A Valence Asymmetry in Pre-decisional Distortion of Information: Evidence From an Eye Tracking Study with Incentivized Choices

DOI:
[10.1037%2Fxm0000711](https://doi.org/10.1037%2Fxm0000711)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Krol, M., & Krol, M. (2019). A Valence Asymmetry in Pre-decisional Distortion of Information: Evidence From an Eye Tracking Study with Incentivized Choices. *Journal of experimental psychology. Learning, memory, and cognition*. <https://doi.org/10.1037%2Fxm0000711>

Published in:

Journal of experimental psychology. Learning, memory, and cognition

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

A Valence Asymmetry in Pre-decisional Distortion of Information:
Evidence From an Eye Tracking Study with Incentivized Choices

Michał Król¹ and Magdalena Ewa Król²

¹ Department of Economics, School of Social Sciences, University of Manchester, Manchester,
United Kingdom

² Wrocław Faculty of Psychology, SWPS University of Social Sciences and Humanities in
Wrocław, Wrocław, Poland

* Corresponding author:

Email: michal.krol@manchester.ac.uk, The University of Manchester, Manchester, UK

OSF pre-registration link:

https://osf.io/t8mbc/?view_only=4a36eb6e609c4728a631f7be949d7d2a

This work was supported by the Polish Ministry of Science and Higher Education grant number
BST/WROC/2016/A/6.

Abstract

Existing research shows that the order in which evidence arrives can bias its evaluation and the resulting decision in favor of information encountered early on. We used eye-tracking to study the underlying cognitive mechanisms in the context of incentivized financial choices based on real world market data. Subjects learned about the presence/absence of a transaction fee, before seeing expert opinions regarding an investment prospect and deciding whether to invest. Although the fee had no effect on the processing of negative opinions, we found that positive ones were processed more effortlessly (with lower gaze duration and pupil dilation) when it was absent, i.e. when they were congruent with the positive initial information in the shape of the lack of fees. Despite their more effortless processing in the absence of fees, positive opinions then had a greater impact on the subjects' beliefs. In addition to an initial study with N=100 subjects, these findings were replicated in a second, pre-registered experiment with N=103 subjects, in which a positive premium was paid in the event of no fee. Thus, we argue that the valence asymmetry in favor of positive information observed in evaluative priming, person perception, and related tasks (the 'density hypothesis') also plays a crucial role in incentivized economic choice. In fact, rather than being a detrimental bias, the overweighting of initial evidence often observed in decisions could be seen as an adaptive heuristic aimed at reducing the cost of processing later, similar information.

Keywords

valence asymmetry; density hypothesis; pre-decisional information distortion; eye-tracking; economic decision-making

As shown by extensive research on pre-decisional information distortion, the order in which information arrives can bias its processing, whereby encountering early evidence supporting a particular choice option shifts the interpretation of subsequent, ambiguous evidence in its favor (see DeKay, 2015, for an overview). It was suggested that the distortion is driven by maximizing the consistency between old and new information (J Edward Russo, Carlson, Meloy, & Yong, 2008; Simon, Pham, Le, & Holyoak, 2001). However, little attempt has been made to study the underlying cognitive mechanisms via process-tracing techniques, despite existing research suggesting that such an analysis might be fruitful. In particular, there is substantial evidence of a positive feedback loop between eye-movements and preferences (Shimojo, Simion, Shimojo, & Scheier, 2003), of gaze patterns consistent with bidirectional links between information and choice options (Glöckner & Herbold, 2011), and of considerable predictive power of attentional evidence accumulation models that allow for the ‘primacy’ (overweighting) of early information (Ashby, Jekel, Dickert, & Glöckner, 2016).

What is more, existing studies in domains related to, but not strictly within the decision-making domain, suggest that what might explain the pre-decisional distortion is the fact that early evidence induces an initial sentiment in people which hinders the processing and interpretation of subsequent data incongruent with that sentiment, while facilitating the processing of congruent information. For example, research on ‘epistemic Stroop effects’ (Gilead, Sela, & Maril, 2018; Richter, Schroeder, & Wöhrmann, 2009) demonstrates that people take longer to give a positive answer to a question about a piece of textual information when having a negative rather than positive pre-existing sentiment towards it. For instance, they take longer to confirm that a sentence ‘Internet makes you lonely.’ is grammatically correct when

they disagree with this statement than when they agree. Interestingly, however, the converse is not observed, i.e. a negative initial sentiment does not result in faster negative answers.

This kind of positive-negative valence asymmetry is in line with the ‘density hypothesis’ (Unkelbach, Fiedler, Bayer, Stegmüller, & Danner, 2008), which posits that positive information is, in general, relatively similar to other positive information (that is, ‘densely packed’ in the brain, hence the name of the hypothesis). In contrast, negative information tends to be considerably less similar to other negative information. It has been argued that this similarity asymmetry is a robust and general characteristic of the environment humans live in (Koch, Alves, Krüger, & Unkelbach, 2016). As a result, early exposure to positive information facilitates the processing of subsequent positive data, as the two are readily linked together, as opposed to different pieces of negative information. This is observed not only in evaluative priming, which is stronger for positive than for negative primes, but also in the perception of other people. For instance, ‘halo effects’ are stronger for positive than for negative traits (Gräf & Unkelbach, 2016), i.e. ‘being honest makes you industrious (in others’ eyes), but lying does not make you lazy’.

However, despite the apparent potential of the density hypothesis to explain a wide range of phenomena, so far it has not been examined in a setting in which decisions have real economic consequences for the study participants (see Alves, Koch, & Unkelbach, 2017b for a recent overview of the scope of the related literature). In contrast, in this paper, we used eye-tracking to study the cognitive processes underlying the pre-decisional information distortion, with a particular focus on the interplay between positive vs. negative pieces of evidence.

To this end, we conducted an experiment modeled on a real-world scenario in which financial investors first access readily accessible, easy-to-understand data on stocks’ past returns

and potential transaction fees, before reading more nuanced expert opinions about the considered investments. Such a sequence of information processing is naturally imposed by most web portals for investors (e.g., seekingalpha.com), where browsing for a stock brings up numerical summary information before one can click through to access relevant articles. Importantly, in this type of choices there is considerable evidence of behavioral bias consistent with a pre-decisional distortion (e.g. Chang, Solomon, & Westerfield, 2016; Frazzini, 2006; Frydman, Barberis, Camerer, Bossaerts, & Rangel, 2014; Park et al., 2013), while an asymmetry in learning from positive vs. negative information has also been reported (Kuhnen, 2015). However, the vast majority of behavioral finance studies provide subjects with numerical data alone, despite the fact that the role of textual information in financial markets is increasingly recognized (e.g. Da, Engelberg, & Gao, 2015; Gerard, Gordon, & Nagpurnanand, 2013; Hendershott, Livdan, & Schürhoff, 2015; Manela & Moreira, 2017). In contrast, we used a mixture of numerical and textual data that is more likely to be encountered in complex real-world choices, and which allowed us to study the way in which people interpret the often ambiguous textual information.

More specifically, in our first experiment, we presented one hundred student subjects with investment opportunities based on real-world stock market data. They were first presented with information about the historical return of a randomly chosen stock in a randomly chosen past period, as well as on whether a transaction fee is payable in the event of choosing to invest. Next, subjects would see a word cloud of expert opinions about the stock sourced from seekingalpha.com, a leading crowd-sourced content service for investors, before choosing if they want to invest, in which case they would accumulate real monetary rewards according to the stock's return in the subsequent period. As shown by existing research (Chen, De, Hu, & Hwang,

2014), aggregating data from several seekingalpha.com articles and evaluating the sentiment of individual words included therein can predict the subsequent stock returns. This motivated our use of this data to construct experimental stimuli that subjects would find credible.

Crucially, but unknown to subjects, each investment opportunity was shown twice over several trials of the study, once with and once without the transaction fee. We hypothesized that the presence/absence of the fee, by inducing an initial negative/positive sentiment towards investment, would affect the processing of positive vs. negative words (defined as per Loughran & McDonald, 2011). Specifically, we expected positive words to be easier to process and interpret in those trials in which they were congruent with the initial positive information in the form of the absence of the fee. This should manifest in decreased measures of mental effort in the gathered eye-data, but an increased influence of positive words on subjects' beliefs, which we elicit via an innovative paradigm based on anticipatory eye-movements (Santos & Kowler, 2017). At the same time, the density hypothesis would suggest that in case of negative information the analogous effect might be weaker or non-existent, i.e. that a negative initial information in the form of the presence of the fee would not facilitate the processing of negative information to the same extent, because different negative pieces of evidence (discouraging investment) are not as readily associated with each other as positive ones.

