, 2015). Together, these datasets include more than 23,000 images that highlight complex everyday scenes containing common objects in their natural context.[8] Sources for these images include Microsoft COCO, Flickr, and LabelMe. The images in these datasets consist of natural indoor and outdoor scenes and were originally used to train machine learning algorithms for object identification. SAM is trained and validated against the eye-tracking fixations of human subjects on these images to predict visually salient regions in out-of-sample images, such as the price charts we base our experiments on.

The baseline SAM trained on images of natural indoor and outdoor scenes identifies features of the images that correlate highly with a higher number of fixations made by human observers. The identified features could be low-level (bottom-up) such as horizontal and vertical lines, or brighter colors, or shifts in contrast and shapes, but could also be high-level (top-down) like faces, trees or animals etc.[9] SAM identifies 512 features from the images in a training set, and continually refines these features using an attentional LSTM to identify an optimized set of features that are correlated with eye fixations (or lack thereof). The trained SAM then attempts to identify these refined features in out-of-sample images in the test set to predict which pixels

[8] Some categories used to separate the images in these datasets include: Action, Affective, Art, Black & White, Cartoon, Fractal, Indoor, Inverted, Jumbled, Line Drawing, Low Resolution, Noisy, Object, Outdoor Man-made, Outdoor Natural, Pattern, Random, Satellite, Sketch, and Social. The only images in the training set that bear a resemblance to price paths are a very small subset of the Sketch category. Two examples are provided in Online Appendix Section 5.

[9] It should be noted that some features identified by the CNN in SAM are a matrix of numbers that do not have any obvious interpretation in the English language, which makes the underlying algorithm somewhat "black-box" in nature.

within an image are more likely to be looked at, and which pixels are likely to be ignored. SAM makes predictions about portions of an image that will be fixated on within the first few seconds of interaction. We hypothesize that these initial seconds are enough for agents to form a heuristic about potential future returns when shown images of historical price charts. This is very similar to the ideas outlined in Barberis et al. (2016).

The SAM output after processing an input of a price chart is a corresponding greyscale image of the same resolution as the input image, with brighter patches (values closer to one) showing visually salient pixels, and darker patches (values closer to zero) showing visually non-salient pixels. We then identify the pixels within each input price chart, and add up the visual salience numbers of the corresponding pixels in the output image. These visual salience numbers are then normalized into decision weights so that they add up to one (see Section 2.1).

3.2 SAM and Price Paths

SAM is trained on images that span a variety of domains, because the algorithm is intended to capture general features of what is visually salient during perception of a huge variety of images people see every day. Our approach tests the rather bold hypothesis that decision weights derived from general visual salience might also be predictive in a specialized domain such as financial price paths (even though no stimuli resembling price paths are in the SAM training and test sets). However, because the SAM algorithm is not specially trained to predict how investors look at price paths and invest, how well it predicts is a lower bound on how well more precisely-tuned algorithms might work.

To get more evidence about whether SAM is useful for perceptions of price paths specifically (independent of top-down investment goals), we collected eyetracking data showing human subject fixations for a sample of price path images. We recorded N=57 Caltech student participants' eye-movements and fixations while they looked at price paths on a computer monitor. This is an important step. Success is far from certain. If the participants' fixations do not match up with SAM predictions, there is little hope that off-the-shelf SAM can predict investment choices.

Following the original training procedure for MIT300 images (Judd et al., 2012), participants were provided with the following instructions, "You will see a series of 60 price path images. Look closely at each price path image. After every three seconds, a new price path image will automatically appear. After viewing the price path images you will have a memory test: you will be asked to identify whether or not you have seen particular images before." This task is called "free gaze" because the subjects are not asked to make any choices which depend on what they see. The memory test afterward is simply a device to guarantee that they are attending.

Images were only shown for three seconds each to capture the parts to which initial attention is allocated. Participants were presented with a total of 60 price path images. Thirty of of these paths were selected to match those in Study III to allow for a more thorough analysis of price path images with obvious up-down and down-up trends. The other 30 paths were drawn from actual empirical price paths with distinct visual properties like the ones in Study I.

The participants were tested individually using an EyeLink 1000 Plus (SR Research, Osgoode, ON, Canada) eye-tracker. Monocular eye movements were recorded at 500Hz and fixations were identified by the eye tracker using velocity algorithms. The experiment was displayed on a widescreen Dell monitor (1280 x 1024 resolution). Participants were seated approximately 28 inches away from the monitor to allow for accurate eyetracking and comfortable gameplay. All subjects underwent a 13 point calibration and validation cycle of the eyetracker before proceeding with the experiment trials. Stimulus presentation was controlled by MATLAB using Psychtoolbox extensions (Brainard, 1997; Pelli, 1997).

We obtain a continuous fixation map for each price path image from the eye tracking data by convolving a Gaussian filter across fixation locations of all observers. A Gaussian filter is used to spread out each fixation to the pixels that are within 1 degree of visual angle (or approximately 35 pixels) to match with the area that an observer sees at high focus around the point of fixation. Figure 5 gives two examples of the groundtruth density maps from our experiment on the left, and the SAM-predicted heatmaps overlaid with the price path participants looked at on the right. By visual inspection, the figures confirm that predictions and fixations look very similar. Moreover, we use these groundtruth density maps along with the binary fixation data as inputs for different metrics to evaluate quantitatively how well SAM predicts salient points within price paths.

For an overview of suitable metrics, we refer to Le Meur and Baccino (2013) and Bylinskii et al. (2018). We use three different metrics. The first metric belongs to the group of metrics based on receiver operating characteristics (ROC). In signal detection theory, the ROC measures the trade-off between true and false positives at various discrimination thresholds. The area under the ROC curve (AUC) gives us the degree of similarity of two salience maps. A value of one indicate a perfect classification while a value of 0.5 indicates classification as by any random picture. The simplest measure we include is the pearson correlation coefficient (CC). Lastly, as a measure of dissimilarity we take KL-divergence which is a measure from information theory capturing the overall dissimilarity between two probability distributions. Table 2 presents the results. We compute metrics based on the participants' fixations maps for price path images and the salience maps predicted by SAM. As a benchmark, we compare this to SAM's performance on the set of domain-neutral images on which it was originally trained and validated. Finally, we compare participants' salience maps with heatmaps where (for each price path image) fixations were randomly assigned. While there is some drop-off in SAM's performance when applied to the domain of price path images, it is not significant. Also, SAM performs significantly better than random assignment of fixations when making predictions. Thus we conclude that SAM, although not trained on price path images, captures well what people perceive as salient within a price path.

Figure 5: Heatmaps

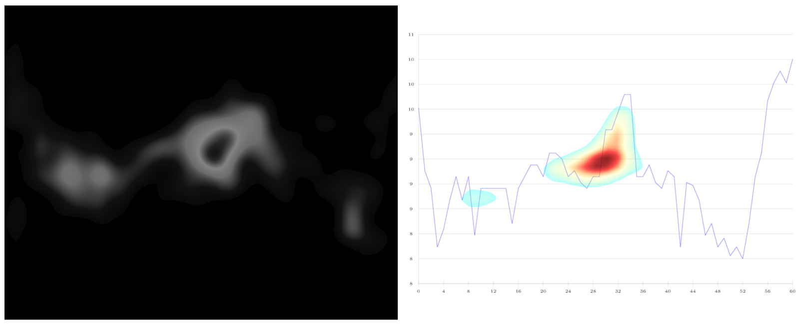

(a) Eye-Tracking Heatmap (b) SAM prediction

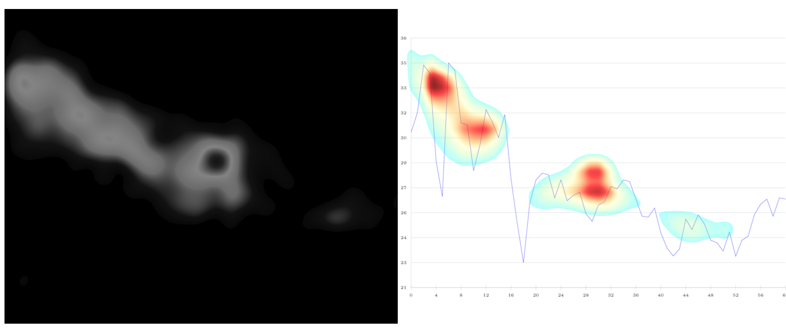

(c) Eye-Tracking Heatmap (d) SAM prediction

The figure shows two examples of the groundtruth density maps from our experiment on the left, and the SAM-predicted heatmaps where we also show the underlying price paths participants looked at on the right.

3.3 SAM and Statistical Features of Price Paths

The previous section shows that the off-the-shelf SAM has the ability to predict salient features of price paths. This justifies our ultimate goal of using the attention paid to salient features as predicted by SAM to model meaningful decision weights. However, using SAM would be redundant if the attention weights generated by SAM can be substituted by statistical features of price paths. Although we believe that SAM captures intuitive and emotionally evoked attention to salient parts of a price path while statistical features are related to active and more cognitive

Table 2: Evaluation metrics

	AUC-J.	CC	KL
SAM (domain-neutral)	0.87	0.78	1.27
SAM vs fixations (price paths)	0.81	0.52	1.58
Random vs fixations (price paths)	0.50	0.07	12.57
Range (worst-best)	0-1	0-1	∞-0

The table reports three evaluation metrics. We compare SAM's performance on the domain-neutral images it was trained on against SAM's performance against human eye-fixations for price path images. We also use random assignment of fixations as a comparison to see how well SAM is performing for price path images. We include one metric using receiver operating characteristics analysis and calculating the area under the curve: AUC-J. (Judd et al., 2009). We also include the Pearson Correlation Coefficient (CC), and Kullback-Leibler divergence (KL).

calculations, we cannot rule out that visual salience merely acts as a substitute for statistical moments of the price paths. To eliminate this alternative explanation, we explore the relation of visual salience weights and statistical features of price paths. We analyze pairwise correlations between several measures and SAM weights in a first step, and regress SAM weights on common features of the statistical measures identified by principal component analysis (PCA). To find a meaningful set of statistical features to start with, we draw on the asset pricing literature (e.g., George and Hwang, 2004; Mizrach and Weerts, 2009; Bali et al., 2011; Chen et al., 2001; Chabi-Yo et al., 2018; Raghubir and Das, 2010; McLean, 2010; Stambaugh et al., 2015; Jegadeesh and Titman, 1993; Conrad et al., 2014) . A detailed description of the methodology and the chosen measures can be found in Section 6 of our Online Appendix.

The main takeaway from performing these analyses is that a maximum of 20% of the variance in SAM decision weights can be explained by these statistical features. We conclude that while there are some features of the return distribution that are correlated with visual salience, SAM's algorithm captures properties of the price path that cannot be reproduced by merely combining additional statistical features.

4 Experiments

We collect experimental data to evaluate the performance of the visual salience model. Two studies (Study I and II) are settled on an experiment run on Amazon Mechanical Turk (MTurk). The third study (Study III) is based on a lab experiment conducted at a German university. The studies vary with respect to the degree of realism of the depicted price charts and the degree of control over the predictions of the different models.Study I explore whether the visual salience model can predict experimental investments based on price charts (from CRSP return) like those investors often see. Study II uses the same price series, but artificially rearranges the temporal order of the returns within a price series. Rearrangement changes the predictions of VS but not the other two theories (excluding recency effects). Study III considers the simplest possible depiction of price charts controlling for many which could influence investment decisions like image and task complexity. Subjects are also told the data generating process behind the price paths such that an optimal choice can be computed. Moreover, in this study we have perfect control over the environment and attention of the participants. In all three studies, the start price is fixed at 100 monetary units to keep the scaling of the y-axis constant and to limit psychological framing effects (e.g., Glaser et al., 2007; Huber and Huber, 2018).

4.1 MTurk Experiment

For Study I and II, we recruited 500 participants from MTurk which is an online platform that allows researchers to post small tasks that require a human to perform. Potential participants can browse through the lists of small tasks and can self-select into the tasks of their interest. The large available population on MTurk ensures heterogeneity in participant characteristics and a speedy data collection for an acceptable amount of money. Our task offers the participants a fixed and a variable completion fee, the latter depending on their choices within the experiment. The fixed fee was set to 2 US-$ and the variable part was on average 0.94 US-$. Participants took a mean of 12 minutes to complete the experiment, while 95% of all completion times were in the interval between 5 and 30 minutes. After completing the experiment, each participant received an individual, alpha-numerical code to verify their participation. The participant typed in the code on the MTurk Webpage and after having verified their participation, we authorized the payment.

The experiment consists of three parts: an introduction with a small tutorial to ensure that participants understand the general task, the task itself and a final questionnaire.

We show each participant ten different pre-selected price charts. Eight price charts stem from a set based on stocks from the Center for Research in Security Prices (CRSP) in 2017 presenting the input data for Study I. Two charts are based on constructed price charts simulated via a Geometric Brownian Motion (GBM), presenting the input data for Study II. For each price chart, participants decide how attractive the price chart is, what return they expect over the next 12 months, how risky they deem an investment into this stock, and finally and most importantly, how much they want to invest over the next year from their initial endowment of 1000 monetary units into this stock vs a safe bank account.[10] We emphasize that the chart can give the participant an idea about the risk and return of investing over the next year but that the return over the next year can also be much higher or lower. Participants make their investment decisions by using a slider which allows them to invest every amount between zero and 100% of their initial endowment. To avoid anchoring effects, there is no pre-selected slider position at the beginning of each period. Instead, participants need to click at any position of the slider to activate it; afterwards, they can move the slider position freely to decide on their investment. The variable payment is based on the investment task, i.e., we incentivize this task separately.[11] The participant learns that the computer will randomly select one of his 10 investment decisions

[10]For a detailed description of the questions and an example of the interface, we refer to Section 7 of the Online Appendix.

[11]This is one of the provisions we include to account for the recent MTurk data quality debate (Kennedy et al., 2018).

after having completed the final questionnaire. The participants final wealth is the value of the selected stock investment after one year plus the money in the bank account. The return of the selected stock investment after one year is based on this stock's actual return in the subsequent year, i.e., CRSP data for 2018 for Study I and simulated data for Study II. The participant received 0.001 US-$ for every monetary unit of his final wealth of the selected decision in the investment task.

For Study I, we preselect 1000 price charts based on all 8,453 companies inferred from CRSP for 2017-2018. We first remove double entries, firms with incomplete price data, firms with negative prices, penny stocks (i.e., prices below 5 US-$), and lastly incomplete stock/year combinations. In total, that leaves us with 4,246 stocks. We then form ten times ten portfolios based on two criteria. The first criteria is the stock's return within 2017. The second criteria is the degree of convexity and concavity of the price path. Paths that are classified as highly convex (concave) by our applied score tend to exhibit specific visual features, most prominently troughs (peaks) followed by rather long winning (losing) streaks. Previous experiments indicate that historical price paths with similar characteristics attract the attention of investors (Grosshans and Zeisberger, 2018; Nolte and Schneider, 2018). For a detailed description of the applied score, we refer to Online Appendix Section 8. First each stock is allocated to one decile based on its return in 2017. Second, within each decile, every stock is assigned to a convexity score decile. We then randomly pick 10 stocks from each bucket. We ensure that each participant chooses on eight different stock charts where the order is randomized. In total, every price chart is evaluated four times ensuring that each participant evaluates each chart no more than once.

For Study II, we select 300 constructed price charts. We start constructing our price paths by simulating 252 daily returns based on a GBM. We implement 3 different parameter constellations for the GBM where the base case implies a mean of 6% and a volatility of 20%, then we decrease the mean in one specification to -6%, and fix realized skewness of the drawn path to zero in the third specification.[12] More precisely, we draw one possible path realization for each parameter constellation from the underlying lognormal distribution. The drawn path is the base for constructing 100 different price paths shapes: We split the price path into twelve subperiods of equal length mirroring twelve months and rearrange the subperiods randomly. Thus, all shapes of price paths are identical with respect to daily returns and returns of the twelve subperiods. Moreover, they are also identical for any sub-periods that are multiples of twelve. This implies that individuals having CPT-preferences as described in the basic model of Barberis et al. (2016), or Salient-preferences as in Bordalo et al. (2013) or Cosemans and Frehen (2017) would assign the same value to every price path.

Besides demographic questions, we include questions on risk taking, and general traits of decision making. Moreover, we control for attention and understanding by including a set of quiz questions in the aftermath of the experiment. For details on the questionnaire, we refer to the Online Appendix Section 7.

As Online Appendix Section 12 shows, the recruited MTurk sample matches the US population exhibiting a somewhat overrepresentation of high-education groups and younger individuals. This is consistent with previous literature documenting that MTurkers are actually quite representative of the population of U.S. internet users (Ross et al., 2010; DellaVigna and Pope, 2017).

4.1.1 Results: Study I

Study I addresses how the visual salience model performs compared to the CPT-framework and the salience-framework in predicting invested amounts of our participants relying on empirical price charts. Table 3 presents the regression results of valuations on invested amounts for all

[12] In the experiment, we include two additional parameter specifications with an extreme high volatility and an extreme low volatility. To leave one parameter fixed and increase the homogeneity in the price paths, we exclude these in this analysis.

