

Top-Down or Bottom-Up?

Disentangling Channels of Attention in Risky Choice*

Jan B. Engelmann¹, Alejandro Hirmas¹, and Joël J. van der Weele¹

¹University of Amsterdam, Center for Experimental Economics and political Decision making (CREED), and the Tinbergen Institute

February 14, 2022

Abstract

Economists have become increasingly interested in using attention to explain behavioral patterns both on the micro and macro level. This has resulted in several disparate theoretical approaches. Some, like rational inattention, assume “top-down” cognitive control over attention. Others, like salience theory, assume a “bottom-up” influence where attention is driven by contextual factors. This distinction is fundamental for the economic implications of attention, but so far there is little understanding of their *relative* importance. We propose a multi-attribute random utility model that unifies prior theoretical approaches by distinguishing between the impact of top-down and bottom-up attention. We accomplish this by separating agent-specific and decision-specific variation in attention and verify our framework in an eye-tracking experiment on risky choice. We find that both top-down and bottom-up attention are connected to important choice variables: both are associated with the weighting of the attributes of choice options, while top-down attention is additionally associated with measures of loss aversion. We discuss the insights regarding the nature of attention and its role in economic theory.

Keywords: Attention, Random, Utility Models, Eye-tracking, Loss Aversion.

JEL Codes: D81, D83, D87, D91.

*We thank Yangyang Xu and O’Jay Medina for help with data collection. This work was supported by startup funds from the Amsterdam School of Economics, awarded to Jan Engelmann. Joël van der Weele gratefully acknowledges funding by the NWO in the context of VIDI grant 452-17-004.

1 Introduction

Over the last decades, economists have become increasingly interested in attention. For instance, on the microeconomic level, researchers have proposed that attention may explain behavioral biases such as the endowment effect, the attraction effect or the phenomenon of motivated cognition. On the macroeconomic level, limits to attention may explain how economic agents react to news shocks, form expectations about future prices and how this affects business cycles. Alongside these applications, several prominent new theories try to incorporate the role of attention in economic behavior. “Salience theory” explains how prominent features among potential payoffs attract attention and sway decisions, leading to behavioral biases (Bordalo et al., 2012, 2013). Theories of “Rational Inattention” propose that decision makers direct limited attentional resources to information that is deemed to be most useful (Sims, 2010; Gabaix, 2019). Finally, sequential sampling models offer a descriptive framework of how processes of information acquisition translate into decision making (Ratcliff, 1978; Krajbich et al., 2012; Fudenberg et al., 2018).

These theoretical approaches differ fundamentally in their description of economic agents. On the one hand, rational inattention maintains the traditional assumptions of an optimizing agent with executive control over her choices. On the other hand, salience theory views the agent’s attention and her choices as determined by her environment. This discrepancy mirrors a prominent distinction in psychology and neuroscience, where researchers distinguish between “top-down” (also referred to as “endogenous”) and “bottom-up” (or “exogenous”) attention processes (Posner et al., 1980; Egeth and Yantis, 1997; Corbetta and Shulman, 2002). Here, “top-down” refers to the control of attention by internal factors related to predetermined goals and expectations. As an example, consider going to the supermarket with a shopping list that contains items that your partner underlined to signal their importance for their culinary projects. By contrast “bottom-up” refers to attentional control by factors external to the observer, such as the physical salience of different stimuli. In terms of our supermarket example, the shopper might be tempted to

make unplanned purchases of highly salient items, for instance those that are prominently advertised. Importantly, and as illustrated by our example, these two attentional processes jointly influence behavior, suggesting that both rational inattention and salience theory are both valuable but incomplete explanations of attention and behavior.

The distinction between bottom-up vs. top-down attention is key in understanding the important role of attention not just for developing economic theory, but also regarding practical applications. Top-down attentional processes reveal information about the agents intentions and goals, providing insights into agent-specific variables that affect how decisions are made, such as cognitive ability and personality. Simultaneously, bottom-up processes can disrupt ongoing decisions, thereby potentially producing behavioral biases. This causal impact can have interesting applications for policy-makers, who can exploit such knowledge of attentional mechanisms to nudge decisions. While some papers in psychology and neuroscience have tried to quantify the influence of these different channels on behavior, as we discuss in more detail below, there has been little work to understand their relative importance.

In this paper, we propose a formal framework to disentangle bottom-up and top-down attention within the decision-making process. To do so, we adapt a multi-attribute random utility model that considers multiple attentional channels and allows us to identify their relative influence on choice. Specifically, to allow for both types of attention. Like in the attentional drift diffusion model (Krajbich et al., 2012), we assume attention can influence choice by affecting the weighting of different attributes. Our key assumption is that the drivers of attention are separable into agent-based factors like preferences (the top-down channel) and contextual factors or “salience” (the bottom-up channel). We show that under some plausible assumptions, this framework can be used to decompose the variance in attention. The between-subject variation in attention reflects differences between agents’ decision-making processes, and hence captures the top-down channel. By contrast, within-person variation in attention across trials is driven by the salience of specific choice options

on a given trial that influences choice via bottom-up processes.