Accordingly, our analysis plan was split into two parts. First, we tested the overall 'congruency effect' of the transaction fee, namely that positive words should be processed faster relative to negative ones (in the sense of shorter eye fixation durations) when the former are congruent and the latter incongruent with the positive early information in the form of the absence of the fee, rather than when the fee is present, making negative words congruent and the positive ones incongruent. Second, we decomposed the overall congruency effect of the fee on

fixation durations, testing it separately for positive and negative words, and expecting to find it in case of the former but not the latter, as suggested by the density hypothesis. Additionally, although existing research focused on the effect of congruency on the speed of processing, we conducted an exploratory analysis to see whether or not our findings in terms of fixation duration might be supported by pupil dilation measurement, a well-known alternative indicator of mental effort (Beatty, 1982). Similarly, our exploratory analysis of subjects' beliefs, inferred via anticipatory eye-movements, was designed to investigate if the absence of fees would make the decision-makers more sensitive to subsequent positive information, in the sense that the proportion of positive words in the word cloud would have a stronger positive impact on their inferred optimism about the subsequent return on the considered investment.

To further strengthen our findings, we also conducted a pre-registered replication of the initial study, in which a positive premium was paid in the absence of the fee, in order to ensure that such an event is indeed interpreted as 'positive' by our subjects, and that the congruency effect of the fee, as well as the valence asymmetry in this respect, still holds in those circumstances. A robust demonstration of this effect would imply that the pre-decisional distortion of information is driven by the early information facilitating the processing of subsequent congruent evidence, but that this process depends on the similarity between old and new information. Thus, the pre-decisional 'distortion' could, in fact, be viewed as an adaptive heuristic reducing information processing costs, rather than a detrimental decision bias.

Experiment 1

Method

Subjects. We recruited 106 students (mean age 27.9, 62 females) with normal or corrected-to-normal eyesight at a large private university. Six subjects were excluded due to poor eye-tracking calibration or data quality (no eye fixations registered in more than 50% of choice trials).

Stimuli and Design. We used a custom-built Wolfram Mathematica script to scrape and process 15337 ‘single-ticker’ expert opinion articles published on seekingalpha.com (SA) between January 2014 and October 2017 on the 20 largest S&P500 stocks. Such articles explain whether a particular stock should be invested in and why.

For each stock and each monthly period within the overall timespan, we collected articles on that stock from this period and extracted from them words classed as positive/negative according to the Loughran and McDonald (2011) financial sentiment lexicon, which eliminated words identifying the stock (e.g. ‘iPad’). As a significant majority of words in the lexicon are negative, we also included words that were not included there but were classed as positive according to the alternative and widely used Harvard Psychosociological Dictionary (Harvard-IV-4). This ensured that the proportions of positive vs. negative words were on average approximately equal across all word clouds shown to subjects (see below).

As shown by existing research (Chen et al., 2014), the overall proportion of negative words in SA articles published about a stock in the past can predict its return in the subsequent trimester. More specifically, future abnormal returns (net of average market returns) were found to be 0.379% lower when the fraction of negative words was 1% higher. Here, our aim was not to predict returns, but to give subjects a sample of textual evidence that they might consider

181 useful for making such a prediction by themselves. Due to the practical requirements of an eye-
182 tracking analysis, we wished to present subjects with relatively condensed stimuli, thus exposing
183 them to several pieces of relevant information within a short time-span of a single decision trial.
184 Accordingly, from each set of positive/negative words (extracted from SA articles about a given
185 stock published in a given month), we selected the most representative 50 words according to the
186 ‘term frequency-inverse document frequency’ metric, commonly used by internet search engines,
187 whereby a word is ranked high if it appears often in a text sample relative to its frequency in the
188 whole corpus of data (in our case, all the SA articles we scraped). We matched the resulting set
189 of words to actual returns of the stock in NYSE in the previous and subsequent trimesters. For
190 instance, the set of sentiment words in March 2017 was matched to the returns in the first and
191 second trimesters of 2017.

192 In each of the 80 trials of the study, each subject was offered an investment opportunity
193 drawn from the above set, i.e. was shown the *previous* return and 50 representative expert
194 opinion words corresponding to some stock during a certain time period. The returns were shown
195 as whole numbers (‘points’), each percentage point converted to 10 points. Additionally, the
196 subject was told if a transaction fee of 20 points must be paid on investment. If so, then a
197 decision to invest resulted in getting the point-equivalent of the return of the stock in the
198 *subsequent* trimester, minus the fee (otherwise, no fee was paid). A decision not to invest yielded
199 a fixed one-point reward, representing risk-free return, and deliberately set at a very low level to
200 represent the fact that interest rates on secure deposits in world’s largest economies have been
201 close to zero in recent years.

202 Subjects began with 1000 points and were paid an equivalent of 3 USD per 1000 points
203 accumulated on completion. The average payoff was 7 USD (subjects also received university

course credits), and the study took around 25 minutes. We randomly drew the set of 80 investment opportunities for each subject, ensuring that the average previous/subsequent returns and the proportion of positive words across all trials were within 0.1 SD of their averages for the whole set of seekingalpha.com data, i.e. all subjects received broadly similar opportunities representative of the whole set of acquired data.

Crucially, we also ensured that each investment offered to a subject appeared twice over the 80 trials, once with and once without the transaction fee, where the fee appeared in the earlier/later of the two matched trials in exactly half of the trial-pairs, and at least 30 other trials separated every two matched trials. While the repetition was unknown to subjects, we carefully explained to them that the fee is drawn randomly, independently of expert opinions or returns. Finally, we ensured that each subject had a chance to invest in each of the 20 stocks in our dataset 4 times, with no overlap between the involved three-monthly periods across non-matched trials.

Subjects learned the previous return and the fee, before seeing a cloud of sentiment words and deciding to invest or not, moving to subsequent screens by pressing a key (Figure 1). Compared with word clouds that subjects will have seen in day-to-day life, ours was standardized to eliminate factors such as font size, color, or orientation that might have added noise to the eye-tracking data. Specifically, the 50 words were all printed in the same font and randomly arranged in a fixed-sized ellipse ('cloud'), the height/width of which was approximately 80% of the screen.

We used numerical optimization to distribute the words in a way that minimized the variance of the distances between adjacent words, i.e. to ensure that they were approximately evenly distributed. On average, the distance between adjacent words was greater than in typically

seen word clouds, so as to allow for a reliable identification of the exact word a subject is looking at.

[Figure 1 here]

Following the decision, the subsequent return was revealed in a way that enabled inferring the subject's expectations by studying their anticipatory eye-movements. Specifically, we first displayed a horizontal axis, and 800 ms afterward a collection of characters above it, where the position of the only character that was not upside down indicated the return (see Figure 2). The reason for having all but one characters upside down, rather than the other way round, was that this made the task of inferring the return harder for subjects. This, in turn, motivated them to focus their search efforts on those sections of the axis where the correct character was most likely, in their view, to occur.

[Figure 2 here]

Procedure. The stimulus presentation software was programmed in Wolfram Mathematica. Each subject was seated at a laptop with a 15.4-inch, 1280x720px screen, with an SMI-RED250 eye-tracker attached underneath, set to 250Hz frequency. We conducted a five-point semi-automatic calibration and validation with maximum allowed deviation 0.5°. A headrest ensured a distance between the subject's eyes and the device of approximately 70cm. We used a luxometer to check that light intensity was equal across experimental sessions (all conducted in the same lab location without natural light). The study was approved by the local faculty research ethics committee. All words used in the study were translated from English into the local language by a professional translator, and we verified that this preserved the original

word sentiment by asking 50 subjects in an online pilot survey to classify the individual translated words as positive or negative.

Results

Manipulation checks. In the first instance, we wanted to check if subjects understood the task and if the various parameters of the decision problem had the desired effect. To this end, we estimated a mixed-effects binary logistic regression model with the investment decision as the dependent variable (1 = 'invest'), and random subject intercept and slope effects to allow for different observations of the same subject being correlated.

The model estimates in Table 1 indicate that the tendency to invest more with experience was insignificant, i.e. there was no significant relationship between the likelihood of investing and the number of the trial ($\beta_{n\text{-trial}}=0.165, p=.338$)¹, despite the positive average return from investment (+37). Similarly, the subjects were not more likely to invest when having already seen the investment yield positive subsequent returns in a previous matching trial ($\beta_{\text{seen-positive}}=-0.067, p=.540$); nor were they less likely to invest if they have seen the opportunity yield a negative return ($\beta_{\text{seen-negative}}=-0.029, p=.815$). This suggests that subjects did not notice the repetition of investment opportunities.

The presence of the fee significantly reduced investment likelihood ($\beta_{\text{fee-present}}=-0.366, p<.001$), while observing a larger previous return significantly increased it ($\beta_{\text{prev-return}}=1.152, p<.001$).

¹ To facilitate the assessment of the relative strengths of the different effects, all variables were rescaled to [0;1] prior to estimation of all regressions presented in the paper. In addition, to allow for comparisons with subsequent analyses that include eye-data, trials in which no eye fixations on words were recorded while the word cloud was shown were removed from all analyses (less than 5% of all trials). We used the R (version 3.3.3) lme4 and lmerTest packages to estimate all regressions and compute the coefficient p-values via Satterthwaite's approximation.