Table 3: Regressions for IA, Study I: Correlation Measure

	(1) IA [%]	(2) IA [%]	(3) IA [%]	(4) IA [%]	(5) IA[%]
Corr(Return, VS weights)	0.670***			0.635***	0.706***
	(0.235)			(0.237)	(0.249)
Corr(Return, CPT weights)		0.289		0.984**	1.147**
		(0.242)		(0.434)	(0.534)
Corr(Return, Salience weights)			-0.0271	-0.249**	-0.310**
			(0.0689)	(0.121)	(0.146)
Return	0.297***	0.294***	0.295***	0.291***	0.367***
	(0.0206)	(0.0203)	(0.0204)	(0.0206)	(0.0160)
Stdv	-0.130***	-0.137***	-0.131***	-0.140***	-0.184***
	(0.0152)	(0.0158)	(0.0152)	(0.0164)	(0.0159)
Skewness	0.00859	-0.0183	0.0159	-0.0303	-0.0340
	(0.0112)	(0.0257)	(0.0193)	(0.0266)	(0.0349)
Constant	0.00617***	-0.0144	-0.00142	-0.0562**	-0.117**
	(0.00216)	(0.0121)	(0.00362)	(0.0275)	(0.0488)
Observations	4000	4000	4000	4000	4000
R^2	0.162	0.160	0.160	0.163	

* p $<$ 0.1, ** p $<$0.05, *** p$<$0.01

The table reports invested amounts on the correlation between current return, visual salience decision weight, CPT decision weight, and the salience decision weight. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account for accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. All variables are standardized for easy comparison of coefficient sizes and to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

1000 CRSP-based price charts included in the MTurk study. We approximate the value for all three models by the correlation of the historical return and the different decision weights as outlined in Section 2. This isolates the effect of different value functions for the different models from the implications of the decision weights. Model (1) of Table 3 only includes the correlation measure for the visual salience decision weight. As controls we include in each regression the return, standard deviation and skewness of the different price charts to account for heterogeneity in stock characteristics. The results confirm that there is a significant and positive relation between our measure and the invested amount. For CPT weights, this relation is insignificant and positive (Model (2)) while for salience weights the relation is insignificant and negative (Model (3)). Model (4) presents the results of a regression where we include all three measures. Note that the correlation between CPT weights and salience weights is above 90%, thus we face multicollinearity. Despite this, the coefficient of the visual salience measure remains positive and significant. Model (5) is a Tobit regression to account for some clustering of invested amounts at 0 which has almost no impact on the visual salience coefficient. In our MTurk sample, we see that in 13% of all decisions, nothing was invested and 0.2% of the participants never invested. In our Online Appendix Section 11, we run further regression analyses to account for MTurk specific issues. One regression excludes participants who fail attention and comprehension filters, and another regression uses the attractiveness question as a control to check the consistency of the answers. The results stay qualitatively the same. It is also noteworthy that in line with economic intuition, the return has a positive and significant coefficient while the standard deviation has a negative and significant coefficient for all models.

As additional analyses, we present in Table 4 the same five regressions as in Table 3 but for a CPT value function. Table 4 shows that with a CPT value function, the visual salience model is the only model which positively and significantly predicts invested amounts. For CPT and salience decision weights, all coefficients remain insignificant changing between positive and

Table 4: Regressions for IA, Study I: CPT Value Function

	(1) IA [%]	(2) IA [%]	(3) IA [%]	(4) IA [%]	(5) IA [%]
Visual Salience	0.0955***			0.106**	0.109**
	(0.0357)			(0.0415)	(0.0474)
CPT		0.0184		-0.0489	-0.0588
		(0.0590)		(0.0809)	(0.0955)
Salience			-0.0117	0.00768	0.00821
			(0.0242)	(0.0285)	(0.0346)
Return	0.252***	0.289***	0.297***	0.264***	0.342***
	(0.0255)	(0.0296)	(0.0206)	(0.0315)	(0.0327)
Stdv	-0.0453	-0.115**	-0.136***	-0.0776	-0.126*
	(0.0364)	(0.0567)	(0.0180)	(0.0599)	(0.0723)
Skewness	0.0139	0.00406	0.0187	0.0235	0.0241
	(0.0111)	(0.0211)	(0.0209)	(0.0247)	(0.0280)
Constant	-4.60e-09***	-4.68e-09***	-4.66e-09***	-4.55e-09***	-0.0501
	(2.31e-11)	(5.55e-11)	(1.96e-11)	(9.01e-11)	(0.0356)
Observations	4000	4000	4000	4000	4000
R^2	0.162	0.160	0.160	0.162	

* p < 0.1, ** p <0.05, *** p<0.01

The table reports invested amounts on the Visual Salience model, the CPT model (Barberis et al., 2016) and the Salience model (Bordalo et al., 2012) where we assume gain-loss utility with curvature $\alpha = 0.88$, a reference point ρ of zero, and a loss aversion coefficient of $\lambda = 2.25$. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account for accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. All variables are standardized for easy comparison of coefficient sizes and to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

negative sign across regressions.

To sum up, regardless of the specification of the value function, the VS model predicts invested amounts significantly and positively. This is tentative experimental evidence that when investors look at price charts, their gaze is drawn towards salient regions which impact investment. Visual attention should therefore be studied further as an aspect of models of invetment decision making when expectations are formed taking the past realizations as a proxy.

4.1.2 Results: Study II

In Study II, we aim at disentangling the effect of visual salience from pure temporal ordering like a recency effect, i.e., recent historical returns are more likely to be representative of returns in the immediate future than returns from a long time ago (e.g., Bordalo et al., 2019). Recency effects capture a fundamental psychological factor that approximates the speed of memory decay, and has some empirical support in extrapolation based models (e.g., Cassella and Gulen, 2018). While CPT-values and salience-values to not vary across price paths in Study II, if those weights are not temporally influence, Barberis et al. (2016) suggest a modification of the original CPT to accomodate recency effects. They introduce a recency parameter ρ which underweights (geometrically) returns that are further in the past. This kind of recency effect was first established by Ebbinghaus (1880) in learning word lists, and is seen everywhere in psychological and neuroscientific models of learning (Jones and Sieck, 2003; Baddeley and Hitch, 1993).

More specifically, t_k is the number of months ago that return r_k was realized, the multiplicative term in front of the value function is ρ^{t_k} normalized by the weighted sum of all ρ. In the robustness test of Barberis et al. (2016), ρ equals $0.9, 0.85, 0.8$. Barberis et al. (2016) also argue that investor's quick, passive reaction to a chart is likely to be based on the chart as an integral whole, with the early part of the chart affecting the investor just as much as the

Table 5: Regressions for IA, Study II: Recency Effect

	(1) IA[%]	(2) IA [%]	(3) IA [%]	(4) IA [%]	(5) IA [%]
Visual Salience	0.205*				0.217**
	(0.107)				(0.107)
CPT (0.95)		0.00417			
		(0.0418)			
CPT (0.85)			0.0362		
			(0.0393)		
CPT (0.50)				0.0961*	0.104*
				(0.0583)	(0.0585)
6, 20	0.0139	0.0556	0.0496	0.0706	0.0252
	(0.113)	(0.116)	(0.111)	(0.105)	(0.109)
6, 20, skew	0.0985	0.130	0.113	0.128	0.0909
	(0.103)	(0.116)	(0.112)	(0.0997)	(0.102)
Constant	-0.0710	-0.0627	-0.0521	-0.0473	-0.0530
	(0.0618)	(0.0713)	(0.0701)	(0.0665)	(0.0648)
Observations	600	600	600	600	600
R^2	0.030	0.011	0.015	0.024	0.045

Standard errors in parentheses

* p<0.1, ** p<0.05, *** p<0.01

The table reports invested amounts on the Visual Salience Model with CPT value function (Model(1)) and a CPT model with recency parameter $\rho = 0.95, 0.85, 0.50$, Models (2)-(4). In Model (5) we run the regression with the Visual Salience value and the CPT value ($\rho = 0.5$) as explanatory variables. We also include controls for the parameter constellation of the constructed price charts, with $\mu = -6\%$ and $\sigma = 20\%$ as reference category.

later part. The visual salience decision weights basically capture the weight investors assign to returns which are visually salient, while ρ captures a recency effect in the weighting of the returns. To see whether a modified CPT-value has predictive power for our constructed price charts, we run additional regressions. Table 5 presents the results where in Models (2)-(4) we include $\rho = 0.95, 0.85, 0.5$. Model (1) includes the visual salience weights with a CPT-value function. Interestingly, only for a very strong recency effect does the CPT-value become significant. Including both CPT- and the visual salience-value, shows that neither model is redundant. To determine the exact parameter value ρ for which the CPT value becomes significant, we run regressions varying ρ from 0.95 to 0.05. Figure 6 depicts the results by illustrating the coefficient and confidence intervals for regressions of CPT with varying recency parameter on invested amounts. For very low parameter values, the coefficient becomes insignificant again. Moreover, the visual salience model remains significant also when including the CPT model with recency weighting. This gives rise to the question whether in fact investors exhibit a recency and a visual attention/salience bias. Empirical research has shown that the 52-week high has predictive power for subsequent returns (George and Hwang, 2004). Translated to the visual salience model, this could mirror a salience weighted event, while the general tendency to overweight most recent events is an independent effect. A further investigation of this relation is left for future research. But our results clearly demonstrate that to fully understand investment choices, a combination of different approaches could be fruitful.

4.2 Laboratory Study

4.2.1 Experimental Setup

In Study III, we run a laboratory experiment to analyze how well the visual salience model fares in a controlled environment with highly simplified investment options. The goals of this study are threefold: First, we reduce task complexity to a minimum by making the available

Figure 6: CPT with Recency weighting

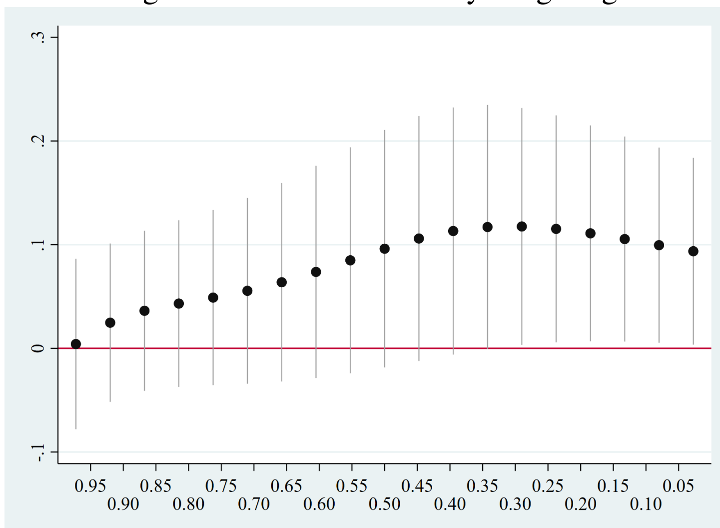

The figure shows the coefficient and confidence intervals for regressions of CPT with varying recency parameter on invested amounts.

asset a binary lottery which yields a positive or negative return with known outcomes and probabilities. This implies that participants face no uncertainty regarding the underlying processes that determine future asset prices and returns. Moreover, the investment task and information provided are comparably easy to process. Consequently, any biases or distortions of the investment decision caused by the complexity of the task are also minimized. Second, we make the visualization process easy to understand. Participants receive graphical information about the lottery outcomes in form of a price chart. The price chart is a depiction of previous periods' realizations of the binary lottery, and participants see how the price path updates after each realization. This makes the fact that a price path is merely a visualization of lottery outcomes (asset returns) apparent to participants. Given that the participants attended advanced statistic classes during their studies, they should be able to fully understand the mechanisms behind how a price path develops over time. Third, we run the experiment with a homogeneous student sample in a controlled setting. Compared to Studies I and II, this adds the necessary control regarding available information, distractions, focus on the task and all other benefits that a lab environment provides. The experiment for Study III is fully computer-based and consists of three parts: comprehensive instructions, the main experimental task, and a concluding questionnaire. The course of the experiment is visualized in Section 9 of the Online Appendix.

In the main part of the experiment, participants have to repeatedly decide how much of their endowment to invest in a risky asset. The part of the endowment that is not invested earns no return. As described above, the risky asset is resembled by a binary lottery that has one positive and one negative outcome.[13] Expected return, standard deviation, and skewness of outcomes is constant for a given asset, but varies between assets to reduce repetitiveness. Each participant sees a total number of 18 different assets in randomized order. For 15 of the 18 assets the participants first observe 15 periods in which a price path develops based on the outcomes of the lottery.[14] In these first 15 periods, they cannot make any investment decisions. Over the subsequent 15 periods (periods 16-30), they make their investment decisions. Immediately after they confirmed their decision, the outcome of this period is realized and the price path continues to develop. All return realizations in the individual periods are independent from one another, and all return properties are clearly communicated to participants.

The price charts for the first 15 periods are preconstructed for 2/3 of the 15 assets. In these preconstructed scenarios, distributions of period returns are identical but only differ with respect to the order in which returns occur. In preconstructed convex (concave) price charts, negative (positive) returns occur earlier than positive (negative) returns. We included these preconstructed charts to analyze how the price path shape affects the viability of the visual salience measure.[15] Random (non-preconstructed) charts are random realizations of the binary

[13]For a detailed specification of the set of binary lotteries, we refer to the Instructions of the experiment in Section 10 of the Online Appendix.

[14]For the three remaining assets, no price path was shown. The analyses for this situation can be obtained from the authors upon request, but does not impact our results but serves as a benchmark w.r.t. the replication of previous results of Nolte and Schneider (2018).

[15]There were no qualitative differences in performance between scenarios, so we do not pursue this analysis further in this paper. Additional data is available from the authors upon request.

lottery for the first 15 periods, which can result in equal, higher, or lower overall return than in the convex and concave scenarios. For periods 16-30, period returns for all scenarios are truly random outcomes determined by the given parameters of the asset.

Participants use a slider to communicate their investment decisions. To avoid anchoring effects, there is no pre-selected slider position at the beginning of each period. Instead, participants need to click at any position of the slider to activate it; afterwards, they can move the slider position freely to decide on their investment. The current endowment and possible changes in wealth resulting from the chosen investment are depicted on screen to facilitate accessibility of the decision. After the final decision is made and confirmed, a pop-up window informs participants about the realized return and the wealth consequences.

Similar to the design in Study I and II, participants are paid according to one of the 18 decision rounds. This round is determined randomly after the experiment. For this, each participant draws a numbered ping-pong ball out of a concealed urn with 18 balls. This aims at ensuring trust among participants about how their remuneration is determined. Participants then receive a variable payment of 0.001€ for each monetary unit of their final wealth in this round, plus a fixed payment of 8.00€. The experiment concludes with a questionnaire on demographics and self-perceived behavior, the Cognitive Reflection Test (Frederick, 2005), a financial literacy test (Lusardi and Mitchell, 2011), and a risk aversion elicitation task (Holt and Laury, 2002).

We recruited 275 undergraduate and graduate business students at an university in Germany. The experiment was conducted under controlled conditions in the computer labs of the university in 18 separate sessions. On average, it took participants 61 minutes to complete the entire experiment (including instructions and questionnaire), and average remuneration was 16.60€ (about 18.68 USD-$ at the time of the experiment). For comparison, a student research assistant earned 10.00€ per hour at that time.

4.2.2 Results: Study III

To analyze our data, we use all 15 investment decisions of each decision round as our dependent variable. This allows us to incorporate the effect of an unfolding price path in our analyses. Since our asset only has binary outcomes, analyzing the correlation between weights and returns as for the previous studies is not a viable option. Instead, we focus on the CPT value function for visual salience, salience, and CPT weights.

We include expected return, standard deviation, skewness, and actual realizations as control variables. The first three are defined by the characteristics of the risky asset and determine the (normative) attractiveness of the asset. Since we only vary these parameters marginally between decision rounds to prevent repetitiveness, we do not expect to see large effects on any of these measures. Actual realizations are the result of a random process and determine the shape of the respective price path in each round and period. They are operationalized as the average realized returns per period.[16] While expected return, standard deviation, and skewness exclusively account for participants' beliefs, realized returns might influence beliefs and preferences, or both. Participants might believe in autocorrelation (positive realized returns indicate higher future returns) or mean-reversion (positive realized returns indicate lower future returns), ignoring the actual probabilities and outcomes provided. Participants might also act more risk averse (seeking) after positive (negative) realized returns if that pushes the asset in the gain (loss) domain. So while there is no straight-forward expectation we have for the coefficient of this variable, it is important to control for these effects as they would otherwise be captured in the visual salience, salience, or CPT score.

Table 6 reports regression results for each of the three measures individually (Model (1)-(3)) and jointly (Model (4)). We see the same pattern we have observed in Studies I and II

[16]Note that this is synonymous with including the total realized return, realized standard deviation, or total number of periods with a price in-/decrease, since probabilities and outcomes of the binary lottery are constant within each round.

with respect to the visual salience measure. Its coefficient is highly significant and has positive explanatory power for investments. In contrast to the first two studies, we see a similar effect both in strength and significance for CPT, while salience remains insignificant.