[Table 1 here]

Crucially, a larger proportion of positive words in the cloud increased the likelihood of investment ($\beta_{\text{prop-positive}}=2.119$, $p<.001$). This suggests that subjects were able to assess the sentiment of opinions and this informed their decisions in the expected direction.

Finally, it should be noted that the time spent examining the word cloud in the absence of the fee ($Mdn=9.08s$) was not significantly different than in its presence ($Mdn=8.85s$), Wilcoxon two-tailed $Z=-0.784$, $p=.420$. Similarly, the number of words looked at without the fee ($Mdn=15.06$) was not significantly different than in its presence ($Mdn=14.70$), $Z=-0.335$, $p=.738$.

Confirmatory analysis of the effect of the fee on gaze duration on positive vs. negative words. Having verified that our experimental manipulation worked as intended, we now proceed to test our hypotheses. Specifically, we hypothesized that the presence of fees would induce subjects to process opinions differently, depending on whether an opinion's positive or negative sentiment is congruent with the presence or absence of the fee, in the sense that both influence the decision in the same direction. Based on existing research, congruent opinions should be processed faster, resulting in shorter gaze durations on positive words relative to negative ones when the fee is absent rather than present, i.e. given positive rather than negative early information.

To test this hypothesis, we computed the duration of looking at individual words across all subjects and trials. We defined the looking duration as the total duration of successive eye fixations on a word. Specifically, each word constituted a separate Area-of-Interest, constrained by a rectangle centered around the word, with a constant height of 45px (approximately 1° of a 1280x720 screen at a 70cm viewing distance), and a variable width equal to the word width plus a padding equal to the width of a single letter on each side (we used a monospaced font). The

minimum size of an AOI was 60x45px and the AOIs never overlapped, with the minimum distance between an AOI and its nearest neighbor being at least 10px for 95% of the words. We set the minimum required fixation duration to 120 ms, with a maximum dispersion of 45px⁽²⁾. If a word was re-visited after seeing other words in the interim, we treated this as a separate observation, but the results are robust to only including instances of looking at each word for the first time.

Examining the basic descriptive statistics of gaze duration reveals that the average duration of looking at negative words, across all subjects and trials, was 342 ms both with and without the fee, while for positive words it equaled 336 ms in the absence of the fee vs. 340 ms when it was present. In other words, at the aggregate level, the fee seems to increase the duration of looking at positive words, while having no effect on the negative ones.

To assess the statistical significance of this observation, we analyzed the effect of the fee on the duration of looking at positive and negative words, while controlling other factors that might influence the time spent looking at individual words, such as their length or on-screen position. This was to verify that the effect of the fee was not caused by a change in the information search strategy, i.e. in how people decide which words to look at (e.g., by creating a tendency to look at longer words, words that are closer to the center of the screen, etc.). Thus, we

² There is much debate among eye-tracking researchers as to the optimal value of the minimum fixation duration threshold, with thresholds ranging from 50ms (Inhoff & Radach, 1998) to 200ms (Manor & Gordon, 2003) being widely applied. In our case, we did not want to set the threshold too low in order to focus on those instances of looking at words that were long enough for the subject to actually read and understand the word, and particularly its positive vs. negative sentiment. Although some studies reported that people might require less than 100ms to successfully read simple words, this occurs only when successive words are displayed at a fixed point in the center of the screen, eliminating the need for saccadic eye movement (Rubin & Turano, 1992). In our studies, saccades were not only necessary but the arrangement of the words on the screen was irregular, unlike in standard, stationary text, making reading more complicated.

wished to find out if the fee will have a different effect on the duration of looking at a word depending on whether it is a positive or negative word, but given that other features of the word that we control, such as its length, are the same in each case.

To this end, we estimated a mixed-effects linear regression in which the dependent variable was the duration of looking at an individual word, defined as above. The model included hierarchically nested random intercept and slope effects. Specifically, as each subject was presented with an independently drawn random selection of investment opportunities, we treat stimuli (trials) as nested under subjects, thus allowing for correlation between measures (e.g., durations) of fixations made by a given subject in a given decision trial. Apart from indicators of the sentiment of the word and of the presence of the fee, we aimed to include as controls all variables that might influence the processing of individual words.

As seen in Table 2, the control variable effects were largely as expected. Subjects spent more time looking at words that are longer, $\beta_{\text{length}}=11.612*10^{(-3)}$, $t(158179)=47.935$, $p<.001$, less common in the whole corpus of seekingalpha.com articles, $\beta_{\text{frequency-in-corpus}}=-3.028*10^{(-3)}$, $t(156913)=-16.716$, $p<.001$, have not already been seen in the same trial, $\beta_{\text{seen-before}}=-1.768*10^{(-3)}$, $t(158791)=-15.309$, $p<.001$, and were located further from the center of the screen, $\beta_{\text{distance-from-center}}=1.024*10^{(-3)}$, $t(158618)=5.387$, $p<.001$. These results are consistent with the prevalent view in eye-tracking research on reading that ‘readers make longer pauses at points where processing loads are greater’ (Just & Carpenter, 1980; see also Rayner, 1998). The fact that words were examined longer if seen later during a trial, $\beta_{\text{order-seen}}=0.061*10^{(-3)}$, $t(88235)=19.770$, $p<.001$, is, in turn, consistent with studies reporting increased fixation time with more exposure to a stimulus (Król & Król, 2018).

However, the frequency in the cloud of words with the same sentiment as the current word had no significant effect on gaze duration, $\beta_{\text{sentiment-prevalence}} = -0.267 \times 10^{-3}$, $t(5993) = -0.922$, $p = .357$. That is, words that were sentiment outliers in the current word cloud (representing a relatively uncommon sentiment) were not processed differently, e.g. positive words were not looked at longer or shorter when most of the words in the current word cloud were negative.

Finally, the reduced durations in later trials, $\beta_{\text{n-trial}} = -3.763 \times 10^{-3}$, $t(6538) = -21.114$, $p < .001$, may signify a reduction in processing load with more experience, a phenomenon observed across different domains (e.g. Konstantopoulos, Chapman, & Crundall, 2010). Thus, consistent with existing research, looking durations seemed well-aligned with processing load.

[Table 2 here]

More importantly for our hypotheses, in the absence of the fee (fee-present = 0) the duration of looking at negative words was significantly larger than for positive ones, $\beta_{\text{negative}} = 0.776 \times 10^{-3}$, $t(584) = 5.196$, $p < .001$. This translates to a difference of approximately 9 ms before rescaling, which is larger than the 6 ms difference between the corresponding raw averages reported above, because it measures the ceteris paribus effect, whereas positive words were, on average, slightly longer than negative ones, thus taking longer to read. However, the difference in the duration of looking at negative vs. positive words was significantly reduced when the fee was present, $\beta_{\text{negative*fee-present}} = -0.348 \times 10^{-3}$, $t(116) = -2.529$, $p = .013$. Put it differently, the effect of the fee on the duration of looking at negative words was significantly smaller (more negative) than its effect on positive words. In particular, in case of positive words (negative = 0), the effect of the fee was significantly positive, $\beta_{\text{fee-present}} = 0.301 \times 10^{-3}$, $t(5622) = 2.465$, $p = .014$, translating to approximately 4 ms. In contrast, re-estimating the regression in

Table 2 with a ‘positive’ word sentiment dummy variable used instead of the ‘negative’ one revealed that the analogous effect of the fee on the duration of looking at negative words was not significantly different from zero, $\beta_{\text{fee-present[negative words]}}=0.050*10^{(-3)}$, $t(5429)=-0.330$, $p=.7417$.

Exploratory analysis of the effect of the fee on pupil dilation

Although existing evaluative priming studies reported valence asymmetries in terms of processing times, we also explored the possibility of obtaining analogous results using other measures of cognitive effort. In particular, we estimated a second model, identical to the one in Table 2, except that instead of using looking duration as the dependent variable, we used peak pupil dilation while fixating on a word, computed net of a baseline calculated for the 500 ms white screen preceding the word cloud. This is a common measure of cognitive effort in reading and listening studies (e.g. Hyona, Tammola, & Alaja, 1995; Zekveld, Heslenfeld, Johnsrude, Versfeld, & Kramer, 2014). On average across all subjects/trials, the peak pupil dilation relative to baseline when looking at negative words was 0.132 mm both with and without the fee, while for positive words it equaled 0.133 mm with- and 0.124 mm without the fee.