One drawback of the simplified setup we test in Study III is that we cannot analyze correlations, and therefore our results may be dependent on the exact specification of the value function. We account for this unavoidable deficit with two additional analyses: first, we include the value of a CPT value function with equal decision weights for each realized return as a control variable. That way, we create a baseline effect that can be attributed to the particular specification of the value function itself. Any explanatory power that the three models retain in this setup has to stem from the decision weights. The effect of including this control variable is depicted in Model (5) of Table 6. The visual salience score retains its significance, and the coefficient remains almost unchanged. CPT on the other hand loses significance as well as strength of relation, indicating that this measure is much more dependent on the exact specification of the value function than visual salience. In Model (6), we also run a Tobit regression to accommodate for accumulation of participants who did not invest in that stock. The results remain qualitatively the same.

For our second analysis to investigate the impact of value function specification, we tested a large number of α-λ combinations. Namely, we varied α between 0.7 and 1 in increments of 0.01, and λ between 0.8 and 3 in increments of 0.05. We then ran Model (4) of Table 6 for each possible alpha-lambda combination within this range. In total, this procedure resulted in 1,350 regression analyses. Figure 7 reports histograms of resulting coefficients for each of the measures on the left hand side, and histograms of p-values on the right hand side. Results confirm that the effect of visual salience is relatively independent of the chosen value function, while the effect of CPT varies strongly with the parameters of the value function. For salience, the choice of value function also doesn't affect results, although coefficients remain small and

Table 6: Regressions for IA, Study III: CPT Value Function

	(1)	(2)	(3)	(4)	(5)	(6)
	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]
Visual Salience	0.0358***			0.0257***	0.0287***	0.0469***
	(0.00814)			(0.00867)	(0.00838)	(0.0119)
CPT		0.0386***		0.0245*	0.0183	0.0163
		(0.0128)		(0.0140)	(0.0153)	(0.0140)
Salience			0.00804	0.00737	0.00524	0.00719
			(0.00766)	(0.00763)	(0.00801)	(0.00794)
Expected return	0.0834	0.0685	0.0778	0.0730	0.0775	0.113*
	(0.0794)	(0.0790)	(0.0796)	(0.0793)	(0.0798)	(0.0601)
Expected Stdv.	-0.00913	0.000798	-0.0242*	0.00271	-0.0261	-0.0409
	(0.0133)	(0.0154)	(0.0129)	(0.0153)	(0.0394)	(0.0270)
Expected skewness	-0.251***	-0.287***	-0.256***	-0.287***	-0.275***	-0.438***
	(0.0410)	(0.0428)	(0.0443)	(0.0457)	(0.0495)	(0.0338)
Average realized return	-0.339***	-0.346***	-0.234***	-0.363***	-0.224	-0.355***
	(0.0427)	(0.0520)	(0.0389)	(0.0547)	(0.167)	(0.120)
CPT equal weights					-0.0488	-0.0665*
					(0.0589)	(0.0403)
Constant	0.115**	0.0940*	0.160***	0.0871	0.172	0.469***
	(0.0478)	(0.0532)	(0.0468)	(0.0530)	(0.119)	(0.114)
Observations	61875	61875	61875	61875	61875	61875
N	275	275	275	275	275	275
R^2	0.024	0.024	0.024	0.024	0.024	

* p < 0.1, ** p <0.05, *** p<0.01

The table reports regressions of invested amounts on the Visual Salience model, the CPT model (Barberis et al., 2016) and the Salience model (Bordalo et al., 2012) respectively. Control variables are expected return of the binary lottery, its standard deviation, its skewness, and the average realized return. In Model (5) we also include a CPT value function with equal (1/n) decision weights to control for the effect of value function specification. In Model (6) we use a Tobit regression to account for accumulation of participants who did not invest in that stock. We use fixed effects to control for individual characteristics of participants. All variables are standardized for easy comparison of coefficient sizes and to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

insignificant.[17] Although there are some parameter constellations where CPT outperforms visual salience in terms of effect strength, we can conclude from Study III that attention driven by visual salience has high explanatory power for investor decision making, even in the simplest of choices and in a highly controlled environment.[18]

Figure 7: Parameter sensitivity analysis

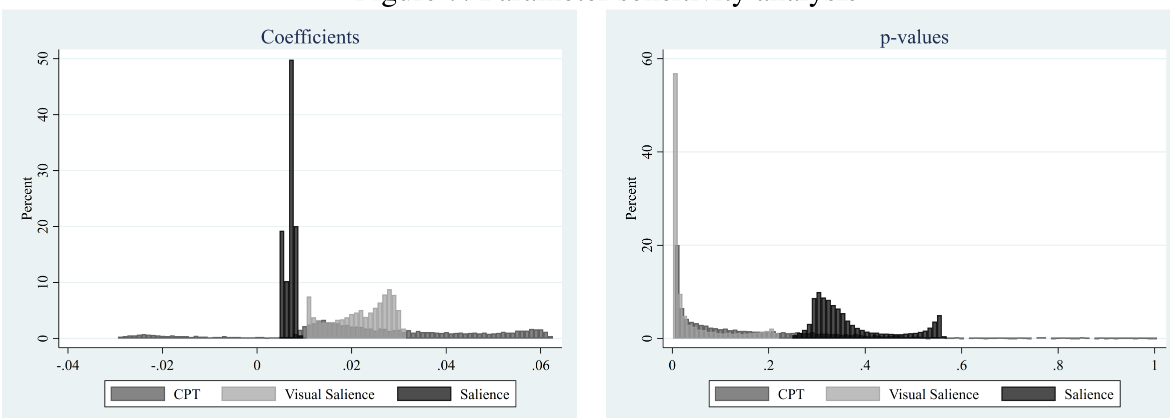

The histogram on the left (right) shows the three models' distribution of coefficients (p-values) for all analyzed curvature (α) and loss aversion (λ) combinations. α is varied in increments of 0.01 between 0.7 and 1, λ in increments of 0.05 between 0.8 and 3. In total, the graphs include coefficients and p-values from 1,395 regressions.

5 Conclusion

Human perception is highly adapted to notice certain kinds of visual features. Those attention-grabbing features are called "salient". Attention toward salient features is reliably predicted by their properties, which are often context-dependent (such as contrast, novelty, familiarity, Bayesian surprise...). We use a particular algorithm to predict visual salience of prices in a sequential price path. In the visual salience model, investors assign decision weights to past returns based on how visually salient the respective area of the corresponding price path is.

[17]Interestingly, we find strongest results for CPT when λ is between 1.2 and 2. This is consistent with results from recent meta studies on loss aversion (Walasek et al., 2018).

[18]In Section 12 of the Online Appendix, we also show the parameter sensitivity analysis for the cumulative distribution function of coefficients and p-values.

Put differently, investors tend to pay more attention to and consequently overweigh returns that "stick out" in some way. Li and Yu (2012) show that investors indeed pay more attention to returns that are near natural anchors, for example the height of a peak (George and Hwang, 2004), the bottom of a trough (Huddart et al., 2009), or during a particularly long streak (Raghubir and Das, 2010).

We repurpose machine learning techniques, adapted to predict what people look at in a wide variety of natural images, to predict salient points in price paths. These predictions may work or may fail. To the extent that human perception of visual stimuli is universal and consistent across many domains, an all-purpose algorithm may also predict the salient aspects of financial stock price paths. However, the salience models are tuned to predict the first few seconds of "bottom-up" (goal-free) perception. This does not mean the same models will predict salience over longer periods of deliberation, of the influence of that kind of salience on longer decisions.

We can confirm the adequacy of this algorithm for price paths by providing evidence that its predictions match actual eye-tracking data closely. We then see in three experimental studies that the visual salience model can actually forecast investment decisions to some extent. This holds for real stock price charts, for artificially constructed charts that only vary in the order of the returns but not the returns themselves, and for the easiest possible depiction of artificial price charts in repeated two-outcome lotteries. Compared to other economic models that focus on different aspects of attention and salience like perception of probabilities and return differences, the visual salience model has some higher explanatory power and often makes more accurate predictions about investment decisions by the human subjects.

An interesting puzzle (which we hope to solve in upcoming revisions to this paper) is why there is such a large difference in predictive strength between the CPT and Salience theories. Salience theories do not predict actual subjects' investments well. This is odd because there is a high correlation between the two sets of weights (an average of .80, across 1,000 Study I price

Figure 8: CPT and Salience Rank Differences

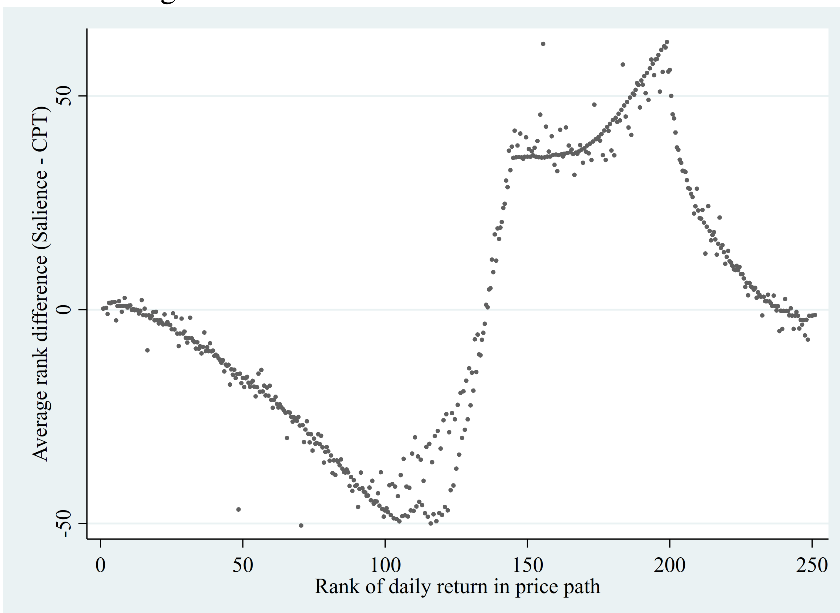

The figure shows underlying differences between the rankings of the CPT and the Salience weights. We sort all 250 returns from highest to lowest along the x-axis. The y-axis plots the average differences (across all 1,000 CRSP images) in the ranking assigned to the weights of these returns by the CPT and Salience theories.

paths). Both tend to overweight extreme returns. What is going on?

To dig further into this surprise, first remember that returns are ranked so that a low number (e.g. #1 rank) is a positive return. A high number is a poor return. Figure 8 shows the (averaged) differences in rankings (the CPT rank minus the Salience theory rank), for the price paths from Study I. The x-axis is the rank of each return of the 250 returns in a particular sequence (one is highest and 250 is lowest). You can see that extreme and middle returns are ranked similarly in both theories (that is, the CPT-Salience rank difference is close to zero). But above-average returns (e.g. rank 50-100) and below-average returns (150-200) are ranked quite differently by the two weighting theories. Look at returns that are numerically ranked around 100 (i.e. these are a little higher than he median out of 250 returns). On average, these returns have a CPT rank which is much higher than their Salience rank (around 50 ranks higher). In the opposite direction, lower returns that are higher-ranked numerically (150-200) get lower CPT ranks than Salience ranks. The CPT weights are pushing the above-average returns up and the below-average returns down, compared to Salience weighting. Put in the opposite way, Salience ranking tends to compress return ranks toward the middle (compared to CPT). It's as if the Salience weights are good at overweighting very high and low returns, but do not differentiate sufficiently among returns that are a little above or below the median. We will continue to explore whether this tendency is a diagnostic for why salience weights are not predicting decisions well.

Salience theory has two free parameters, ν, θ which have psychophysical meaning. We have not carefully explored whether regions of this two-dimensional parameter space make much more sensible and accurate predictions about the experimental data. It may be that there are more sensible configurations of such parameters which produce a much more plausible model that fits these data better, while also honoring the intuitions and structure underlying Salience theory.

Another plausible source of surprising inaccuracy of the Salience theory is our particular selection of reference level. Bordalo et al. (2013) and related papers like Cosemans and Frehen (2017) suggest that a sensible reference level for returns is the risk-free rate or a market return. Bordalo et al. (2013) are particularly interested in how simple versions of their framework could generate effects like preference for skewness and apparent time-varying risk premia, from a more basic common source. For that purpose, a simple choice of reference point such as market return is a good approximation. In understanding path-specific investments as in our experimental price charts (where the average market return is not visible) there are likely to be better specifications for the reference point. We chose a path-specific reference return (the average) because participants have no information about any market or risk-free rate, and are only evaluating the attractiveness of one asset at a time. We are exploring the effects of using other plausible reference levels such as the path-specific median return, or zero, to check what effect this has on the performance of Salience theory in our experiments. The neuroscientific idea that salience is strongly affected by contrast, and adaptive coding to previous experience, may be useful in sifting through sensible specifications.

We are cognizant of the fact that investors likely look at a number of additional sources for information before making choices in the real world. Two of our experiments present a stylized setup where the price path is the *only* source they are exposed to, because we wanted to isolate the effect of visual salience. It is likely that when considering these other sources, probability weighting and return differences like in the models of Barberis et al. (2016) and Bordalo et al. (2012, 2013) play a larger role in decision-making. As mentioned earlier, we believe that integrating visual salience and recency effects into one model highlighting probability weighting or return differences could deepen our understanding of the underlying mechanisms at play.

An interesting question is whether the highly salient prices and returns in the price charts tend to correspond to statistical features. For example, do short periods with high or highly variable prices grab attention? Do sharp ”jumps” in prices, or periods of up-and-down negative return autocorrelation, tend to grab attention? Section 6 of the Online Appendix discusses some preliminary results of this type. Statistical features, reduced in dimension by principal component analysis, seem to explain about 20% of salience. We will study these relations more carefully in future work.

Our data suggest that economic models seem to benefit from incorporating visual salience features of decision problems, above and beyond existing concepts of salience based on numerical contrast and probability weighting. Since visual salience is grounded in basic human visual perception, it is directly measurable via eye-tracking and predictable via algorithms. This allows analyzing the impact of visual salience in other economic domains where visual features of the provided information and alternatives also may play a role. The attention that aspects of the decision problem receive based on their visual salience potentially affects several layers of the decision making process, for example reference points, outcomes, or even lead to ignoring or considering entire alternatives. Future research will shed further light on these additional applications of visual salience.

References

Abdellaoui, M. (2000), Parameter-free elicitation of utility and probability weighting functions, *Management Science* **46**(11), 1497–1512.

Baddeley, A. D. and Hitch, G. (1993), The recency effect: Implicit learning with explicit retrieval?, *Memory & Cognition* **21**(2), 146–155.

Bali, T. G., Cakici, N. and Whitelaw, R. F. (2011), Maxing out: Stocks as lotteries and the cross-section of expected returns, *Journal of Financial Economics* **99**(2), 427–446.

Barber, B. M. and Odean, T. (2007), All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *The Review of Financial Studies* **21**(2), 785–818.

Barberis, N., Mukherjee, A. and Wang, B. (2016), Prospect theory and stock returns: An empirical test, *The Review of Financial Studies* **29**(11), 3068–3107.

Barberis, N., Shleifer, A. and Vishny, R. (1998), A model of investor sentiment, *Journal of Financial Economics* **49**(3), 307–343.

Baucells, M., Weber, M. and Welfens, F. (2011), Reference-point formation and updating, *Management Science* **57**(3), 506–519.

Benartzi, S. and Thaler, R. H. (1995), Myopic loss aversion and the equity premium puzzle, *The Quarterly Journal of Economics* **110**(1), 73–92.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2012), Salience theory of choice under risk, *Quarterly Journal of Economics* **127**(3), 1243–1285.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2013), Salience and asset prices, *American Economic Review* **103**(3), 623–28.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2019), Memory, attention, and choice, *Technical report*, National Bureau of Economic Research.

Borji, A. and Itti, L. (2015), Cat2000: A large scale fixation dataset for boosting saliency research, *arXiv preprint arXiv:1505.03581*.

Brainard, D. H. (1997), The Psychophysics Toolbox, *Spatial Vision* **10**, 433–436.

Bylinskii, Z., Judd, T., Oliva, A., Torralba, A. and Durand, F. (2018), What do different evaluation metrics tell us about saliency models?, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(3), 740–757.

Camerer, C. F., Johnson, E., Rymon, T. and Sen, S. (1993), Cognition and framing in sequential bargaining for gains and losses, *Frontiers of Game Theory* **104**, 27–47.

Cassella, S. and Gulen, H. (2018), Extrapolation bias and the predictability of stock returns by price-scaled variables, *The Review of Financial Studies* **31**(11), 4345–4397.

Chabi-Yo, F., Ruenzi, S. and Weigert, F. (2018), Crash sensitivity and the cross section of expected stock returns, *Journal of Financial and Quantitative Analysis* **53**(3), 1059–1100.