The resulting mixed model estimates are relegated to the Appendix, Table A1, due to their exploratory and supplementary nature vis-à-vis the main analysis of looking duration. It is, however, worth noting that they were generally similar to the results in Table 2. Their most important aspect was that the effect of the fee on the processing of positive words was reproduced when using pupil dilation instead of looking duration, $\beta_{\text{fee-present}}=1.203*10^{(-3)}$, $t(6731)=2.423$, $p=.015$. However, the interaction between the fee and word sentiment was, in this case, not significant, $\beta_{\text{negative*fee-present}}=-0.945*10^{(-3)}$, $t(5677)=-1.681$, $p=.093$. As in the case of looking duration, re-estimating the regression with a ‘positive’ dummy variable instead of the ‘negative’ one revealed that the analogous effect of the fee on pupil dilation while looking at

negative words was not significantly different from zero, $\beta_{\text{fee-present[negative words]}}=0.261*10^{(-3)}$, $t(6367)= 0.467$, $p=.640$.

Exploratory analysis of the effect of the fee on opinion interpretation and belief updating. Finally, we wished to make sure that the observed changes in the perception of positive words were indeed a sign of them becoming harder to interpret when incongruent, rather than simply more important for the decision process. Thus, for each subject/trial, we computed the average horizontal position of the eye in the 800 ms during which the return axis (but not yet the return) was shown (Figure 2). We used this as the dependent variable in our final ‘by-trial’ mixed-effects model (Table 3).

[Table 3 here]

Based on existing research on anticipatory eye-movements, we assumed that prior to the return being shown subjects would look further to the right if this is where they expect to find it, i.e. when they are more optimistic about the stock’s subsequent return. Thus, if the fee was indeed making positive words harder to interpret rather than more important, then the effect of the proportion of positive words in the word cloud on the dependent variable should be smaller with the fee present rather than absent.

By way of a manipulation check, our measure appeared to accurately reflect optimism, with subjects looking further towards the right of the axis (more positive returns) when the previous return was larger, $\beta_{\text{prev-return}}=0.021$, $t(7255)= 3.348$, $p<.001$, when the proportion of positive words was higher in the absence of fees, $\beta_{\text{prop-positive}}=0.053$, $t(106)= 4.524$, $p<.001$, and when the trial occurred later in the study, $\beta_{\text{n-trial}}=0.015$, $t(7354)= 2.907$, $p<.004$, possibly reflecting the growing experience of positive (average) returns.

The key result, however, was that under the presence of fees the impact of the proportion of positive words was significantly reduced, $\beta_{\text{fee-present}*\text{prop-positive}}=-0.046$, $t(97)=-2.890$, $p=.005$, suggesting that positive opinions no longer translated into optimistic beliefs to the same extent as without fees. While we did not include the investment decision among controls (as it is likely co-determined with the dependent variable), similar results are obtained when only including in the analysis the trials in which subjects actually chose to invest. Thus, the effect is not an artefact of the fee influencing the decision and not a mere reflection of ‘wishful thinking’ or lack of interest in returns when not investing.

Discussion

We presented the results of an eye-tracking study in which subjects learned about the presence or absence of a transaction fee before viewing expert opinions about the given stock, sourced from an online financial platform, and finally choosing whether or not to invest. In the event of investment, they received real monetary rewards determined by the actual subsequent return of the stock in the stock market. Each investment opportunity was seen twice: with and without the transaction fee, and we classified the opinion words as positive or negative based on widely used sentiment lexicons.

We hypothesized that, in line with existing research on pre-decisional information distortion, the presence of the fee would affect the way in which opinions are processed, facilitating the processing of opinions congruent with the sentiment towards investment induced by the fee. Based on the density hypothesis, we also anticipated a potential valence asymmetry in this respect, whereby positive words might be particularly strongly ‘primed’ by positive early information (no fee). More precisely, consistent with existing research, we hypothesized that subjects will take longer to examine positive words in the opinion word cloud when the fee is

present rather than absent, i.e. when they are initially less positively inclined towards investment. Given an anticipated negligible impact of the fee on the processing of negative words, the duration of looking on positive words should be higher relative to that on negative words when the fee is present rather than absent.

The effect of the fee on gaze duration on positive vs. negative words. We conducted a mixed-model regression analysis of the duration of looking at individual words. This made it possible to examine the ceteris paribus impact of the fee on the duration of looking at positive vs. negative words, while controlling other word features that might influence the gaze duration. We found that the fee did, indeed, increase the relative duration of looking at positive vs. negative words. This could be seen as consistent with existing studies of pre-decisional distortion of information, where exposure to early information favoring one of the choice options over the other was found to influence the interpretation of subsequent evidence in favor of the initially preferred option (e.g. Miller, DeKay, Stone, & Sorenson, 2013). In our case, the fact that subjects spend more time examining positive words relative to negative ones when aware of a fee that discourages investment could be a sign of positive words then becoming ‘less positive’, and therefore harder to interpret. At the same time, this effect was found to be driven by a change in the processing of positive words, while that of negative ones was largely unaffected by the fee. This, in turn, is in line with the valence asymmetries in evaluative priming reported by existing research in other contexts, as discussed in more detail in the General Discussion section.

The effect of the fee on pupil dilation and inferred beliefs

We also conducted two exploratory analyses to further support and help interpret the above findings. First, as an alternative measure of cognitive effort, we used peak pupil dilation while looking at individual words instead of looking duration, in an otherwise unchanged mixed-

model structure. We found that the fee increased pupil dilation while looking at positive words but, once again, had no impact on the processing of negative ones, giving a further indication of more effortful processing of positive words when incongruent with the fee. Nevertheless, these supplementary findings should be interpreted with caution, subject to caveats which we later discuss in the ‘Scope and Limitations’ section.

Second, through a trial-level mixed-model, we showed that the fee weakened the impact of the proportion of positive words in the cloud on the subjects’ optimism about the stocks’ subsequent returns inferred via anticipatory eye-movements. This suggests that, when incongruent, positive words were harder to process and interpret, rather than more important for the decision process (in which case their more effortful processing would yield greater, not smaller, effect on beliefs). When more words in the cloud were positive rather than negative (the proportion of positive words increased), these additional positive words were harder to interpret in the presence than in the absence of fees, and thus contributed less to positive expectations of future returns.

All in all, the results seemed to confirm our hypotheses, and are consistent with the idea that early positive information can facilitate the processing of later positive evidence. Specifically, positive information in the form of the absence of the fee decreased the gaze duration and pupil dilation while later looking at positive opinions about the stock, at the same time increasing the influence of these opinions on the subjects’ optimism about subsequent decision outcomes.

Experiment 2

Motivation for an additional replication study

Despite our different Experiment 1 measures and tests converging into a consistent picture, it should be noted that the magnitudes of the observed effects were quite small. In particular, the fee increased the duration of looking at positive words by just 4 ms on average, i.e. by only slightly more than 1%. In the same vein, on the individual word-level, the estimated impact of the fee on the duration of looking on positive words was just a small fraction of the difference in this respect between a very long and a very short word, or roughly a third of the difference between a word located at the center of the screen and one placed at the peripheries of the word cloud.

On the one hand, this is not surprising, since physical, objective, and readily accessible features of the stimuli are bound to have more impact on their visual processing than factors that might have engendered subjective psychological predispositions towards the stimuli in some of the observers. On the other hand, the small observed sizes of the hypothesized effects made it essential to replicate our initial findings in another experiment, possibly with slight adjustments in the design to eliminate potential confounds. In particular, we could not be entirely sure that the absence of the fee was perceived by the subjects as a positive (rather than neutral) event. Similarly, the fact that negative early information stemmed from the presence of an event, while positive information was based on its absence could also be considered a problem.

Accordingly, we conducted a replication of Experiment 1, pre-registered on OSF, (link: https://osf.io/t8mbc/?view_only=4a36eb6e609c4728a631f7be949d7d2a), in which in place of the absence of the fee subjects received a positive premium of the same value as the fee. That is, while in Experiment 1 the adjustment applied to the stock return in the event of investment was

either '0' (fee-present = 0) or '-20' (fee-present = 1), in Experiment 2 it was either '+20' or '-20' (where, for consistency, in describing the results we use the same dummy variable notation as before, except now fee-present = 0 means that a +20 adjustment was in place, whereas fee-present = 1 still means a -20 adjustment). In both experiments, the value of the adjustment in a given trial was communicated to subjects in exactly the same way, prior to seeing the word cloud. That is, the only change compared with Experiment 1 was that '0' was replaced with '+20'.

Our pre-registered hypotheses comprised the replication of the effects of the fee reported in Experiment 1, via an unchanged set of mixed-model analyses.

Method

Subjects. A statistical power analysis was performed for sample size estimation, based on data from Experiment 1. However, due to difficulties in conducting power analysis in a mixed-model setting, we based it on a simple test of our main effect on the aggregate (subject) level. Specifically, for each subject, we calculated the difference in average gaze duration on positive vs. negative words, separately for when the fee was present vs. absent. The resulting paired Wilcoxon test of the effect of the fee yielded an effect size of $d = 0.295$. With an $\alpha = .05$ and power = 0.80, the projected sample size needed with this effect size (based on GPower 3.1) was 97. With this in mind, we aimed to recruit up to 120 students for Experiment 2, so that the final sample after exclusions would not fall below this threshold.