Chen, J., Hong, H. and Stein, J. C. (2001), Forecasting crashes: Trading volume, past returns, and conditional skewness in stock prices, *Journal of Financial Economics* **61**(3), 345–381.

Conrad, J., Kapadia, N. and Xing, Y. (2014), Death and jackpot: Why do individual investors hold overpriced stocks?, *Journal of Financial Economics* **113**(3), 455–475.

Cornia, M., Baraldi, L., Serra, G. and Cucchiara, R. (2018), Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model, *IEEE Transactions on Image Processing* **27**(10), 5142–5154.

Cosemans, M. and Frehen, R. (2017), Salience theory and stock prices: Empirical evidence, *Working Paper*.

DellaVigna, S. and Pope, D. (2017), What motivates effort? Evidence and expert forecasts, *The Review of Economic Studies* **85**(2), 1029–1069.

Diecidue, E., Schmidt, U. and Zank, H. (2009), Parametric weighting functions, *Journal of Economic Theory* **144**(3), 1102–1118.

Ebbinghaus, H. (1880), *Urmanuskript" Ueber das Gedächtniss" 1880*, Passavia Univ.-Verlag.

Frederick, S. (2005), Cognitive reflection and decision making, *Journal of Economic Perspectives* **19**(4), 25–42.

George, T. J. and Hwang, C.-Y. (2004), The 52-week high and momentum investing, *The Journal of Finance* **59**(5), 2145–2176.

Glaser, M., Langer, T., Reynders, J. and Weber, M. (2007), Framing effects in stock market forecasts: The difference between asking for prices and asking for returns, *Review of Finance* **11**(2), 325–357.

Goldstein, W. M. and Einhorn, H. J. (1987), Expression theory and the preference reversal phenomena., *Psychological Review* **94**(2), 236.

Grosshans, D. and Zeisberger, S. (2018), All's well that ends well? On the importance of how returns are achieved, *Journal of Banking and Finance* **87**, 397–410.

Hillert, A., Jacobs, H. and Müller, S. (2014), Media makes momentum, *The Review of Financial Studies* **27**(12), 3467–3501.

Holt, C. A. and Laury, S. K. (2002), Risk aversion and incentive effects, *American Economic Review* **92**(5), 1644–1655.

Huber, C. and Huber, J. (2018), Scale matters: Risk perception, return expectations, and investment propensity under different scalings, *Experimental Economics* pp. 1–25.

Huddart, S., Lang, M. and Yetman, M. H. (2009), Volume and price patterns around a stock's 52-week highs and lows: Theory and evidence, *Management Science* **55**(1), 16–31.

Jegadeesh, N. and Titman, S. (1993), Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* **48**(1), 65–91.

Jiang, M., Huang, S., Duan, J. and Zhao, Q. (2015), SALICON: Saliency in Context, *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Jones, M. and Sieck, W. R. (2003), Learning myopia: An adaptive recency effect in category learning., *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**(4), 626.

Judd, T., Durand, F. and Torralba, A. (2012), A benchmark of computational models of saliency to predict human fixations, *Technical Report, MIT*.

Judd, T., Ehinger, K., Durand, F. and Torralba, A. (2009), Learning to predict where humans look, *2009 IEEE 12th International Conference on Computer Vision*, IEEE, pp. 2106–2113.

Kahneman, D. (2003), Maps of bounded rationality: Psychology for behavioral economics, *American Economic Review* **93**(5), 1449–1475.

Kahneman, D. and Tversky, A. (1972), Subjective probability: A judgment of representativeness, *Cognitive Psychology* **3**(3), 430–454.

Kahneman, D. and Tversky, A. (1973), On the psychology of prediction., *Psychological Review* **80**(4), 237.

Kennedy, R., Clifford, S., Burleigh, T., Waggoner, P. and Jewell, R. (2018), The shape of and solutions to the MTurk quality crisis, *Available at SSRN*.

Krajbich, I. and Rangel, A. (2011), Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions, *Proceedings of the National Academy of Sciences* **108**(33), 13852–13857.

Krajbich, I., Armel, C. and Rangel, A. (2010), Visual fixations and the computation and comparison of value in simple choice, *Nature Neuroscience* **13**(10), 1292.

Le Meur, O. and Baccino, T. (2013), Methods for comparing scanpaths and saliency maps: strengths and weaknesses, *Behavior Research Methods* **45**(1), 251–266.

Li, J. and Yu, J. (2012), Investor attention, psychological anchors, and stock return predictability, *Journal of Financial Economics* **104**(2), 401–419.

Lusardi, A. and Mitchell, O. S. (2011), Financial literacy around the world: An overview, *Journal of Pension Economics and Finance* **10**(4), 497–508.

McLean, R. D. (2010), Idiosyncratic risk, long-term reversal, and momentum, *Journal of Financial and Quantitative Analysis* **45**(4), 883–906.

Mizrach, B. and Weerts, S. (2009), Highs and lows: a behavioural and technical analysis, *Applied Financial Economics* **19**(10), 767–777.

Mohrschladt, H. and Schneider, J. C. (2018), The Idiosyncratic Volatility Puzzle and its Interplay with Sophisticated and Private Investors, *Available at SSRN 3131703*.

Mormann, M. M. and Frydman, C. (2017), The role of salience and attention in choice under risk: An experimental investigation, *Working Paper*.

Mullett, T. L. and Stewart, N. (2016), Implications of visual attention phenomena for models of preferential choice., *Decision* **3**(4), 231.

Nolte, S. and Schneider, J. C. (2018), How price path characteristics shape investment behavior, *Journal of Economic Behavior and Organization* **154**, 33–59.

Pelli, D. G. (1997), The VideoToolbox software for visual psychophysics: transforming numbers into movies, *Spatial Vision* **10**, 437–442.

Prelec, D. (1998), The probability weighting function, *Econometrica* **66**(3), 497–527.

Raghubir, P. and Das, S. R. (2010), The long and short of it: why are stocks with shorter runs preferred?, *Journal of Consumer Research* **36**(6), 964–982.

Ross, J., Irani, L., Silberman, M., Zaldivar, A. and Tomlinson, B. (2010), Who are the crowdworkers?: shifting demographics in mechanical turk, *CHI'10 Extended Abstracts on Human factors in Computing Systems*, ACM, pp. 2863–2872.

Stambaugh, R. F., Yu, J. and Yuan, Y. (2015), Arbitrage asymmetry and the idiosyncratic volatility puzzle, *The Journal of Finance* **70**(5), 1903–1948.

Tversky, A. and Kahneman, D. (1992), Cumulative prospect theory: An analysis of decision under uncertainty, *Journal of Risk and Uncertainty* **5**(4), 297–323.

Wakker, P. P. (2010), *Prospect theory: For risk and ambiguity*, Cambridge University Press.

Walasek, L., Mullett, T. L. and Stewart, N. (2018), A meta-analysis of loss aversion in risky contexts, *Working Paper*.

Wu, G. and Gonzalez, R. (1996), Curvature of the probability weighting function, *Management Science* **42**(12), 1676–1690.

Online Appendix: Decision weights for experimental asset prices based on visual salience

Devdeepta Bose,[1] Henning Cordes,[2] Sven Nolte,[3]
Judith C. Schneider,[4] Colin F. Camerer[5]

[1]California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA
Email: devbose@caltech.edu
[2]University of Münster, Universitätsstr. 14-16, 48143 Münster, Germany
Email: henning.cordes@wiwi.uni-muenster.de
[3]Institute for Management Research, Radboud University, Nijmegen, The Netherlands
Email: s.nolte@fm.ru.nl
[4] Leuphana University Lüneburg, Universitätsallee 1, 21335 Lüneburg, Germany
Email: judith.schneider@leuphana.de
[5]California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA
Email: camerer@hss.caltech.edu

May 7, 2020

This Online Appendix provides additional material and analyses for ”Decision weights for experimental asset prices based on visual salience”.

- Section 1 provides details on the various strands of literature whose concepts are utilized in this paper.
- Figure S1 of Section 2 lines up with Figure 1 of the main analyses depicting additional specifications for CPT and Salience.
- Section 3 presents the statistics on the three decision weights for price paths used in Study I, with Section 3.1 summarizing the distance between the decision weights and empirical frequencies, Section 3.2 summarizing the distance between the decision weights, and Section 3.3 summarizing the correlation between the decision weights.
- Section 4 presents the price paths used for the proof of concept in Section 2.1 of the main paper.
- Section 5 contains two examples from the Sketch category of images that were used to train SAM, that most resemble the shape of price paths.
- Section 6 presents the statistical measures and analyses on the relation of visual salience and statistical features of price paths.
- Section 7 provides the full set of instructions for the MTurk experiment where the convexity score for picking the CRSP stocks is described in Section 8.
- Section 9 provides an overview over the course of the experiment used for Study III and Section 10 shows the full set of instructions for Study III.
- Section 11 presents additional analyses for the MTurk studies, including a comparison of the MTurk sample with the US population and additional robustness tests.

- Section 12 compares the sample of Study III with a sample representative for German stockholders and provides cumulative density functions for distributions of coefficients and p-values of our parameter sensitivity analysis.

- Section 13 repeats all the analyses for all three studies in the main paper using price differences as our basis for analysis instead of returns.

1 Literature Review

This paper adds to four different strands of literature: First, we contribute to the literature that explores how the framing or presentation of historical returns can stimulate different biases and heuristics within investors, and influence their decision making. Recent experimental evidence suggests that price path shapes influence asset attractiveness and investment behavior (Grosshans and Zeisberger, 2018; Nolte and Schneider, 2018). Moreover, the representation of the price chart, i.e., the visual properties (e.g., scaling) of charts influence risk perception and return expectations (Huber and Huber, 2018). Most information brochures display a visualization of an asset's price path. The average investor is easily overwhelmed by the complexity and variety of offered products. Thus, an easily accessible graphical presentation is expected to have a significant impact on decision making. There are several experimental studies which address the impact of price path presentation on investment behavior. Glaser et al. (2019) analyze how representation of returns vs. prices shape individual forecasts of asset performance. Weber et al. (2005) find that the presentation format of historical asset returns influences expected asset volatility. Diacon and Hasseldine (2007) find that the presentation format of asset returns has a significant effect on subjects' perceptions of risk and return. Glaser et al. (2007) associate beliefs in mean reversion or trend chasing with framing effects. Asset forecasts vary depending on whether investors are presented with returns or with prices, and the authors attribute these differences to the representativeness heuristic. They conclude that surveys on asset forecasts should rely on price presentation rather than returns, see also Czaczkes and Ganzach (1996). Grosshans and Zeisberger (2018) present subjects with variously shaped price paths, and ask them to record their levels of satisfaction if they had hypothetically held such an asset. They conclude that the shape of the historical asset price path might shape future investment behavior through reference point effects. The visual salience model shows that visual salience of certain returns in a historical price path generate attention towards those returns, which is what drives the predictability of invested amounts in all three experimental studies.

Second, this paper adds to the literature applying visual salience to economic decision making. Outputs from visual salience algorithms have been shown to predict economic choice in other domains. Towal et al. (2013) find that objective visual properties such as saliency, and subjective properties such as value, are both interactively useful in predicting decisions in multiple-alternative forced-choice environments, e.g., choosing snacks from a vending machine. Li and Camerer (2019) find that visual salience of focal points in images are useful in predicting winners in coordination games and hider-seeker games. Thus, visual salience has been used in fields such as consumer choice and empirical game theory. This paper extends the scope of its use by introducing this concept to the field of behavioral finance.

Third, this paper adds to the literature which proposes different theoretical models on how to integrate past returns into forming expectations about future returns. A centerpiece of any investment choice is to form expectations about the future performance of investment opportunities at the time of judgment. These expectations crucially depend on available information and how this information is processed. Barberis et al. (2001, 2016) argue that prospect theory values based on past returns predict future returns. One key assumption in their model is that the presentation format of the price charts is irrelevant. The ordering of returns is only captured in an extension of their model, by an exponential overweighting of the most recent returns, i.e., only temporal ordering of the returns is relevant. In another class of models by Bordalo et al. (2012, 2013) and applied in Cosemans and Frehen (2017) and Gödker and Lukas (2017), the magnitude of the return difference to some reference value determines the weight assigned to the return. Thus, this version of salience is context dependent, i.e., dependent on the other available assets. The salient (extreme) returns attract attention and predict future returns. One limitation of these models is that the salience rank of a past return is independent of the actual probability of its occurrence. Moreover, the temporal ordering of returns and presentation format are irrelevant in these models. In contrast, the proposed model based on the Salience Attentive Model (Cornia et al., 2018) assumes that the weighting depends only on visual features. Thus, presentation format and temporal ordering of returns both matter. The returns close to salient points are overweighted. While in the models of Barberis et al. (2016) and Bordalo et al. (2012, 2013), the weighting depends on specific functional forms, the weighting implied by visual salience is parameter free and depends on feature extraction by a machine learning algorithm.

Lastly, the paper adds to the literature on the role of attention in financial decision making. Given the almost infinite amount of stock market information readily available to any investor, paying attention to every piece would exceed individuals' cognitive abilities (Kahneman, 1979). Only a few stocks end up in the choice set of investors. A firm which catches more investor attention is more likely to be included in the investors' portfolio (Odean, 1999; Busse and Green, 2002; Barber and Odean, 2007). A stock can attract attention through media coverage (e.g., Fang and Peress, 2009), lottery-like features (e.g., Kumar, 2009), high daily trading volume (e.g., Barber and Odean, 2007), extreme daily returns (e.g., Bali et al., 2011; Ungeheuer, 2017), daily newpapers winners and losers (e.g., Kumar et al., 2019), or related events (e.g., Koester et al., 2016, for earnings surprises). All these papers study indirect measures of attention. An exception is Da et al. (2011) who provide evidence that stock-related Google search volume reflects the attention paid by private (not sophisticated) investors as they gather information using Google. Thus, Google search volume indices provide a timely measure of firm-level, private investor attention. Which features of these stocks determine their attractiveness also depends on which feature grabs the most attention. Inattention to information affects choices of private investors as well as analysts (Koester et al., 2016). While these approaches are various their common idea is to determine the choice set of stocks investors are paying attention to.

In contrast, within this choice set, we look at attention towards path characteristics of an asset measured by using eye-tracking, i.e., by using a specific computational model of where attention is predicted to be allocated to.

In summary, we propose a model that estimates salience based on actual measurements (eye-tracking) while previous models on salience miss the visual impact of return representation.

2 Additional Visualizations for CPT and Salience Weights

Figure S1: CPT and Salience

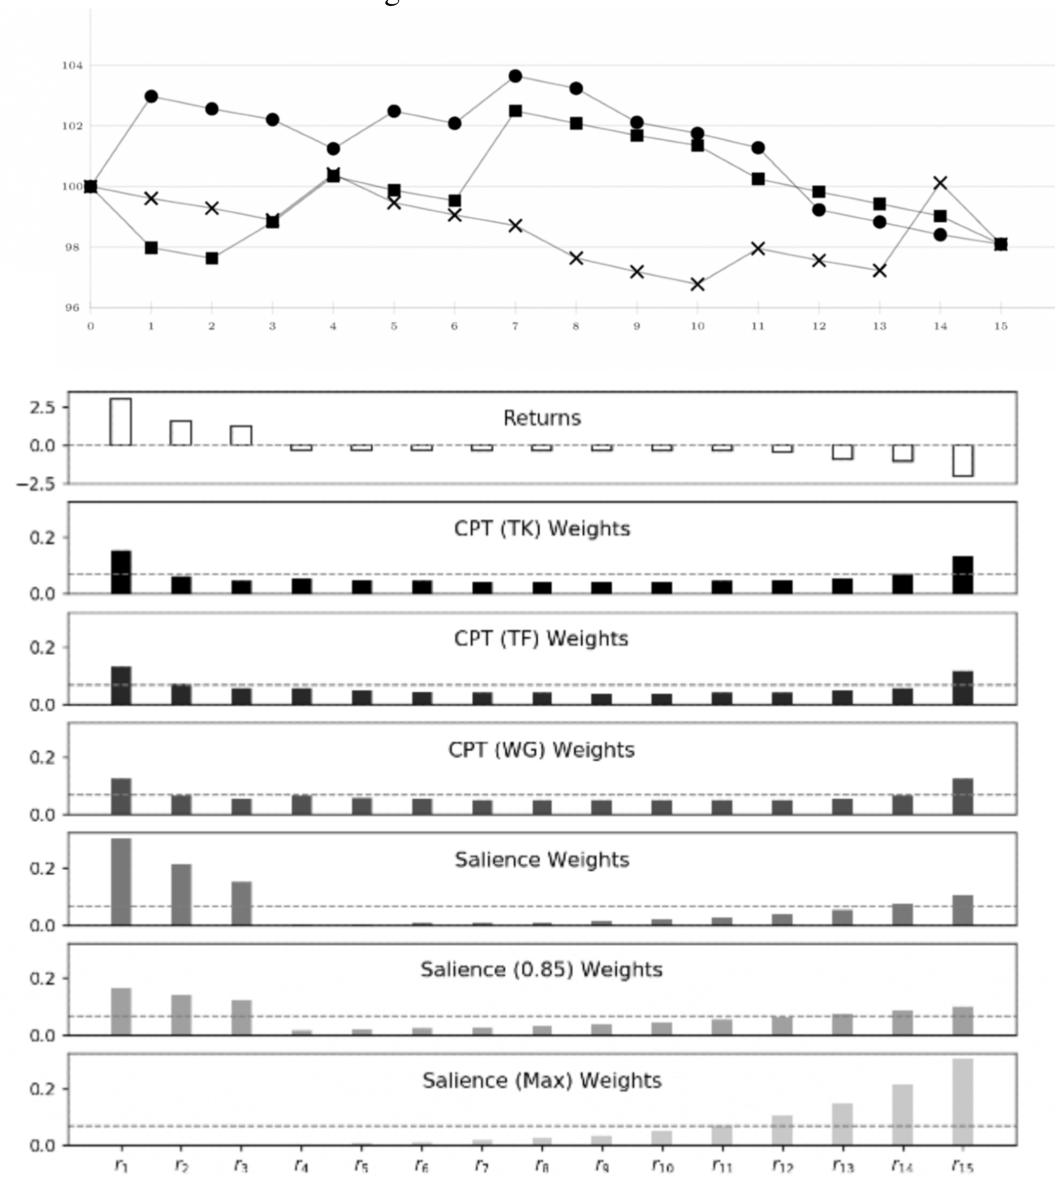

The figure shows three price paths of 15 periods historical realizations constructed from the returns in the graph below. The returns are the result of stationary process $x_{t+1} = 0.6x_t + \varepsilon_t$. The bar charts present decision weights based on CPT and salience theory.

The bar chart shows the returns for the historical realizations for the same three price paths in the main paper, arranged in descending order. It also shows the decision weights associated with the returns when using the CPT value function along with the Tversky and Kahneman (1992) (TK) weighting function, the Tversky and Fox (1995) (TF) weighting function, and the Wu and Gonzalez (1996) (WG) weighting function, shown below in Equations 1, 2, and 3 respectively. The figure also shows the Bordalo et al. (2013) (Salience) decision weights (shown below in Equation 4), along with variations where the ν parameter is changed from 0.7 to 0.85, and the reference level $\bar{x}_k$ is changed from the mean of the historical returns to the maximum of the historical returns. The key idea here is that the decision weights for all variations of CPT and Salience decision weights are the same for all three price paths. Furthermore, the different weighting functions for CPT do not generate meaningfully different decision weights for these price paths, but Salience decision weights are quite sensitive to parameter and reference level specifications.

- Tversky and Kahneman (1992) (TK) weighting function

$$w^+(p) = \frac{p^{\delta_+}}{(p^{\delta_+} + (1-p)^{\delta_+})^{\frac{1}{\delta_+}}} \text{ and } w^-(p) = \frac{p^{\delta_-}}{(p^{\delta_-} + (1-p)^{\delta_-})^{\frac{1}{\delta_-}}} \quad (1)$$
$$\text{and } \delta^+ = 0.69, \delta^- = 0.61$$

- Tversky and Fox (1995) (TF) weighting function

$$w^+(p) = \frac{\gamma^+ p^{\delta^+}}{\gamma^+ p^{\delta^+} + (1-p)^{\delta^+}} \text{ and } w^-(p) = \frac{\gamma^- p^{\delta^-}}{\gamma^- p^{\delta^-} + (1-p)^{\delta^-}} \quad (2)$$
$$\text{and } \delta^+ = 0.6, \gamma^+ = 0.65, \delta^- = 0.65, \gamma^- = 0.84$$

- Wu and Gonzalez (1996) (WG) weighting function

$$w(p) = \exp(-(-\ln(p))^{\delta}) \quad \text{and } \delta = 0.74 \quad (3)$$

- Bordalo et al. (2013) (Salience) decision weights

$$\pi_k^{Sal} = p_k \times h_k \text{ with the salience weight } h_k = \frac{\nu^{\kappa_k}}{\sum_i \nu^{\kappa_i} p_i} \quad \nu = 0.7 \text{ or } 0.85 \quad (4)$$

where κ_k denotes the salience rank of x_k, which is measured by:

$$\kappa_k = \sigma(x_k, \bar{x}_k) = \frac{|x_k - \bar{x}_k|}{|x_k| + |\bar{x}_k| + \theta} \text{ where } \bar{x}_k = mean(x_k) \text{ or } max(x_k)$$

3 Statistics on decision weights

3.1 Distance from empirical frequencies

This section provides summary statistics and distributions of distances between decision weights and empirical frequencies for price paths from Study I. Distance is calculated on a path level by taking the sum of absolute differences between decision weights and empirical frequencies across all 250 daily returns. In total, there are 1,000 paths, so summaries and density functions are based on 1,000 distance values. Salience Theory is based on rank, so the sum of all differences are identical for each price path. The procedure is similar for CPT, where different numbers of positive and negative returns result in small differences between paths. Distances for the Visual Salience measure are independent of return size but only rely on the degree of visual salience, which results in a wider distribution of distances.

Distance from empirical frequencies

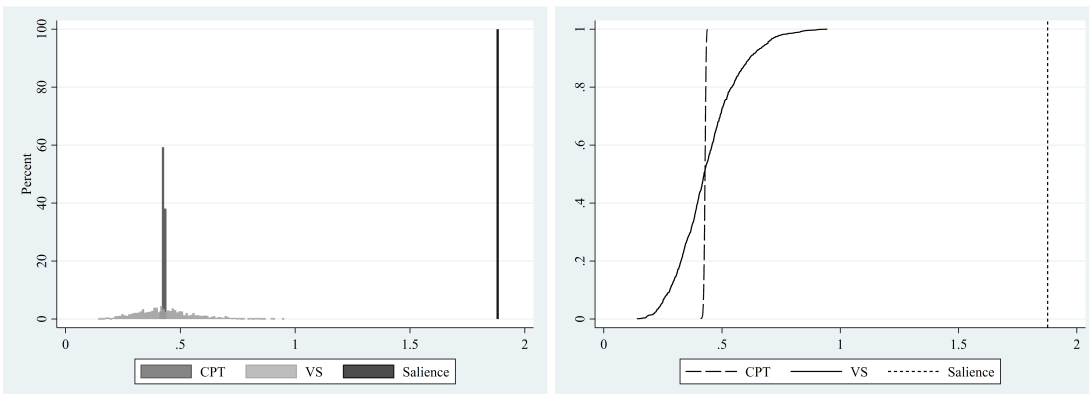

	Obs.	Mean	Median	Std. Dev.	Min	Max
CPT	1000	0.4279	0.4280	0.0041	0.4096	0.4380
Salience	1000	1.8766	1.8766	0.0000	1.8766	1.8766
VS	1000	0.4389	0.4239	0.1329	0.1423	0.9432

Figure S2: Distance of decision weights from empirical frequencies

The figure shows the distribution of absolute distances between empirical return frequencies and associated decision weights per theory. The left graph displays probability density functions of distances, the right graph the corresponding cumulative density functions. The table reports mean, median, standard deviation, minimum, and maximum of distances for each theory.

3.2 Distance between theories

This section provides summary statistics and distributions of distances between theories for price paths from Study I. Distance is calculated on a path level by taking the sum of absolute differences between decision weights across all 250 daily returns. In total, there are 1,000 paths, so summaries and density functions are based on 1,000 distance values. Because Salience Theory has a lot of weights at or very close to zero, distances are generally large between Salience Theory and the other two theories.

Distance between theories

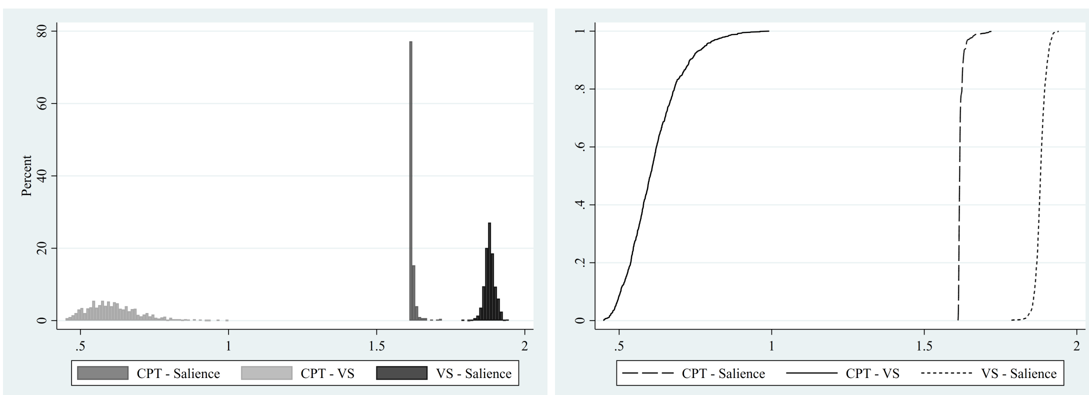

	Obs.	Mean	Median	Std. Dev.	Min	Max
CPT - Salience	1000	1.6193	1.6159	0.0121	1.6101	1.7197
CPT - VS	1000	0.6133	0.6019	0.0906	0.4496	0.9904
VS - Salience	1000	1.8815	1.8810	0.0178	1.7859	1.9403

Figure S3: Distance of decision weights between theories

The figure shows the distribution of absolute distances of decision weights between theories. The left graph displays probability density functions of distances, the right graph the corresponding cumulative density functions. The table reports mean, median, standard deviation, minimum, and maximum of distances between each pair of theories.

3.3 Correlation between theories

This section provides summary statistics and distributions of correlations between decision weights for price paths from Study I. Correlations are calculated on a path level between decision weights across all 250 daily returns. In total, there are 1,000 paths, so summaries and density functions are based on 1,000 correlations. Salience Theory and CPT exhibit strong correlations of around 80% on average, as both theories overweigh more the extremes of the return distribution. Correlation between Visual Salience and CPT or Salience Theory decision weights is indistinguishable from zero, supporting our claim that visual salience adds a level of salience beyond that captured in other theories.

Correlations between theories

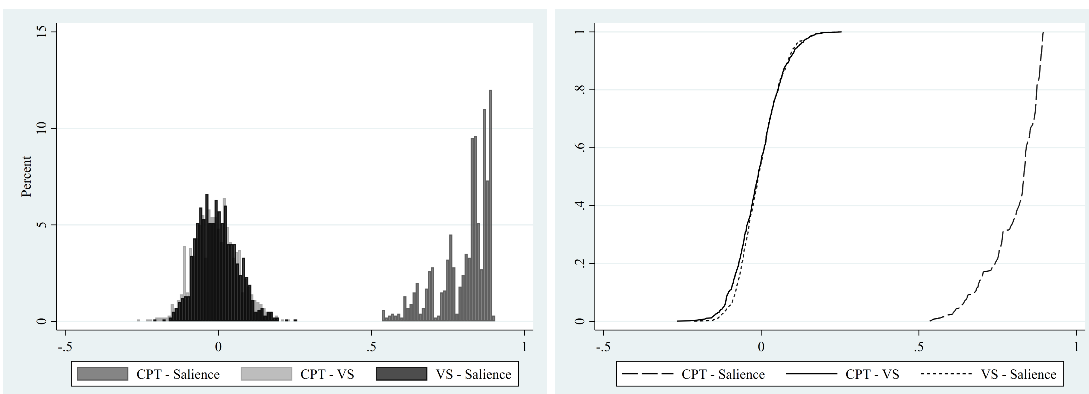

	Obs.	Mean	Median	Std. Dev.	Min	Max
CPT - Salience	1000	0.8037	0.8325	0.0835	0.5342	0.8948
CPT - VS	1000	-0.0086	-0.0100	0.0735	-0.2658	0.2536
VS - Salience	1000	-0.0039	-0.0087	0.0661	-0.2124	0.2482

Figure S4: Correlation of decision weights between theories

The figure shows the distribution of correlations of decision weights between theories. The left graph displays probability density functions of correlations, the right graph the corresponding cumulative density functions. The table reports mean, median, standard deviation, minimum, and maximum of correlations between each pair of theories.

4 Price paths from Nolte and Schneider (2018)

Figure S5: Price paths for proof of concept

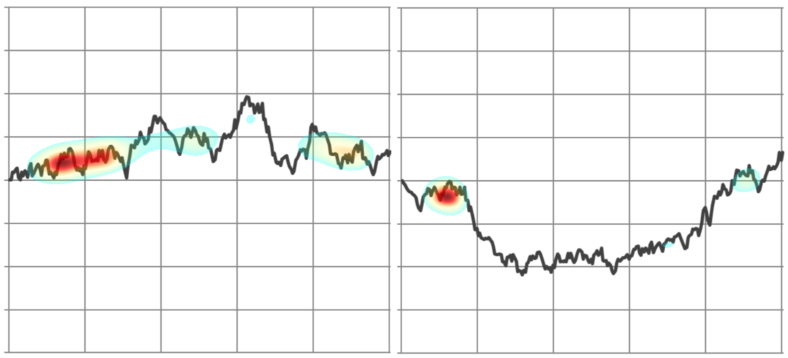

(a) Highest VS value (b) Lowest VS value

The figure shows the two price paths from Nolte and Schneider (2018) which have the same set of returns, but the highest value (left) and the lowest value (right) based on visual salience decision weights. Subjects' investments are ten percentage points higher for the chart on the left.

5 SAM Training Images

Figure S6: SAM training images that most resemble price paths

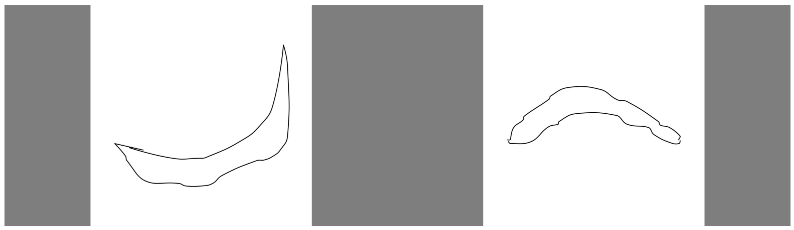

(a) Sketch One (b) Sketch Two

The figures are from the Sketch category of SAM training images, that most resemble the shape of price paths.

6 Statistical Features and Visual Salience

We rely on the large literature exploring properties of past returns and prices to predict subsequent returns to identify relevant statistical features that may act as substitutes for visual salience. Most of these measures capture the notion that retail investors often buy attention grabbing, lottery-like stocks with high total volatility. These measures encompass the 52-week-high (George and Hwang, 2004), other price highs (Mizrach and Weerts, 2009), the maximum return of the previous months (Bali et al., 2011), the run-length (Raghubir and Das, 2010), historical skewness (Chen et al., 2001; Conrad et al., 2014), momentum (Jegadeesh and Titman, 1993), and high idiosyncratic volatility (Stambaugh et al., 2015). Additionally, we include measures from the literature on the perception of price sequences, which include average, maximum, and minimum price (Baucells et al., 2011), if the price is above or below the starting price (Nolte and Schneider, 2018), or the difference in standard deviation between the first and second half of the (sub-)period (Borsboom and Zeisberger, 2020). We adopt these measures to correlate them with saliency predictions by SAM. Table 1 provides an overview.