The experiment was conducted at the same location as Experiment 1. A total of 118 students volunteered for the study, of which we excluded 15 due to having previously taken part in Experiment 1, poor eye-tracking calibration or data quality (no eye fixations in more than 50%

of choice trials). This left a final sample of 103 subjects (mean age 26.7, 66 females), all of whom had normal or corrected-to-normal eyesight and did not take part in Experiment 1.

Stimuli and Design. Experiment 2 was identical to Experiment 1, apart from a single exception. Specifically, in those trials in which the fee was absent (fee-present = 0), subjects who chose to invest received a payoff adjustment on top of the stock's returns equal to '+20' (compared with '0' in Experiment 1 and '-20' in the fee-present = 1 condition in both experiments). In line with this change, the 'transaction fee:' caption in the initial decision screen (Figure 1, top) was replaced with a more general 'payoff adjustment' caption. In all other respects, the adjustment was still communicated to subjects in the same way and, in particular, prior to the word cloud being shown.

Results

In terms of the overall descriptive statistics, the average duration of looking at negative words, across all subjects and trials, was 329 ms both with and without the fee, while for positive words it equaled 327 ms in the absence of the fee (the +20 condition) vs. 332 ms when it was present. Thus, compared with Experiment 1, subjects' fixations were slightly shorter, but once again the fee seemed to increase the duration of looking at positive words, while having no effect on the negative ones. In addition, the overall number of fixations per trial was reduced by approximately 15% compared with Experiment 1. This was probably caused by the fact that with the payoff adjustment now being either -20 or +20 (instead of 0 or +20), learning which of these alternatives occurred provided subjects with a stronger cue as to which choice is optimal, thus making the subsequent word cloud less important. This resulted in a tendency to read fewer words and spend less time reading those that were looked at.

We estimated three mixed regression models (reported in the Appendix), identical in structure to the ones used to analyze the data from Experiment 1 (with the exception of excluding the insignificant ‘sentiment-prevalence’ control variable). The first model (Table A2) replicated the previously established finding that the fee had a positive effect on the duration of looking at positive words, $\beta_{\text{fee-present}}=0.550*10^{(-3)}$, $t(2599)=3.209$, $p=.001$, translating into approximately 5 ms before rescaling³. Additionally, the duration of looking at positive words relative to negative ones was larger when the fee was present rather than absent, $\beta_{\text{negative*fee-present}}=-0.559*10^{(-3)}$, $t(5667)=-2.360$, $p=.018$. As before, re-estimating the model with a ‘positive’ rather than ‘negative’ dummy variable showed no significant effect of the fee on the duration of looking at negative words, $\beta_{\text{fee-present[negative words]}}=-0.009*10^{(-3)}$, $t(3322)=-0.046$, $p=.963$.

The second model (Table A3) replicated the analogous findings for the alternative, peak relative pupil dilation measure of cognitive effort. Specifically, the fee had a positive effect on pupil dilation while looking at positive words, $\beta_{\text{fee-present}}=3.081*10^{(-3)}$, $t(111)=2.141$, $p=.034$, and pupil dilation while looking at positive relative to negative words was larger when the fee was present rather than absent, $\beta_{\text{negative*fee-present}}=-2.206*10^{(-3)}$, $t(5344)=-3.566$, $p<.001$. Once again, re-estimating the model with a ‘positive’ dummy variable showed no significant effect of the fee on pupil dilation while looking at negative words, $\beta_{\text{fee-present[negative words]}}=0.874*10^{(-3)}$, $t(114)=0.604$, $p=.547$.

Finally, the third model (Table A4) replicated the effect of the fee on inferred optimism about the stock’s subsequent return. Specifically, while in the absence of the fee optimism increased with the proportion of positive words in the word cloud, $\beta_{\text{prop-positive}}=0.032$,

³ Note that the proportional increase of the regression coefficient relative to Experiment 1 is greater than the corresponding increase of its millisecond equivalent. This is because the rescaling of Experiment 2 data is based on different extreme values of recorded variables than it was the case for Experiment 1.

$t(138) = 3.365, p = .001$, this effect was significantly smaller when fees were present, $\beta_{\text{fee-present} * \text{prop-positive}} = -0.037, t(637) = -2.933, p = .003$. As in Experiment 1, the findings were robust to only using observations from trials in which subjects chose to invest in the stock.

General Discussion

The overall picture that we obtained is that the presence of the fee influences the processing of subsequent positive vs. negative information. This occurred regardless of whether a positive premium was paid in the absence of the fee (Experiment 2) or not (Experiment 1). Either way, the absence of the fee was apparently seen by the subjects as a positive event and facilitated the processing of subsequent opinion words congruent with its valence (that is, positive) relative to the incongruent negative words. This was manifested in the fact that, in the absence of the fee, positive words were read faster than when it was present (both in absolute terms and relative to negative words), and yet had a greater impact on the subjects' beliefs. At the same time, no effect of the fee on the processing of negative words was found.

In our view, these findings create a link between existing research showing the pre-decisional distortion as the product of maximizing the consistency between old and new information (J Edward Russo et al., 2008), and the work centered around the density hypothesis (Unkelbach et al., 2008). On the one hand, research on information distortion demonstrated that early information supporting a particular choice option can distort the interpretation of subsequent evidence, with evidence in favour of the leading option being seen as stronger and more unambiguously supportive of that option ('pro-leader distortion'), and evidence supporting the trailing option being seen as weaker and less strongly in its favour ('anti-trailer distortion'). The two types of distortion are typically symmetric (Blanchard, Carlson, & Meloy, 2014;

DeKay, Miller, Schley, & Erford, 2014), with some evidence of the anti-trailer distortion dominating in certain contexts (Nurek, 2014). It seems likely that evidence that is more ambiguous and weaker would also be more difficult to process, requiring more cognitive effort. Thus, the fact that, in our study, we see positive words being processed more effortfully in the presence of fees, with a smaller effect on beliefs, could mean that we observe the attentional correlates of information distortion. What this contributes to the information distortion literature is that most, if not all of this existing research is based on tracing the subjects' cognitive processes by directly and repeatedly asking them about their preferences and interpretation of each piece of evidence. As acknowledged by Russo (2014), it cannot be ruled out that this belief elicitation procedure could itself drive the distortion, e.g. subjects who volunteered an opinion favourable to an option could feel bound to interpret subsequent evidence accordingly, to avoid openly contradicting their previous judgments. In contrast, in our case direct belief elicitation is absent, and yet we do observe patterns consistent with information distortion in the subjects' eye-data.

What should also be noted is that the distinction between pro-leader and anti-trailer distortions is not analogous to our positive/negative dichotomy. Specifically, the equivalent of a pro-leader distortion in our case would be if positive words become easier to interpret ('more positive') in the absence of fees, while negative ones become easier to interpret with the fees present. In contrast, an anti-trailer distortion would occur if negative words become 'less negative' in the absence of fees and positive words become 'less positive' in their presence. Thus, observing an effect consistent with both types of distortion for positive words, but no effect for negative words, neither supports nor contradicts the previous reports of a symmetry between pro-leader and anti-trailer distortions, or of asymmetries in favour of the latter.

What our results do suggest, however, is that certain properties of the early-encountered evidence, and particularly its valence, could determine its potential to cause an information distortion. Bringing the two mentioned strands of literature together, this role of information valence is, in turn, well explained by existing research on evaluative priming, and specifically the density hypothesis. As argued by Alves, Koch, and Unkelbach (2017a), human preferences towards most attributes relevant to their life are single-peaked (that is, a positive range is located in the middle of an attribute dimension, flanked by two negative ranges toward the two ends of the dimension). With extremity being, in general, negative, and moderation positive, the moderate (positive) pieces of information tend to lie closer together on average than the extreme (negative) ones.

The consequence of this tendency is that positive information ends up being, loosely speaking, more densely packed in the associative network of the mind (hence the name of the hypothesis), allowing for easier and faster associations between different pieces of positive information. As shown by Unkelbach et al. (2008), preceding a positive target stimulus with a positive prime object facilitates classifying the target as positive, but this priming effect is stronger than when preceding a negative target with a negative prime to elicit a negative response. In our case, a positive initial information in the form of the absence of fees (and the positive early sentiment to investment that it induces) might facilitate classifying positive opinion words as positive. Introducing the fee (i.e., negative early information) might take this advantage away from positive words, without transferring it to negative ones, because a negative initial sentiment is not as readily connected to or associated with negative expert opinions.

This asymmetry could have important consequences for our understanding of the pre-decisional information distortion. It suggests that the goal of achieving consistency between old

and new information, previously shown to be a major driver of this phenomenon, could be more readily achieved by the brain when positive rather than negative information arrives early on.

Importantly for both the evaluative priming and information distortion literatures, existing research in these areas is based predominantly on tasks in which the chosen answers have little or no direct consequence for the subjects, like rating pictures or statements (even in studies of information distortion in risky choices, e.g. J.E. Russo & Yong, 2011, subjects typically receive a certain, fixed payment). In contrast, here, we showed that the same human biases continue to hold in incentivized economic decisions based on real-world data, despite subjects then being motivated to behave in a thoughtful, non-heuristic manner. The fact that this occurs in a financial context could help explain a number of well-documented phenomena in this domain, like the fact that people underreact to negative news about investments they previously made based on earlier positive signals (Frazzini, 2006; Odean, 1998), or that investors update their beliefs more strongly and more accurately based on positive rather than negative information (Kuhnen, 2015).

But perhaps the most important insight from our results is that the pre-decisional distortion could be interpreted and explained via the ‘error management theory’ (Johnson, Blumstein, Fowler, & Haselton, 2013), which posits that cognitive biases can be advantageous, having evolved as the optimal way to manage errors under cognitive and ecological constraints. In particular, we found that, when the fees were present, the proportion of positive words among opinions about the stock had less influence on the subjects’ inferred optimism about the subsequent investment return. Thus, as coherently evidenced by the looking duration, pupil dilation, and inferred beliefs data, positive initial information in the form of the absence of fees appeared to facilitate the processing and interpretation of subsequent positive opinions, which

were processed faster, with less cognitive effort, but more influence on beliefs. However, having no analogous adverse effect on the processing of negative opinions, the positive early information increased the overall sensitivity of the subjects' beliefs to word cloud composition. Thus, our work offers further process-tracing support for the view that information distortion processes may be adaptive (DeKay, 2015). In particular, the primacy of early information in determining decision outcomes, on which existing work on pre-decisional distortion focuses, could, in fact, be only a by-product of a mechanism which evolved to reduce the cost of the decision process, and in which valence asymmetries play a key part.

Scope and Limitations

Despite their interesting potential implications, our design and analyses come with significant caveats and limitations that must be considered. To begin with, in real-world financial markets, transaction fees are usually higher for investments with higher average returns. In contrast, before the start of both of our experiments, we carefully explained to subjects that the presence of the fee was determined at random, independently of the returns. Despite this, we cannot completely rule out that some subjects would nevertheless expect poor returns when the fee was absent. In this scenario, the less effortful processing of positive words in the absence of fees could be due to positive opinions being dismissed by subjects as contrary to their negative expectations. This, however, could not explain the increased sensitivity of inferred beliefs to the proportion of positive opinions in the cloud. Thus, while we cannot rule it out completely, we consider this scenario to be both unlikely and, in contrast with the density hypothesis, unable to account for all of our results.

At the same time, an interesting question for future research would be to try to separate the direct, *ceteris paribus* effect of the fee on the propensity to invest from its indirect effect due

to moderating the processing of subsequent evidence. Existing literature on information distortion approaches this via mediation analyses, with the effect of initial information on final choices mediated by measures of information distortion that occurred ‘in between’ (DeKay, Stone, & Miller, 2011; Miller et al., 2013). In our case, such an analysis is prevented by the fact that we would need to compute a single numerical measure of how distorted the processing of a given word cloud has been. As subjects’ scanpaths are highly idiosyncratic and endogenous, it is impossible to acquire a benchmark indicating how the same sequence of words would have been examined in the absence of early information about the fees (equivalent to average ratings of each piece of evidence provided by control group subjects in existing information distortion studies). At the same time, our analysis of inferred expectations does suggest that the distortion of subsequent information (words) could mediate the effect of early information (fee) on choice. Specifically, subjects are clearly informed that fees are determined at random, irrespective of future returns. Thus, the only way in which the presence of the fee could influence expected returns is, in theory, via its effect on the processing of the words. The fact that we do observe a significant relationship between the fee and subjects’ expectations suggests, therefore, that the fee could influence choices via an indirect as well as direct route, causing a distortion in the processing of subsequent information affecting expectations on which choices, in turn, are based. Nevertheless, allowing for a full-blown mediation analysis within the current setting, i.e. without direct elicitation of beliefs, would be a potentially very useful design improvement.

Other issues that should be considered are of a more technical nature, and are related to the pupil dilation analysis. First, there is a question of whether the pupil can respond to the sentiment of a word before the gaze is transferred to the next one. Classic studies reported pupil latencies under 300 ms in cognitive tasks (Ahern & Beatty, 1979), while recent experiments in

reading and lexical decision tasks demonstrated that the peak pupil latency can be significantly higher (note, however, that this may be due to the need to execute a response after each word, unlike in our study; see e.g. Haro, Guasch, Vallès, & Ferré, 2017). On the one hand, approximately half of the looking durations that we registered were below 300 ms. On the other hand, it appears that the number of long fixations was sufficiently large to allow for significant pupil dilation results despite the noise brought about by the uninformative short fixations. At the same time, the fact that we might have been able to register only the very early phase of pupillary response could explain its small magnitude relative to classic studies in which exposure to the stimuli is much longer, and the differences in pupil dilation are closer to 0.1mm (e.g. Beatty, 1982).

The second, closely related issue affecting the pupil dilation analysis is that pupil dilation could ‘lag behind’ the eye-movements, leading to order-dependence and autocorrelation between the present and past observations. From the statistical point of view, this is controlled by the clustering of observations by trial within the mixed models. The fact that subjects cannot infer the sentiment of the word prior to reading it, and hence cannot choose the order in which to read positive vs. negative words (which are thus effectively sampled at random), ensures that the issue in question increases noise rather than constituting a systematic confound. In connection with the pupil latency issue above, it may be that seeing two or more words of the same positive or negative valence in succession triggers a pupillary response that only becomes registered and assigned to the words that are close to the end of the sequence. This could be enough to lead to overall differences in pupillary responses to positive vs. negative words, depending on the presence of the fee.

More generally, our experimental paradigm reflects the common tradeoff between realism and control. On the one hand, the simultaneous presentation of all opinions in a word cloud makes it possible to study the way in which people examine evidence when able to freely explore its various elements and choose the duration of each examination, as they do in the real world. On the other hand, this gives us less control over the order and timing of the processing of different pieces of information by our subjects, making it harder to interpret the obtained process data. An improved balance between these two aspects of the tradeoff may be found in future research.

Conclusions

We used eye-tracking in a laboratory stock trading experiment to study the cognitive mechanisms behind the phenomenon of pre-decisional distortion of information. We found evidence suggesting that transaction fees inducing a negative initial sentiment towards investment made subsequent positive opinions about stocks harder to process, with increased cognitive effort manifested in larger gaze duration and pupil size. Despite this increased effort, positive opinions then had a smaller effect on beliefs. In a pre-registered follow-up study, we replicated these findings in a setting in which, in the absence of fees, the payoff adjustment was strictly positive rather than equal to zero.

Interestingly, our process-tracing analysis also demonstrated that a positive vs. negative valence asymmetry, widely documented in evaluative priming, semantic, person perception, and related tasks, extends to incentivized economic choices. In particular, the presence of the transaction fee affected the processing of positive, but not negative opinions, in line with the density hypothesis, which posits greater associative links between positive than between negative pieces of information.

733 Importantly, the fact that the processing of positive information could be facilitated by
734 earlier exposure to positive evidence without hindering that of negative information suggests that
735 the overweighting of early evidence seen in studies of pre-decisional distortion might be a
736 signature of an adaptive heuristic rather than a detrimental decision bias. More specifically, the
737 pre-decisional distortion might be driven by a tendency to reduce the information processing
738 costs, by exploiting similarities, between or within certain categories of data, prevalent in the
739 information ecology that humans operate in. From this perspective, our demonstration of the fact
740 that valence asymmetries matter for information integration in fully incentivized choices is also
741 significant. It suggests that the said focus on the processing costs is present not just in choices of
742 no direct consequence for the decision-makers (like rating pictures or words), but also in ones in
743 which they have a vested interest and an incentive to choose carefully. Thus, valence
744 asymmetries present in pre-decisional integration of information could have important real-world
745 implications.