Table 1: Statistical features: description, summary statistics, and correlation with SAM weights

Statistical feature	Description	mean	std. dev.	min	max	correlation with π^{VS}
period	number of period numbered early (1) to most recent (10)	5.5	2.8724	1	10	-0.4044
runlength	no. of previous returns with the same sign as current return	4.9181	1.6792	2	16	-0.1326
autocorrelation r_t, r_{t-1}	average correlation between subsequent daily returns within a period	-0.0735	0.2007	-0.7872	0.7379	-0.1156
distance from starting price	average absolute differences of prices within a period from starting price of price path	15.2128	20.6039	0.0185	312.40	-0.1129
spread in % of path spread	difference between maximum and minimum price in a period. relative to difference between maximum and minimum price of price path	0.2761	0.1424	0.0232	1	-0.1064
spread in % of max price	difference between maximum and minimum price in a period. relative to maximum price in that period	0.0837	0.0613	0.0006	0.6364	-0.0927
std. dev. in % of path std. dev.	standard deviation of daily returns in a period. relative to standard deviation of daily returns of price path	0.9642	0.2794	0.1969	2.7246	0.0845
min price in % of path min price	minimum price in a period. relative to minimum price of price path	1.1574	0.2139	1	3.5259	0.0807
max return in % of path max return	maximum daily return in a period. relative to maximum return of price path	0.5518	0.2543	0.0348	1	0.0758
Relative Strength Index (RSI)	number of positive daily returns in a period divided by number of negative daily returns in a period	1.2	0.5029	0.25	11.5	-0.0744
spread	difference between maximum and minimum price in a period	9.8445	9.3912	0.0635	161.11	-0.0612
no. of gains	number of positive daily returns in a period	13.4913	2.3969	5	23	-0.0537
loss domain	1 if any price in a period is below the starting price of the price path. 0 otherwise	0.4416	0.4966	0	1	0.0434
skewness	skewness of daily return distribution in a period	-0.0273	0.9744	-4.4248	4.4089	0.0372
max price in % of path max price	maximum price in a period. relative to maximum price of price path	0.8874	0.1140	0.2409	1	0.035
jump	1 if any daily return in a period belongs to the largest absolute returns of all returns in the sample	0.0212	0.1441	0	1	0.0349
min price	minimum price in a period	103.17	22.302	19.305	374.5	0.0301
min return in % of path min return	minimum return in a period. relative to minimum return of price path	0.5290	0.2625	0.0229	1	0.0217
skewness in % of path skewness	skewness of daily return distribution in a period. relative to skewness of daily return distribution of price path	0.0999	68.571	-2669.6	2624.1	0.0184
average price	average price in a period	108.03	24.283	22.092	412.40	0.0171
std. dev.	standard deviation of daily returns in a period	0.0140	0.0101	0.0001	0.1173	0.0165
max return	maximum return in a period	0.0320	0.0311	0.0001	0.5171	0.0137
average return in % of path average return	average daily return in a period. relative to average daily return of price path	1	63.523	-2178.9	3244.1	0.006
autocorrelation r_t, r_{t-1} in % of path autocorrelation	average correlation between subsequent daily returns within a period. relative to average correlation between subsequent daily returns of price path	1.5915	55.505	-2372.1	2041.56	0.0039
max price	maximum price in a period	113.01	26.960	25.222	451.23	0.0036
momentum relative to path momentum	average return in last quarter of a period. relative to average return in last quarter of price path	1	0.0031	0.9751	1.0263	0.0009
momentum	average return in last quarter of a period	1.0005	0.0032	0.9696	1.0310	0.0008
min return	minimum daily return in a period	-0.0307	0.0294	-0.4838	-0.0003	0.0007
average return	average daily return in a period	0.0006	0.0032	-0.0280	0.0360	0.0002

The table reports descriptions and summary statistics on statistical features used in our analysis. The last column reports the correlation coefficient of the respective feature with SAM weights. Features are sorted by their absolute correlation with SAM weights.

We analyze whether visual salience can be subsumed by statistical features of price paths by checking if our set of statistical moments are highly correlated with the visual salience of areas on the price chart as predicted by SAM. Using the 1000 price paths analyzed in Study I, we split the 250 daily returns in each path into 10 blocks of 25 returns each. For each 25-return-block, we have SAM predict the visual salience of that particular region. We normalize SAM predictions to obtain relative weights for each of the 10 blocks that add up to one. We then explore if certain statistical features for each block of 25 returns can predict the relative weight assigned to that block. In a first step, we report Pearson correlation coefficients between SAM weights and each individual statistical feature in Table 1. The most prominent feature with a correlation of -0.4044 is $period$, i.e., whether the respective subperiod occurs early or late. This makes sense from a visual perspective, as $period$ can also be interpreted as an indicator how far left the respective subperiod is in the price path image. Other features that appear high on the list are measures of return sequences like runlength and autocorrelation, and distance measures like the difference between minimum and maximum price ('spread'). The highest correlation for these measures is a little above 13% though, which can be considered rather low. The second step in our empirical strategy would be to regress SAM weights on all statistical metrics. However, since many of these metrics are based on the same fundamental concept, multicollinearity in such a regression would be extremely high. Therefore, we first run a principal component analyis (PCA) among all statistical features to identify common components that can then be included in a regression analysis. Results of this PCA are displayed in Table 2. Based on Eigenvalues $>$ 1, eight components emerge. Regressing SAM weights on these eight components, we find that the time-period of the blocks (earlier blocks are more salient), low autocorrelation within block (returns fluctuating up and down are more salient), and dispersion of returns within block (larger spreads are less salient) are statistical features that are most predictive of SAM weights (Table 3). This is consistent with the earlier analysis of correlation coefficients. However, only 8.59%

of the total variance of SAM weights can be explained by these features. Even including twelve components to accomodate those with an Eigenvalue > 0.8 only leads to an improvement of R^2 to just below 20%. We conclude that even though there are some correlations between SAM weights and statistical features of the respective area in the price chart, visual salience provides a way of interpreting price charts that goes beyond what is possible with standard statistical metrics.

Table 2: Statistical features: Principal Component Analysis

Statistical Feature	Comp1	Comp2	Comp3	Comp4	Comp5	Comp6	Comp7	Comp8	Comp9	Comp 10	Comp11	Comp12
Eigenvalue	5.6443	4.7121	3.8542	2.1882	1.8790	1.4644	1.0494	1.0058	0.9989	0.9757	0.8988	0.8631
min price	0.4708											
average price	0.4646											
max price	0.4483											
distance from starting price	0.3252											
min price in % of path min price	0.2994											
loss domain	-0.3139											
spread in % of max price		0.4542										
std. dev.		0.4385										
spread		0.3859										
max return		0.3358	0.2627									
min return		-0.3778	0.2634									
skewness			0.4791	-0.3375								
momentum relative to path momentum			0.4081									
momentum			0.4001									
average return			0.3987									
max return in % of path max return			0.2737		0.4907							
min return in % of path min return			-0.2605		0.527							
no. of gains				0.5989								
Relative Strength Index (RSI)				0.587								
std. dev. in % of path std. dev.					0.5967							
max price in % of path max price						0.694						
spread in % of path spread						0.6533						
runlength							0.6934					
autocorrelation r_t, r_{t-1}							0.6778					
period								0.9525				
jump									0.9905			
autocorrelation r_t, r_{t-1} in % of path autocorrelation										0.9979		
skewness in % of path skewness											0.9985	
average return in % of path average return												0.9996

The table reports rotated factor loadings of the 12 factors with Eigenvalues greater than 0.8. Eigenvalues are listed in the first row of the table. Loadings smaller than 0.25 are blanked out to enhance readability of the table. Statistical features are ordered by their loadings on the respective components, prioritizing components with a larger Eigenvalue.

Table 3: Regressions for components of statistical features

	(1) π^{VS}	(2) π^{VS}
Comp1	-0.00227***	-0.00227***
	(0.000535)	(0.000524)
Comp2	-0.000238	-0.000238
	(0.000592)	(0.000543)
Comp3	0.00106**	0.00106**
	(0.000530)	(0.000515)
Comp4	-0.00121**	-0.00121**
	(0.000567)	(0.000528)
Comp5	-0.00926***	-0.00926***
	(0.000548)	(0.000513)
Comp6	-0.0131***	-0.0131***
	(0.000575)	(0.000534)
Comp7	0.00544***	0.00544***
	(0.000668)	(0.000543)
Comp8	0.00345***	0.00345***
	(0.000686)	(0.000509)
Comp9		0.00335***
		(0.000426)
Comp10		0.00668***
		(0.000489)
Comp11		-0.0131***
		(0.000599)
Comp12		0.0126***
		(0.000637)
Constant	0.100***	0.100***
	(0.000572)	(0.000538)
Observations	10,000	10,000
R^2	0.086	0.194

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The table reports regressions of SAM weights (π^{VS}) on standardized components predicted by a principal component analysis including all investigated statistical features. The first column only includes components with an Eigenvalue greater than 1, the second column also adds components with an Eigenvalue greater than 0.8. Robust standard errors are reported in parentheses.

7 Instructions for Study I and II

7.1 Introduction

Welcome and thank you for participating – your assistance is greatly appreciated! Your answers are very important to us. Please take your time to read all instructions and questions carefully. All your answers will be treated anonymously, with strict confidentiality, and will be used for research purposes only. Completing the experiment will take about 20 minutes. At the end, you will receive a code that you have to enter on Amazon's Mechanical Turk website to receive your payment. The payment consists of a fix payment ($2.00) plus a variable payment. The variable payment will depend on your choices and a random component. Participants will receive a variable payment of about $1.00 on average. Please note that your variable payment might be higher or lower than this. Please do not refresh the page at any time or click "back" or "forward" in your browser. This may cause a crash and you will have to start again. A "continue" button will appear as soon as all relevant texts are shown on screen.

General Topic

We are researching financial investment behavior. You do not need expert knowledge on this topic – all relevant information will be given to you before you provide your choices. You will see 10 different charts displaying stock price developments over one year. The price charts are based on the historical stock price development of real companies, but we do not mention names. All graphs are indexed to 100, so the charts always start at a price of 100.

Stock investments

Below, you see an example for a stock price chart (Figure S7) – just like the ones you will see when you make your decisions. The horizontal axis of the chart represents the time frame (January to December, i.e. 1 year). The vertical axis represents the stock price (measured in monetary units). In this example, the stock price starts at 100 monetary units in January. Over the course of the year, the stock price gradually increases up to 121 monetary units in December. Hence, the return over this year is +21 percent. It is not possible to predict precisely how the stock price will develop next year. However, the displayed chart can give you an idea about the risks and returns of investing into this stock.

Figure S7: Example of a price chart

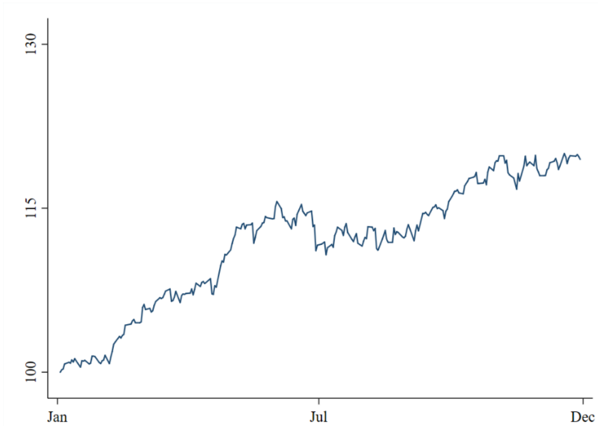

The figure shows a screen shot of the exemplary price chart shown to the participants in the introductory part of the experiment.

Your task

For each price chart you see, we ask you to indicate:

1. How attractive you find an investment into this stock.

2. How risky you deem an investment into this stock.

3. What return you expect over the following 12 months.

4. How much of a 1,000 monetary unit endowment you are willing to invest in the stock for the following 12 months. The remainder is put in a safe bank account.

Detailed explanations on question 4 will be provided on the following page.

Your investment decision

You have 1,000 monetary units. You want to invest this money for the following year. Your task is to decide what part of your endowment you want to invest into a stock. The remainder is put into a safe bank account. The stock price development over the past year is displayed in a chart (like the one you have seen on the previous screens). The chart can give you an idea about the risk and return of investing over the next year. Over the next year, the stock price development can be very similar to the one displayed in the chart, but it can also be very different. Hence, the return of your stock investment might be similar to the one displayed in the chart, but it can also be much higher or much lower. The non-invested part of your endowment is put into a bank account. This money is safe, i.e., it neither increases nor decreases over the next year. Your final wealth after 1 year is the value of your stock investment after one year plus the money in the bank account.

How you will be paid

You will receive $0.001 for every monetary unit of your final wealth in the investment task (i.e. the final value of your stock investment plus the money in the bank account).

Examples:

Imagine that you invest 500 monetary units into a stock. The remaining 500 monetary units are automatically put into the safe bank account.

- If the stock price increases by 50% in the next year, your stock investment is then worth 750 monetary units. Hence, your overall wealth is 1,250 monetary units (750 from stock investment and 500 from bank account), and your variable payment is $1.25 in this example.

- If the stock price decreases by 50% in the next year, your stock investment is then worth only 250 monetary units. Hence, your overall wealth is 750 monetary units (250 from stock investment and 500 from bank account), and your variable payment is \$0.75 in this example.

Additionally to this variable payment, you will receive a fix payment of \$2.00 for sure.

Summary

You will see a total number of 10 price charts. For each price chart, you will judge how attractive and how risky you find the stock and what return you would expect for the following year. You also decide what part of your 1,000 monetary units endowment you invest into the stock for the following year. The remainder is put into a safe bank account where you cannot gain or lose any money. The return of your stock investment is based on this stock's actual return in the subsequent year. The computer will randomly select one of your 10 investment decisions after you complete this questionnaire. You receive \$0.001 for every monetary unit of your final wealth in the selected decision scenario. Additionally, you will receive a fix payment of \$2.00 for sure.

7.2 Main Experiment

Figure S8 presents exemplary screen shots of the four decisions the participants had to take during the main part of the experiment.

7.3 Final Questionnaire

Survey:

Please tell us something about yourself now. Your answers will help us to draw more meaningful conclusions from your choices.

1. How old are you? <dropdown menu with 18-99>

2. Please state your gender? <radio buttons with ”male”, “female”, “other”>

Figure S8: Interface for questions in the main part of the experiment

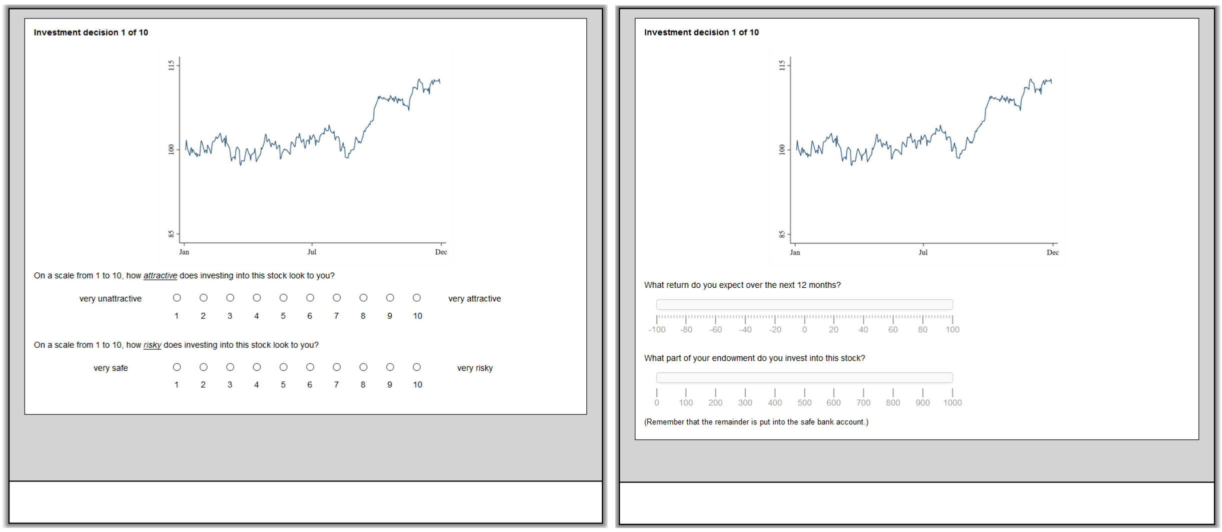

(a) Questions 1 and 2

(b) Questions 3 and 4

The participant had to take 10 decisions on different price paths. The Figure shows how the interface for the four questions in the main part looked like.

3. What's your nationality? <radio buttons with "US-American" and "Other – please specify" (plus textbox for the latter)>

4. What is your highest educational degree attained? <radio buttons with "Some high school", "High school degree", "College degree", "Master degree", "Doctorate degree", "Other">

5. What is your occupation? <textbox>

6. Are you in full employment? <radio buttons with "Yes" and "No">

Questions on the experiment

1. If the stock price increased from 100 to 145 Dollar over one year, what is the return on this stock? <radio buttons with "-45% / +45% / +55% / +145% / you cannot say / don't know" as answer possibilities>

2. If you invested 1,500 monetary units of your endowment of 2,000 monetary units into the stock, how much was put into the bank account? <radio buttons with "0 / 500 / 1,500 / 2,000 monetary units / you cannot say / don't know" as answer possibilities>

3. What is fourteen minus 8? <radio buttons with "0 / 2 / 6 / 10 / you cannot say / don't know" as answer possibilities>

4. If the stocks return was +10% in the last year, what is the stocks return in the subsequent year? <radio buttons with "-10% / 0% / 10% / 110% / you cannot say / don't know" as answer possibilities>

Questions on the experiment – Please evaluate the following statements: (1-7 Likert Scale)

1. Participating in this experiment was fun.

2. I relied on my gut feeling for my decisions.

3. I wanted to play it safe with my decisions.

4. I thought long about my decisions.

5. I believe that the past development of a stock price conveys information about its future development.

6. The understandability of the instructions was good.

Attitude towards risks – Please evaluate the following statements: (1-7 Likert Scale)

1. I do not mind taking risks in general.

2. I do not mind taking risks in financial investments.

3. I try to minimize the chance of losing money in my financial investments.

4. Losing some money with an investments would hurt me more than I would enjoy gaining the same amount of money.

5. Please click on the first button on the scale below.

6. Extreme returns of a stock catch my attention.

7. I do not mind taking risks in my leisure time (e.g. by practicing risky hobbies).

8. I do not mind driving a car without seatbelt fastened.

Your financial situation

1. How much money can you spend each month? <radio button> <textbox> Dollars <radio button> I prefer not to answer.

2. Are you the primary decision maker regarding financial matters in your household? <radio buttons with “Yes” and “No”>

3. How many people are living in your household? <dropdown menu with 1-20> people

4. How many credit cards do you have? <dropdown menu with 0-20> credit cards

5. Did you ever invest money in the stock market? <radio buttons with “Yes” and “No”>

6. Are you currently invested in the stock market? <radio buttons with “Yes” and “No”>

8 Convexity Score

The main idea of the score is to express the degree of the convexity or concavity of a price path. To achieve this, we first fit a line through a price path that connects the starting point with the end point of the path. This line represents the proportional cumulative return over the depicted time period p assuming that the whole period return is evenly distributed over the whole length of the path and abstracting from compounding effects. This line represents the trend of the overall price path. For every period t the proportional cumulative return from 0 to t is defined as

$$\hat{x}_{0,t} = t \cdot \frac{\prod_{i=1}^{p}(1 + x_i) - 1}{p} \qquad \forall t \in [0, p]$$

We are now interested in the differences between the price path and this line (i.e., the area). We calculate this difference between actual and proportional cumulative return separately for positive and negative differences.

$$\bar{x}^{+}_{0,t} = x_{0,t} - \hat{x}_{0,t} \qquad \forall x_{0,t} \geq \hat{x}_{0,t}$$
$$\bar{x}^{-}_{0,t} = \hat{x}_{0,t} - x_{0,t} \qquad \forall x_{0,t} < \hat{x}_{0,t}$$

For the considered time period we then add up all $\bar{x}^{+}_{0,t}$ as well as all $\bar{x}^{-}_{0,t}$ separately to get a proxy for the areas between actual price path and trend line depending on whether the path lies above the line or below it.

$$H^{+} = \sum_{i=1}^{m} \bar{x}^{+}_{0,i}$$
$$H^{-} = \sum_{i=m+1}^{n} \bar{x}^{-}_{0,i}$$

See Figure 1 for an illustrative example. H^+ is represented by the vertically striped area while H^- is the horizontally striped area.