746

References

- Ahern, S., & Beatty, J. (1979). Pupillary responses during information processing vary with Scholastic Aptitude Test scores. *Science*, 205(4412), 1289–1292.
- Alves, H., Koch, A., & Unkelbach, C. (2017a). The “Common Good” Phenomenon: Why Similarities Are Positive and Differences Are Negative. *Journal of Experimental Psychology. General*, 146.
- Alves, H., Koch, A., & Unkelbach, C. (2017b). Why Good Is More Alike Than Bad: Processing Implications. *Trends in Cognitive Sciences*, 21(2), 69–79.
<https://doi.org/10.1016/j.tics.2016.12.006>
- Ashby, N. J. S., Jekel, M., Dickert, S., & Glöckner, A. (2016). Finding the right fit: A comparison of process assumptions underlying popular drift-diffusion models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(12), 1982.
- Beatty, J. (1982). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin*, 91(2), 276–292. <https://doi.org/10.1037/0033-2909.91.2.276>
- Blanchard, S. J., Carlson, K. A., & Meloy, M. G. (2014). Biased Predecisional Processing of Leading and Nonleading Alternatives. *Psychological Science*, 25(3), 812–816.
<https://doi.org/10.1177/0956797613512663>
- Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for Someone to Blame: Delegation, Cognitive Dissonance, and the Disposition Effect. *Journal of Finance*, 71(1), 267–302. <https://doi.org/10.1111/jofi.12311>

- 769 Chen, H., De, P., Hu, Y., & Hwang, B. (2014). Wisdom of Crowds: the Value of Stock Opinions
770 Transmitted Through Social Media. *Review of Financial Studies*, 27(5), 1367–1403.
771 <https://doi.org/10.1093/rfs/hhu001>
- 772 Da, Z., Engelberg, J., & Gao, P. (2015). The Sum of All FEARS Investor Sentiment and Asset
773 Prices. *Review of Financial Studies*, 28(1), 1–32. <https://doi.org/10.1093/rfs/hhu072>
- 774 DeKay, M. L. (2015). Predecisional Information Distortion and the Self-Fulfilling Prophecy of
775 Early Preferences in Choice. *Current Directions in Psychological Science*, 24(5), 405–411.
776 <https://doi.org/10.1177/0963721415587876>
- 777 DeKay, M. L., Miller, S. A., Schley, D. R., & Erford, B. M. (2014). Proleader and antitrailer
778 information distortion and their effects on choice and postchoice memory. *Organizational*
779 *Behavior and Human Decision Processes*, 125(2), 134–150.
- 780 DeKay, M. L., Stone, E. R., & Miller, S. A. (2011). Leader-driven distortion of probability and
781 payoff information affects choices between risky prospects. *Journal of Behavioral Decision*
782 *Making*, 24(4), 394–411.
- 783 Frazzini, A. (2006). The disposition effect and underreaction to news. *Journal of Finance*, 61(4),
784 2017–2046. <https://doi.org/10.1111/j.1540-6261.2006.00896.x>
- 785 Frydman, C., Barberis, N., Camerer, C., Bossaerts, P., & Rangel, A. (2014). Using neural data to
786 test a theory of investor behavior: An application to realization utility. *Journal of Finance*,
787 69(2), 907–946. <https://doi.org/10.1111/jofi.12126>
- 788 Gerard, H., Gordon, P., & Nagpurnanand, P. (2013). Product Market Threats, Payouts, and
789 Financial Flexibility. *The Journal of Finance*, 69(1), 293–324.
790 <https://doi.org/10.1111/jofi.12050>

- Gilead, M., Sela, M., & Maril, A. (2018). That's My Truth: Evidence for Involuntary Opinion Confirmation. *Social Psychological and Personality Science*, 194855061876230. <https://doi.org/10.1177/1948550618762300>
- Glöckner, A., & Herbold, A. K. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, 24(1), 71–98. <https://doi.org/10.1002/bdm.684>
- Gräf, M., & Unkelbach, C. (2016). Halo Effects in Trait Assessment Depend on Information Valence: Why Being Honest Makes You Industrious, but Lying Does Not Make You Lazy. *Personality and Social Psychology Bulletin*, 42(3), 290–310. <https://doi.org/10.1177/0146167215627137>
- Haro, J., Guasch, M., Vallès, B., & Ferré, P. (2017). Is pupillary response a reliable index of word recognition? Evidence from a delayed lexical decision task. *Behavior Research Methods*, 49(5), 1930–1938. <https://doi.org/10.3758/s13428-016-0835-9>
- Hendershott, T., Livdan, D., & Schürhoff, N. (2015). Are institutions informed about news? *Journal of Financial Economics*, 117(2), 249–287. <https://doi.org/10.1016/j.jfineco.2015.03.007>
- Hyona, J., Tammola, J., & Alaja, A. (1995). Pupil Dilation as a Measure of Processing Load in Simultaneous Interpretation and Other Language Tasks. *The Quarterly Journal of Experimental Psychology Section A*, 598–612. <https://doi.org/10.1080/14640749508401407>
- Inhoff, A. W., & Radach, R. (1998). Definition and computation of oculomotor measures in the study of cognitive processes. In *Eye guidance in reading and scene perception* (pp. 29–53). Elsevier.

- 813 Johnson, D. D. P., Blumstein, D. T., Fowler, J. H., & Haselton, M. G. (2013). The evolution of
814 error: error management, cognitive constraints, and adaptive decision-making biases.
815 *Trends in Ecology & Evolution*, 28(8), 474–481.
816 <https://doi.org/https://doi.org/10.1016/j.tree.2013.05.014>
- 817 Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to
818 comprehension. *Psychological Review*, 87(4), 329–354. [https://doi.org/10.1037/0033-](https://doi.org/10.1037/0033-295X.87.4.329)
819 [295X.87.4.329](https://doi.org/10.1037/0033-295X.87.4.329)
- 820 Koch, A., Alves, H., Krüger, T., & Unkelbach, C. (2016). A General Valence Asymmetry in
821 Similarity: Good Is More Alike Than Bad. *Journal of Experimental Psychology. Learning,*
822 *Memory, and Cognition*, 42.
- 823 Konstantopoulos, P., Chapman, P., & Crundall, D. (2010). Driver’s visual attention as a function
824 of driving experience and visibility. Using a driving simulator to explore drivers’ eye
825 movements in day, night and rain driving. *Accident Analysis & Prevention*, 42(3), 827–834.
826 <https://doi.org/10.1016/j.aap.2009.09.022>
- 827 Król, M. E., & Król, M. (2018). “Economies of Experience”—Disambiguation of Degraded
828 Stimuli Leads to a Decreased Dispersion of Eye-Movement Patterns. *Cognitive Science*, 42,
829 728–756.
- 830 Kuhnen, C. M. (2015). Asymmetric Learning from Financial Information. *Journal of Finance*,
831 70(5), 2029–2062. <https://doi.org/10.1111/jofi.12223>
- 832 Loughran, T., & McDonald, B. (2011). When is a Liability not a Liability? Textual Analysis,
833 Distionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65. [https://doi.org/10.1111/j.1540-](https://doi.org/10.1111/j.1540-6261.2010.01625.x)
834 [6261.2010.01625.x](https://doi.org/10.1111/j.1540-6261.2010.01625.x)

- 835 Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of*
836 *Financial Economics*, 123(1), 137–162.
837 <https://doi.org/https://doi.org/10.1016/j.jfineco.2016.01.032>
- 838 Manor, B. R., & Gordon, E. (2003). Defining the temporal threshold for ocular fixation in free-
839 viewing visuocognitive tasks. *Journal of Neuroscience Methods*, 128(1–2), 85–93.
- 840 Miller, S. A., DeKay, M. L., Stone, E. R., & Sorenson, C. M. (2013). Assessing the sensitivity of
841 information distortion to four potential influences in studies of risky choice. *Judgment and*
842 *Decision Making*, 8(6), 662.
- 843 Nurek, M. (2014). Predecisional information distortion in physicians ' diagnostic judgments :
844 Strengthening a leading hypothesis or weakening its competitor ? *Judgment and Decision*
845 *Making*, 9(6), 572–585.
- 846 Odean, T. (1998). Are Investors Reluctant to Realize Their Losses? *The Journal of Finance*,
847 53(5), 1775–1798. <https://doi.org/10.1111/0022-1082.00072>
- 848 Park, J., Konana, P., Gu, B., Kumar, A., Raghunathan, R., Park, J., & Gu, B. (2013). Information
849 Valuation and Confirmation Bias in Virtual Communities : Evidence from Stock Message
850 Boards. *Information Systems Research*, 24(4). <https://doi.org/10.1287/isre.2013.0492>
- 851 Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research.
852 *Psychological Bulletin*, 124(3), 372.
- 853 Richter, T., Schroeder, S., & Wöhrmann, B. (2009). You don't have to believe everything you
854 read: Background knowledge permits fast and efficient validation of information. *Journal of*
855 *Personality and Social Psychology*, 96(3), 538–558. <https://doi.org/10.1037/a0014038>
856

- 857 Rubin, G. S., & Turano, K. (1992). Reading without saccadic eye movements. *Vision Research*,
858 32(5), 895–902. [https://doi.org/10.1016/0042-6989\(92\)90032-E](https://doi.org/10.1016/0042-6989(92)90032-E)
- 859 Russo, J. E. (2014). The predecisional distortion of information. In *Neuroeconomics, judgment,*
860 *and decision making* (pp. 109–128). Psychology Press.
- 861 Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The Goal of Consistency as a
862 Cause of Information Distortion. *Journal of Experimental Psychology: General*, 137(3),
863 456–470. <https://doi.org/10.1037/a0012786>
- 864 Russo, J. E., & Yong, K. (2011). The distortion of information to support an emerging evaluation
865 of risk. *Journal of Econometrics*, 162(1), 132–139.
866 <https://doi.org/10.1016/j.jeconom.2010.07.004>
- 867 Santos, E. M., & Kowler, E. (2017). Anticipatory smooth pursuit eye movements evoked by
868 probabilistic cues. *Journal of Vision*, 17(13), 13. <https://doi.org/10.1167/17.13.13>
- 869 Shimojo, S., Simion, C., Shimojo, E., & Scheier, C. (2003). Gaze bias both reflects and
870 influences preference. *Nature Neuroscience*, 6(12), 1317–1322.
871 <https://doi.org/10.1038/nn1150>
- 872 Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over
873 the course of decision making. *Journal of Experimental Psychology: Learning, Memory,*
874 *and Cognition*, 27(5), 1250.
- 875 Unkelbach, C., Fiedler, K., Bayer, M., Stegmüller, M., & Danner, D. (2008). Why Positive
876 Information Is Processed Faster: The Density Hypothesis. *Journal of Personality and Social*
877 *Psychology*, 95(1), 36–49. <https://doi.org/10.1037/0022-3514.95.1.36>
878

879 Zekveld, A. A., Heslenfeld, D. J., Johnsrude, I. S., Versfeld, N. J., & Kramer, S. E. (2014). The
880 eye as a window to the listening brain: Neural correlates of pupil size as a measure of
881 cognitive listening load. *NeuroImage*, *101*, 76–86.
882 <https://doi.org/10.1016/j.neuroimage.2014.06.069>
883
884

885 Table 1.
 886 *The likelihood of making a decision to invest (Experiment 1). Summary of a mixed-effects*
 887 *regression model ($N = 7633$) with random subject intercept and slope effects.*

independent variable	β	SE	z	p
intercept	−1.057	0.124	−8.541	<0.001
n-trial	0.165	0.172	0.957	0.338
seen-positive	−0.068	0.110	−0.612	0.540
seen-negative	−0.029	0.126	−0.234	0.815
fee-present	−0.366	0.066	−5.554	<0.001
prev-return	1.152	0.106	10.839	<0.001
prop-positive	2.119	0.131	16.074	<0.001

888 *Note.* The variables ‘seen-positive’ and ‘seen-negative’ are not fully linearly dependent (multi-
 889 collinear), as the third, reference category is seeing the investment opportunity for the first time.
 890 All variables were re-scaled to [0;1] prior to estimation.

891

892 Table 2.

893 *The duration of looking at a word (Experiment 1). Summary of a mixed-effects regression model*
 894 *(N = 159346) with random intercept and slope effects nested by subject/trial.*

independent variable	$\beta^*(10^3)$	$SE^*(10^3)$	t	p
intercept	16.251	0.679	23.942	<.001
n-trial	-3.763	0.178	-21.114	<.001
length	11.612	0.242	47.935	<.001
distance-from-center	1.024	0.190	5.387	<.001
frequency-in-corpus	-3.028	0.181	-16.716	<.001
seen-before	-1.768	0.115	-15.309	<.001
order-seen	0.061	0.003	19.77	<.001
sentiment-prevalence	-0.267	0.290	-0.922	.357
negative	0.776	0.157	5.196	<.001
fee-present	0.301	0.122	2.465	.014
negative*fee-present	-0.348	0.186	-2.529	.013

895 *Note.* All variables were re-scaled to [0;1] prior to estimation. Due to their small values, the
 896 displayed coefficient estimates and standard errors were multiplied by 10^3 .

897

Table 3.

The inferred optimism about subsequent returns (Experiment 1), defined as the average horizontal eye position during the 800ms period in which the subsequent return axis was displayed in a given trial prior to the return being shown. Summary of a mixed-effects regression model ($N = 7633$) with random subject intercept and slope effects.

independent variable	β	SE	t	p
intercept	0.456	0.012	36.806	<.001
n-trial	0.015	0.005	2.907	.004
fee-present	0.025	0.009	2.731	.007
prev-return	0.021	0.006	3.348	<.001
prop-positive	0.053	0.012	4.524	<.001
fee-present*prop-positive	-0.046	0.016	-2.89	.005

Note. All variables were re-scaled to [0;1] prior to estimation.

Figure Captions

Figure 1.

An example sequence of decision screens shown in a single decision trial in Experiment 1. First, the subject sees the previous return of the stock and the transaction fee (stated as '0' in case of no fee). Next (preceded by a 500ms central fixation cross and a 500ms white screen), the top 50 sentiment words are shown, randomly arranged in a word cloud. Upon pressing a key, the subject submits the decision in the final screen. The only difference between any two matched decision trials was whether the transaction fee was 0 or -20.

Figure 2.

An example display sequence shown after a decision was made. First, an axis indicating the range of possible subsequent returns was shown for 800ms (top, the same in all trials), preceded by a 500ms central fixation cross. After 800ms, 10 random characters were displayed above the axis (middle). Only one of them, at a location representing the subsequent return (here, -22), was not upside down. The other characters' locations were random but symmetric around 0. The subject was instructed to look at the 'correct' character and press a key, upon which a reminder of the decision was shown together with the resulting payoff (bottom; in this example the fee of -20 was applied).

Appendix

Additional Tables

Table A1.

The peak relative pupil dilation while looking at a word (Experiment 1). Summary of a mixed-effects regression model ($N = 159346$) with random intercept and slope effects nested by subject/trial.

independent variable	$\beta^*(10^3)$	$SE^*(10^3)$	t	p
intercept	391.908	1.473	266.026	<.001
n-trial	-1.095	0.762	-1.435	.151
length	5.713	0.761	7.512	<.001
distance-from-center	-53.866	0.597	-90.161	<.001
frequency-in-corpus	-0.273	0.569	-0.480	.631
seen-before	1.549	0.361	4.288	<.001
order-seen	-0.338	0.010	-33.794	<.001
sentiment-prevalence	-1.602	0.968	-1.655	.098
fee-present	1.203	0.497	2.423	.015
negative	0.387	0.506	0.823	.410
negative*fee-present	-0.945	0.557	-1.681	.093

Note. All variables were re-scaled to [0;1] prior to estimation. Due to their small values, the displayed coefficient estimates and standard errors were multiplied by 10^3 .

933 Table A2.

934 *The duration of looking at a word (Experiment 2). Summary of a mixed-effects regression model*
 935 *(N = 138669) with random intercept and slope effects nested by subject/trial.*

independent variable	$\beta^*(10^3)$	$SE^*(10^3)$	t	p
intercept	19.115	0.711	26.881	<.001
n-trial	-3.970	0.242	-16.386	<.001
length	20.770	0.410	50.705	<.001
distance-from-center	-0.511	0.266	-1.919	.055
frequency-in-corpus	-3.653	0.251	-14.568	<.001
seen-before	-2.809	0.172	-16.312	<.001
order-seen	7.987	0.953	8.380	<.001
fee-present	0.550	0.171	3.209	.001
negative	0.867	0.215	4.040	<.001
negative*fee-present	-0.559	0.236	-2.360	.018

936 *Note.* All variables were re-scaled to [0;1] prior to estimation. Due to their small values, the
 937 displayed coefficient estimates and standard errors were multiplied by 10^3 .

938

Table A3.

The peak relative pupil dilation while looking at a word (Experiment 2). Summary of a mixed-effects regression model ($N = 137678$) with random intercept and slope effects nested by subject/trial.

independent variable	$\beta^*(10^3)$	$SE^*(10^3)$	t	p
intercept	497.654	3.348	148.629	<.001
n-trial	12.869	2.099	6.132	.151
length	8.467	1.021	8.293	<.001
distance-from-center	-82.597	0.668	-123.591	<.001
frequency-in-corpus	0.241	0.625	0.386	.700
seen-before	-1.230	0.426	-2.887	.004
order-seen	-22.795	2.490	-9.154	<.001
fee-present	3.081	1.439	2.141	.034
negative	1.833	0.437	4.199	<.001
negative*fee-present	-2.206	0.619	-3.566	<.001

Note. All variables were re-scaled to [0;1] prior to estimation. Due to their small values, the displayed coefficient estimates and standard errors were multiplied by 10^3 .

947 Table A4.
 948 *The inferred optimism about subsequent return (Experiment 2). Summary of a mixed-effects*
 949 *regression model ($N = 7652$) with random subject intercept and slope effects.*

independent variable	β	SE	t	p
intercept	0.490	0.011	44.082	<.001
n-trial	0.051	0.004	11.927	.004
fee-present	0.012	0.007	1.578	.115
prev-return	0.019	0.005	3.683	<.001
prop-positive	0.032	0.010	3.365	.001
fee-present* prop-positive	-0.037	0.013	-2.933	.003

950 *Note.* All variables were re-scaled to [0;1] prior to estimation.

951