Figure S9: Convexity score example

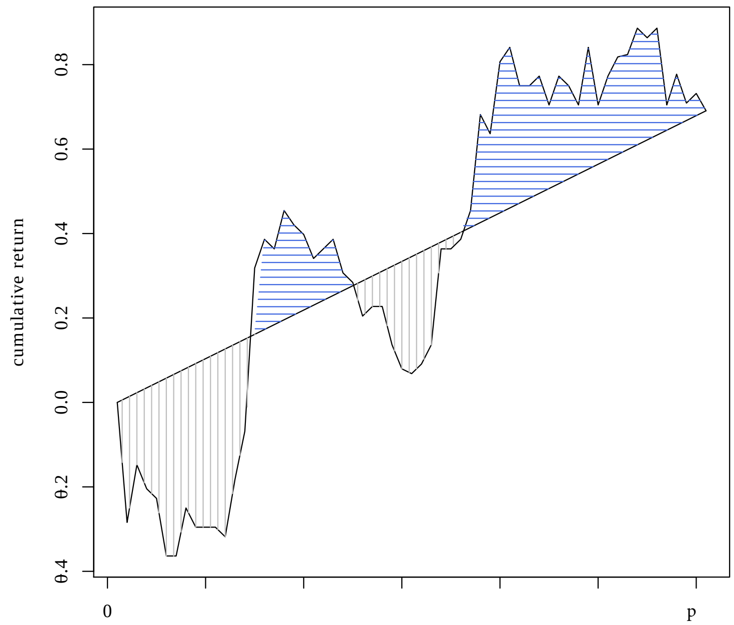

The figure illustrates the construction of our convexity measure.

The score we use now subtracts H^- from H^+ and divides this term by the risk associated with a path. As a risk proxy we use the spread calculated as the difference between the maximum cumulative return and the minimum cumulative return over the entire chart period p. Our convexity score (CS) is thus defined as:

$$CS = \frac{H^+ - H^-}{max_t(x_{0,t}) - min_t(x_{0,t})}(\prod_{i=1}^{p}(1 + x_i) - 1)$$

9 Course of the Experiment (Study III)

Figure S10: Course of the Experiment (Study III)

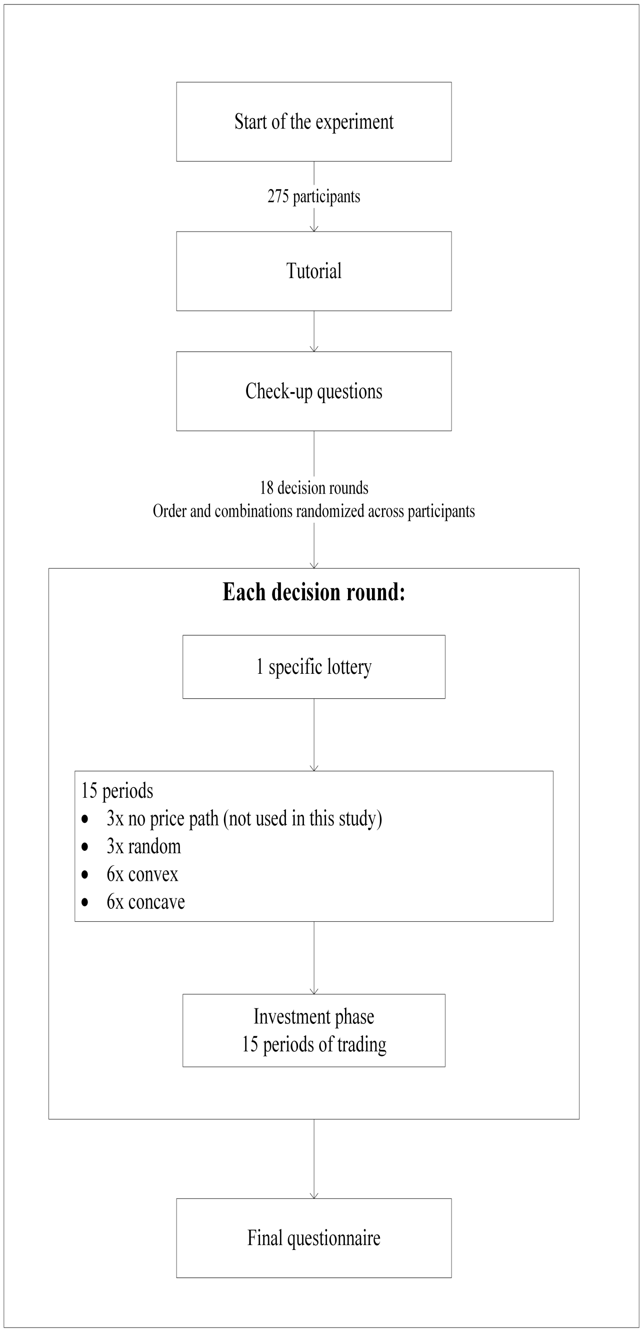

The figure illustrates the course of our lab experiment used for Study III.

10 Instructions for Study III

All paragraphs appeared on screen with appropriate timely delay. Participants were instructed to raise their hands in case they had trouble understanding the instructions. In this case, the experimenter came to assist. Instructions were supported by appropriate graphical displays if necessary (not displayed in this appendix for sakes of brevity).

10.1 Your task in the experiment

You will be playing 18 investment rounds in this experiment. In each investment round you can invest into a stock and a bank account. Each investment round consists of 30 time periods. The first time period is labeled ”period 1” and the last point in time is labeled ”period 30”.

During the first 15 periods (i.e., periods 1-15), you are observing the price development of the stock in these first 15 periods (i.e. periods 1-15). During the last 15 periods (i.e., periods 16-30), you will have to decide how much money you want to invest into this stock and how much money you want to invest into a bank account.

The Stock: For simplicity, we assume that the stock can only generate 2 possible returns. You will be explicitly told about the two possible returns and the probabilities before you will have to make your decisions.

For example: the stock may either generate a positive return of +6 percent with 50 percent probability in a period, or it may generate a return of -4 percent with 50 percent probability in a period. This means that the stock price has a 50 percent chance of going up by +6 percent in the next period, and a 50 percent chance of going down by -4 percent in the next period. However, it does not mean that the stock will always go up and down exactly 15 times each in the 30 periods of the investment round.

The bank account: All money that you don’t invest into the stock will be automatically put into the bank account. In contrast to the stock investment, the money in the bank account will always remain the same, no matter what return the stock generates during the period.

Check-up questions:

1. Which statement about the stock is true:
   - There are only two possible returns that the stock can generate in a period.
   - There is an infinite number of returns that the stock can generate in a period.
   - The stock always generates a safe return.
   - None of the above.
2. Assume that the stock has a 50 percent probability of going up by 6 percent and a 50 percent probability of going down by -4 percent over one period. Which of the following statements is true?
   - The stock price will first go up 15 times and then go down 15 times.
   - The stock price will go up exactly 15 times and down exactly 15 times. The sequence of ”ups” and ”downs” is unknown.
   - You cannot predict the exact number of ”ups” and ”downs” as well as their sequence.
   - None of the above.
3. Which statement about the bank account is true:
   - The money in the bank account always remains the same.
   - The return of the bank account is affected by the return of the stock.
   - There is a 2 percent chance that all money in the bank account will be lost.
   - None of the above.

10.2 How you will be paid

In total, you will be playing 18 investment rounds. After the experiment, you will draw one ping-pong ball out of an urn with 18 balls. The ping-pong balls have the numbers 1 to 18 on them. The number on the ball that you have drawn represents the investment round according to which you will be paid.

For each monetary unit that you own at the end of this investment round, you will receive a payment of 0.001€ after the experiment. In addition, you will receive a show-up fee of 8€ .

Before each investment round, you will solve a ”word scramble”. After having completed the word scramble, you will learn the random amount that will be invested in the stock in period one of the following investment round. Recall that you will be paid according to the amount of money that you own at the end of one investment round (i.e., in period 30). This means: The more money is invested into the stock at the beginning of period 1, the higher is this expected amount of money you own at the end of the investment round.

To sum up: Before each investment round, you will do a ”word scramble”. In periods 16 to 30 of the investment round, you will have to decide how much money you want to invest into a stock and how much money you want to put in a bank account.

The stock will generate one out of two possible returns. You will be told about the possible returns and the probabilities before each investment round. All money in the bank account will always remain the same and is unaffected by the return of the stock.

After the experiment, you will draw one ping-pong ball out of an urn with 18 balls. The number on the ball represents the investment round according to which you will be paid. For each monetary unit that you own at the end of this investment round, you will receive a payment of 0.001€ after the experiment. In addition, you will receive a show-up fee of 8€ .

10.3 Final questionnaire

Cognitive Reflection Test

1. A bat and a ball cost 1.10 in total. The bat costs 1 more than the ball. How much does the ball cost?

2. If it takes five machines five minutes to make five widgets, how long would it take 100 machines to make 100 widgets?

3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

Financial Literacy

1. Suppose you had € 100 in a savings account and the interest rate was 2 percent per year. After 5 years, how much do you think you would have in the account if you left the money to grow? More than € 102 exactly € 102, less than € 102, do not know.

2. Suppose you had 100€ in a savings account and the interest rate is 20 percent per year and you never withdraw money or interest payments. After 5 years, how much would you have on this account in total? More than 200€, exactly 200€, less than 200€, do not know.

3. Imagine that the interest rate on your savings account was 1 percent per year and inflation was 2 percent per year. After 1 year, how much would you be able to buy with the money in this account? More than today, exactly today, less than today, do not know.

4. Assume a friend inherits 10,000€ today and his sibling inherits 10,000€ 3 years from now. Who is richer because of the inheritance? My friend, his sibling, they are equally

rich, do not know.

5. Suppose that in the year 2017, your income has doubled and prices of all goods have doubled too. In 2017, how much will you be able to buy with your income? More than today, the same, less than today, do not know.

Self-stated characteristics (all answered on a 7-point Likert scale)

1. I relied on my gut instinct in my decision-making.
2. I thought long about my decisions.
3. My decisions were well considered.
4. In general I don't mind taking risks.
5. My decisions were influenced by the price paths.
6. The stock price paths from periods 1-15 helped me to form a belief about the future performance of the stock.
7. I perceived the stock to be more risky when there was no price path for periods 1-15.
8. I tried looking for patterns in the presented stock price paths.
9. I believe that it was more likely for the stock price to go down if it went up in the previous period.
10. I believe that it was more likely for the stock price to go up if it went up in the previous period.
11. I had the impression that the price paths were manipulated.

12. To me, it was important to end up with at least as much money as I started with in each investment round.
13. I had trouble understanding the experiment instructions.
14. Taking part in the experiment was fun.
15. I had trouble solving the Word Scrambles.
16. In general, I trust other people.
17. When I made a decision that turned out to be bad, I usually regret that decision for a long time.
18. I would consider myself to be rather patient.

11 Additional Analyses Study I and II

11.1 Descriptives for the MTurk Study

Table 4: Descriptive statistics

Variable	US-Population (2017) $N = 321,004,407$	MTurk Sample $N = 500$
Age [years; median]	37.2	30.0
Gender [female=1]	50.2	32.2
Education [%]		
No degree	12.6	0.2
High School	27.3	23.4
College incl. BA	48.2	64.2
Graduate or higher	11.8	12.2
Full employment [%]	77.2	85.6
Household size [mean]	2.58	3.08

The table reports measures of central tendency of selected demographics for the US population and for our sample. Data for the US population is based on the American Community Survey (ACS) and obtained from the website of the United States Census Bureau.

11.2 Robustness Analyses for Study I

Table 5: Regressions for IA, Study I: Robustness I
(excluding 29.8% of participants who did not pass attention or comprehension filter)

	(1)	(2)	(3)	(4)	(5)
	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]
Corr(Return, VS weights)	0.857***			0.812***	0.975***
	(0.303)			(0.306)	(0.341)
Corr(Return, CPT weights)		0.369		0.984*	1.389*
		(0.312)		(0.562)	(0.732)
Corr(Return, Salience weights)			-0.00889	-0.228	-0.355*
			(0.0883)	(0.156)	(0.199)
Return	0.385***	0.381***	0.383***	0.379***	0.509***
	(0.0267)	(0.0262)	(0.0264)	(0.0266)	(0.0221)
Stdv	-0.165***	-0.174***	-0.168***	-0.175***	-0.251***
	(0.0194)	(0.0203)	(0.0193)	(0.0212)	(0.0217)
Skewness	0.00401	-0.0301	0.00827	-0.0401	-0.0541
	(0.0143)	(0.0343)	(0.0255)	(0.0355)	(0.0487)
Constant	-0.151***	-0.177***	-0.160***	-0.212***	-0.320***
	(0.00279)	(0.0153)	(0.00467)	(0.0351)	(0.0613)
Observations	2808	2808	2808	2808	2808
R^2	0.221	0.218	0.217	0.221	

* p < 0.1, ** p <0.05, *** p<0.01

The table reports invested amounts on the correlation between current return, visual salience decision weight, CPT decision weight, and the salience decision weight. The correlation serves to isolate the effect of different decision weight specifications from the different assumptions about the utility function. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. The dependent variable is coded in percentage values to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

Table 6: Regressions for IA, Study I: Robustness II
(excluding 27.4% of participants with large discrepancy between IA and attractiveness score)

	(1)	(2)	(3)	(4)	(5)
	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]
Corr(Return, VS weights)	0.726***			0.686**	0.814**
	(0.273)			(0.276)	(0.318)
Corr(Return, CPT weights)		0.503*		1.105**	1.238*
		(0.290)		(0.536)	(0.684)
Corr(Return, Salience weights)			0.0363	-0.215	-0.271
			(0.0804)	(0.147)	(0.185)
Return	0.371***	0.366***	0.369***	0.364***	0.492***
	(0.0245)	(0.0243)	(0.0244)	(0.0246)	(0.0208)
Stdv	-0.151***	-0.161***	-0.153***	-0.163***	-0.230***
	(0.0179)	(0.0185)	(0.0178)	(0.0192)	(0.0203)
Skewness	-0.00472	-0.0509*	-0.0109	-0.0621**	-0.0701
	(0.0137)	(0.0296)	(0.0220)	(0.0310)	(0.0437)
Constant	0.156***	0.126***	0.153***	0.0891***	-0.0886
	(0.00271)	(0.0144)	(0.00475)	(0.0339)	(0.0582)
Observations	2906	2906	2906	2906	2906
R^2	0.213	0.211	0.210	0.214	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The table reports invested amounts on the correlation between current return, visual salience decision weight, CPT decision weight, and the salience decision weight. The correlation serves to isolate the effect of different decision weight specifications from the different assumptions about the utility function. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. The dependent variable is coded in percentage values to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

12 Additional Analyses Study III

12.1 Descriptives for the Lab Experiment

Table 7: Descriptive statistics

Variable	All $N = 275$	SAVE $N = 230$
Age [years]	20.47	61.21
	(2.21)	(14.85)
Gender [female=1]	0.29	0.49
	(0.46)	(0.50)
Degree [graduate=1]	0.10	0.19
	(0.30)	(0.40)
Self-stated risk aversion [1,7]	3.83	5.08
	(1.75)	(1.89)
Financial Literacy score [%]	80.28	80.69
	(20.50)	(24.09)
CRT score [%]	54.00	39.24
	(37.60)	(35.70)

The table reports the means of selected demographics and personality traits. Standard deviations are reported in parentheses. ”CRT” indicates the Cognitive Reflection Test by Frederick (2005). The first column represent mean data from our lab study, the last column for stock owners in Germany (data is taken from the German SAVE study in 2009; Börsch-Supan et al., 2008). The SAVE study is representative for the German population aged 18 and above.

12.2 Cumulative Density Functions for Parameter Sensitivity Analysis

Figure S11: Parameter sensitivity analysis - CDF

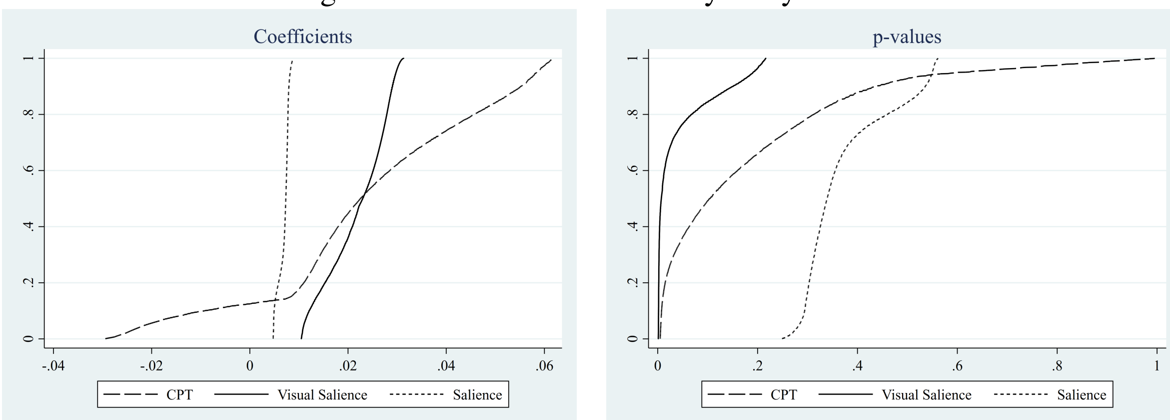

The graph on the left (right) shows the three models' cumulative distribution function of coefficients (p-values) for all analyzed α-λ combinations.

13 Price differences instead of returns

In our three studies, participants look at price path images before making investment choices. It is plausible that participants perceive price differences (change in levels), rather than returns, when looking at these images. We account for this possibility by conducting a robustness test to check how our analyses in the main paper are affected when considering price differences as the basis of analysis, as opposed to returns.

In the following sections, we repeat the analyses in the main paper for all three studies using price differences. Results are qualitatively similar, with the Visual Salience model performing best in all our analyses. The main changes from the analyses using returns as the basis are:

- In Study I (empirical price paths), the CPT model (Barberis et al., 2016) and the Salience model (Bordalo et al., 2013) perform better (using the CPT value function) when predicting investments. We note that both these models are usually defined on returns in the finance literature, so this has interesting potential for further exploration.
- In Study III (laboratory subjects), the the CPT model (Barberis et al., 2016) performs marginally worse when predicting investments, compared to using returns as the basis for analysis.

13.1 Study I

Table 8: Regressions for IA, Study I: Correlation Measure

	(1) IA [%]	(2) IA [%]	(3) IA [%]	(4) IA [%]	(5) IA[%]
Corr(PriceDiff, VS weights)	0.579**			0.537**	0.605**
	(0.236)			(0.236)	(0.243)
Corr(PriceDiff, CPT weights)		0.296		0.760*	0.895*
		(0.239)		(0.420)	(0.467)
Corr(PriceDiff, Salience weights)			-0.00832	-0.195	-0.246*
			(0.0688)	(0.120)	(0.140)
Return	0.295***	0.295***	0.295***	0.292***	0.368***
	(0.0205)	(0.0203)	(0.0205)	(0.0204)	(0.0158)
Stdv	-0.130***	-0.132***	-0.132***	-0.128***	-0.170***
	(0.0152)	(0.0153)	(0.0152)	(0.0151)	(0.0147)
Skewness	0.00902	-0.0190	0.0116	-0.0203	-0.0239
	(0.0112)	(0.0259)	(0.0194)	(0.0267)	(0.0317)
Constant	0.00491**	-0.00842	-0.000600	-0.0311	-0.0881**
	(0.00200)	(0.00678)	(0.00496)	(0.0197)	(0.0421)
Observations	4000	4000	4000	4000	4000
R^2	0.162	0.160	0.160	0.163	

* p < 0.1, ** p <0.05, *** p<0.01

The table reports invested amounts on the correlation between current price difference, visual salience decision weight, CPT decision weight, and the salience decision weight. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. The dependent variable is coded in percentage values to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

Table 9: Regressions for IA, Study I: CPT Value Function

	(1) IA [%]	(2) IA [%]	(3) IA [%]	(4) IA [%]	(5) IA [%]
Visual Salience	0.286***			0.211***	0.238***
	(0.0318)			(0.0347)	(0.0412)
CPT		0.281***		0.122**	0.177***
		(0.0386)		(0.0482)	(0.0497)
Salience			0.143***	0.0559*	0.0217
			(0.0269)	(0.0308)	(0.0346)
Return	0.261***	0.311***	0.415***	0.324***	0.369***
	(0.0202)	(0.0201)	(0.0343)	(0.0345)	(0.0321)
Stdv	0.133***	0.124***	-0.0814***	0.194***	0.216***
	(0.0334)	(0.0384)	(0.0168)	(0.0424)	(0.0404)
Skewness	0.0115	-0.0820***	-0.0389**	-0.0476***	-0.0565***
	(0.0112)	(0.0161)	(0.0152)	(0.0174)	(0.0216)
Constant	1.36e-08***	1.30e-08***	1.33e-08***	1.34e-08***	-0.0496
	(4.25e-11)	(2.96e-11)	(1.44e-11)	(7.32e-11)	(0.0355)
Observations	4000	4000	4000	4000	4000
R^2	0.183	0.179	0.167	0.188	

* p < 0.1, ** p <0.05, *** p<0.01

The table reports invested amounts on the Visual Salience model, the CPT model (Barberis et al., 2016) and the Salience model (Bordalo et al., 2012) where we assume gain-loss utility with curvature $\alpha = 0.88$, a reference point ρ of zero, and a loss aversion coefficient of $\lambda = 2.25$, and price differences instead of returns. Control variables are the return of the price path (Return), its standard deviation (Stdv), and its skewness (Skewness). In Model (5) we use a Tobit regression to account accumulation of participants who did not invest in that stock. We use fixed effects to control for individual specific effects. The dependent variable is coded in percentage values to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

13.2 Study II

Table 10: Regressions for IA, Study II: Recency Effect

	(1)	(2)	(3)	(4)	(5)
	IA[%]	IA [%]	IA [%]	IA [%]	IA [%]
Visual Salience	0.192*				0.177*
	(0.105)				(0.103)
CPT (0.95)		-0.00534			
		(0.0455)			
CPT (0.85)			0.0394		
			(0.0432)		
CPT (0.50)				0.137**	0.127*
				(0.0690)	(0.0678)
6, 20	0.0426	0.0585	0.0573	0.0773	0.0621
	(0.109)	(0.110)	(0.108)	(0.105)	(0.106)
6, 20, skew	0.114	0.136	0.120	0.129	0.112
	(0.101)	(0.110)	(0.108)	(0.0993)	(0.0998)
Constant	0.432	-0.0653	-0.0571	-0.0428	0.413
	(0.289)	(0.0667)	(0.0669)	(0.0665)	(0.282)
Observations	600	600	600	600	600
R^2	0.026	0.011	0.014	0.028	0.041

Standard errors in parentheses
* p<0.1, ** p<0.05, *** p<0.01

The table reports invested amounts on the Visual Salience Model with CPT value function (Model(1)) and a CPT model with recency parameter $\rho = 0.95, 0.85, 0.50$, Models (2)-(4). In Model (5) we run the regression with the Visual Salience value and the CPT value ($\rho = 0.5$) as explanatory variables. We also include controls for the parameter constellation of the constructed price charts, with $\mu = -6\%$ and $\sigma = 20\%$ as reference category. All models are based on price differences instead of returns.

13.3 Study III

Table 11: Regressions for IA, Study III: CPT Value Function

	(1)	(2)	(3)	(4)	(5)	(6)
	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]	IA [%]
Visual salience	0.0213**			0.0226***	0.0262***	0.0430***
	(0.00877)			(0.00846)	(0.00808)	(0.0113)
CPT		0.00926		-0.00253	-0.0121	-0.0255**
		(0.0113)		(0.0120)	(0.0145)	(0.0122)
Salience			0.00240	0.00277	-0.00464	-0.00601
			(0.00754)	(0.00769)	(0.00844)	(0.00826)
Expected return	0.0843	0.0786	0.0796	0.0845	0.0914	0.132**
	(0.0793)	(0.0795)	(0.0793)	(0.0794)	(0.0799)	(0.0600)
Expected std. dev.	-0.0147	-0.0179	-0.0241*	-0.0155	-0.0708	-0.104***
	(0.0136)	(0.0154)	(0.0129)	(0.0155)	(0.0466)	(0.0292)
Expected skewness	-0.248***	-0.255***	-0.245***	-0.247***	-0.228***	-0.372***
	(0.0411)	(0.0434)	(0.0409)	(0.0430)	(0.0473)	(0.0311)
Average realized return	-0.300***	-0.269***	-0.242***	-0.291***	-0.0388	-0.0943
	(0.0376)	(0.0401)	(0.0362)	(0.0421)	(0.190)	(0.126)
1/n Decision Weights					-0.0929	-0.129***
					(0.0687)	(0.0443)
Constant	0.128***	0.140***	0.156***	0.129**	0.297**	0.646***
	(0.0478)	(0.0523)	(0.0469)	(0.0521)	(0.141)	(0.120)
Observations	61875	61875	61875	61875	61875	61875
N	275	275	275	275	275	275
R^2	0.024	0.024	0.023	0.024	0.024	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The table reports regressions of invested amounts on the Visual Salience model, the CPT model (Barberis et al., 2016) and the Salience model (Bordalo et al., 2012) respectively (while using price differences instead of returns). Control variables are expected return of the binary lottery, its standard deviation, its skewness, and the average realized return. In Model (5) we also include a CPT value function with equal (1/n) decision weights to control for the effect of value function specification. In Model (6) we use a Tobit regression to account for accumulation of participants who did not invest in that stock. We use fixed effects to control for individual characteristics of participants. All variables are standardized for easy comparison of coefficient sizes and to enhance readability of the table. Robust standard errors clustered on participant level are reported in parentheses.

References

Bali, T. G., Cakici, N. and Whitelaw, R. F. (2011), Maxing out: Stocks as lotteries and the cross-section of expected returns, *Journal of Financial Economics* **99**(2), 427–446.

Barber, B. M. and Odean, T. (2007), All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *The Review of Financial Studies* **21**(2), 785–818.

Barberis, N., Huang, M. and Santos, T. (2001), Prospect theory and asset prices, *Quarterly Journal of Economics* **116**(1), 1–53.

Barberis, N., Mukherjee, A. and Wang, B. (2016), Prospect theory and stock returns: An empirical test, *The Review of Financial Studies* **29**(11), 3068–3107.

Baucells, M., Weber, M. and Welfens, F. (2011), Reference-point formation and updating, *Management Science* **57**(3), 506–519.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2012), Salience theory of choice under risk, *Quarterly Journal of Economics* **127**(3), 1243–1285.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2013), Salience and asset prices, *American Economic Review* **103**(3), 623–28.

Borsboom, C. and Zeisberger, S. (2020), What makes an investment risky? An analysis of price path characteristics, *Journal of Economic Behavior & Organization* **169**, 92–125.

Börsch-Supan, A., Coppola, M., Essig, L., Eymann, A. and Schunk, D. (2008), The German SAVE study: Design and results, *Mannheim Research Institute for the Economics of Aging (MEA)*.

Busse, J. A. and Green, T. C. (2002), Market efficiency in real time, *Journal of Financial Economics* **65**(3), 415–437.

Chen, J., Hong, H. and Stein, J. C. (2001), Forecasting crashes: Trading volume, past returns, and conditional skewness in stock prices, *Journal of Financial Economics* **61**(3), 345–381.

Conrad, J., Kapadia, N. and Xing, Y. (2014), Death and jackpot: Why do individual investors hold overpriced stocks?, *Journal of Financial Economics* **113**(3), 455–475.

Cornia, M., Baraldi, L., Serra, G. and Cucchiara, R. (2018), Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model, *IEEE Transactions on Image Processing* **27**(10), 5142–5154.

Cosemans, M. and Frehen, R. (2017), Salience theory and stock prices: Empirical evidence, *Working Paper*.

Czaczkes, B. and Ganzach, Y. (1996), The natural selection of prediction heuristics: Anchoring and adjustment versus representativeness, *Journal of Behavioral Decision Making* **9**(2), 125–139.

Da, Z., Engelberg, J. and Gao, P. (2011), In search of attention, *The Journal of Finance* **66**(5), 1461–1499.

Diacon, S. and Hasseldine, J. (2007), Framing effects and risk perception: The effect of prior performance presentation format on investment fund choice, *Journal of Economic Psychology* **28**(1), 31–52.

Fang, L. and Peress, J. (2009), Media coverage and the cross-section of stock returns, *The Journal of Finance* **64**(5), 2023–2052.

Frederick, S. (2005), Cognitive reflection and decision making, *Journal of Economic Perspectives* **19**(4), 25–42.

George, T. J. and Hwang, C.-Y. (2004), The 52-week high and momentum investing, *The Journal of Finance* **59**(5), 2145–2176.

Glaser, M., Iliewa, Z. and Weber, M. (2019), Thinking about prices versus thinking about returns in financial markets, *The Journal of Finance* **74**(6), 2997–3039.

Glaser, M., Langer, T., Reynders, J. and Weber, M. (2007), Framing effects in stock market forecasts: The difference between asking for prices and asking for returns, *Review of Finance* **11**(2), 325–357.

Gödker, K. and Lukas, M. (2017), All that glitters is not gold: Asymmetric investor attention in the stock market, *Available at SSRN 3080332*.

Grosshans, D. and Zeisberger, S. (2018), All's well that ends well? On the importance of how returns are achieved, *Journal of Banking and Finance* **87**, 397–410.

Huber, C. and Huber, J. (2018), Scale matters: Risk perception, return expectations, and investment propensity under different scalings, *Experimental Economics* pp. 1–25.

Jegadeesh, N. and Titman, S. (1993), Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* **48**(1), 65–91.

Kahneman, D. (1979), Prospect theory: An analysis of decisions under risk, *Econometrica* **47**, 278.

Koester, A., Lundholm, R. and Soliman, M. (2016), Attracting attention in a limited attention world: exploring the causes and consequences of extreme positive earnings surprises, *Management Science* **62**(10), 2871–2896.

Kumar, A. (2009), Who gambles in the stock market?, *The Journal of Finance* **64**(4), 1889–1933.

Kumar, A., Ruenzi, S. and Ungeheuer, M. (2019), Daily winners and losers, *Available at SSRN 2931545*.

Li, X. and Camerer, C. (2019), Using Visual Salience in Empirical Game Theory, *Available at SSRN 3308886*.

Mizrach, B. and Weerts, S. (2009), Highs and lows: a behavioural and technical analysis, *Applied Financial Economics* **19**(10), 767–777.

Nolte, S. and Schneider, J. C. (2018), How price path characteristics shape investment behavior, *Journal of Economic Behavior and Organization* **154**, 33–59.

Odean, T. (1999), Do investors trade too much?, *American Economic Review* **89**(5), 1279–1298.

Raghubir, P. and Das, S. R. (2010), The long and short of it: why are stocks with shorter runs preferred?, *Journal of Consumer Research* **36**(6), 964–982.

Stambaugh, R. F., Yu, J. and Yuan, Y. (2015), Arbitrage asymmetry and the idiosyncratic volatility puzzle, *The Journal of Finance* **70**(5), 1903–1948.

Towal, R. B., Mormann, M. and Koch, C. (2013), Simultaneous modeling of visual saliency and value computation improves predictions of economic choice, *Proceedings of the National Academy of Sciences* **110**(40), E3858–E3867.

Tversky, A. and Fox, C. R. (1995), Weighing risk and uncertainty., *Psychological Review* **102**(2), 269.

Tversky, A. and Kahneman, D. (1992), Cumulative prospect theory: An analysis of decision under uncertainty, *Journal of Risk and Uncertainty* **5**(4), 297–323.

Ungeheuer, M. (2017), Stock returns and the cross-section of investor attention, *Available at SSRN 2931547*.

Weber, E. U., Siebenmorgen, N. and Weber, M. (2005), Communicating asset risk: How name recognition and the format of historic volatility information affect risk perception and investment decisions, *Risk Analysis: An International Journal* **25**(3), 597–609.

Wu, G. and Gonzalez, R. (1996), Curvature of the probability weighting function, *Management Science* **42**(12), 1676–1690.