We demonstrate the applicability of our approach in two original experiments on risky choice. Over multiple trials, subjects choose to accept or reject lotteries with equiprobable losses and gains, which vary between trials. While subjects make choices, we record their attention patterns using eye-tracking. In line with the existence of both top-down and bottom-up attention, we find that both between-subject and within-subject variation in attention explain risky choices, with between-subject variations having a stronger impact than within-subject variations in our experimental context. Further supporting the notion that between-subject variation in attention reflect differences in individual preferences, we find that attentional patterns correlate with individual measures of loss aversion.

As we explain in more detail in the next section, we contribute to the literature on attention in economic choice in various ways. First, we translate a core distinction in attentional research (top-down vs. bottom-up control) into the formal framework of economics, and show how it can be used to decompose attentional variation. This offers an integrated way to think about the relation between preferences, context, attention and choice, which can be used to answer a number of follow-up questions, as we elaborate in the conclusion. Finally, we contribute to the literature on risky choice, by showing that both top-down and bottom-up attentional processes drive risk taking. This shows that risk taking is related to both personal, agent-related characteristics involved in deliberate choices, but also to situational factors such as the salience of specific choice options.

2 Related Literature

The fields of psychology and cognitive (neuro-)science have long studied attention as a mechanism that reduces demands on limited visual and other cognitive systems by filtering relevant information (e.g. Posner, 2011). Recent key empirical findings that show a strong link between visual attention and decisions have attracted the interest of the field of decision science. Specifically, choice options that enter the attentional focus more often and for

longer are more likely to be chosen (Krajbich et al., 2010, 2012; Lim et al., 2011; Polonio et al., 2015; Pachur et al., 2018) and choice options with higher values attract attention more than those with lower values (Anderson et al., 2011; Gluth et al., 2018; Amasino et al., 2019; Gluth et al., 2020).

When it comes to characterizing the determinants of attention, the literature makes a fundamental distinction between top-down and bottom-up channels of attention, as defined in the introduction. Bottom-up attention is thought to have a larger influence on explorative decision processes, when individuals do not yet have a specific rule of choice (Fehr and Rangel, 2011; Gottlieb et al., 2013). Nonetheless, a number of studies have provided evidence that both channels of attention play a role in decision-making (e.g. Orquin and Mueller Loose, 2013; Orquin and Lagerkvist, 2015; Corbetta and Shulman, 2002). Moreover, empirical and theoretical considerations in neuroscience, such as by Corbetta and Shulman (2002) and Ungerleider and Kastner (2000), suggest that the brain may process these types of attention in partially separable neural networks. .

In economic theory, similar distinctions have emerged. The bottom-up approach is represented in “salience theory” proposed in Bordalo et al. (2012, 2013) and related models like Kőszegi and Szeidl (2013). These models propose functions that map different choice attributes into “salience”, which reflects the ease by which they are noted by the decision maker. More salient attributes translate into higher weights of these attributes in the decision. In these models, salience operates in a mechanical way, i.e. without any explicit optimization by the decision maker. It is therefore likely to lead to behavioral biases. Indeed, some of the key insights of these models are to account for a variety of behavioral biases such as the Allais’ paradox or the endowment effect (Bordalo et al., 2012).

By contrast, the top-down perspective is reflected in economic models of rational inattention (Sims, 2003, 2010; Gabaix, 2019; Caplin and Dean, 2015; Bartoš et al., 2016). In these theories, the decision maker optimally allocates scarce attention to those information sources or attributes that are most likely to affect the utility of choice. These models offer

an answer to the question how a decision maker can optimally allocate attention before actually knowing the value of the choice (Gabaix, 2014). Applications have emerged in finance (Peng and Xiong, 2006), business cycle theory (Maćkowiak and Wiederholt, 2015), monetary policy (Mackowiak and Wiederholt, 2009), industrial organisation (Dessein et al., 2016; Fosgerau et al., 2020), and consumer theory (Reis, 2006; Matějka and McKay, 2015; Caplin and Dean, 2015).

Our exercise is motivated by the seemingly disparate views of the relative roles of agent and context that is inherent in these theoretical approaches. Our goal here is to capture these attentional processes within one model, even if in a reduced form, and derive a way to distinguish their relative importance empirically. Most closely related in this endeavor are papers that decompose attention using a number of different methods¹. Fisher (2021) investigates the role of attention in intertemporal discounting, and shows that both within- and between-subject variation in attention allocation correlate with decisions. In addition, random variations in exposure time to different attributes explain about 5% to 10% of intertemporal choices. Ghaffari and Fiedler (2018) attempt to disentangle top-down and bottom-up processes in moral choices. Adapting the well-established empirical result that choices are predicted by the last fixation, they experimentally manipulate the last fixation. Their results indicate that the attribute fixated last is predictive of choice, indicating an effect of bottom-up attention, which they estimate to be responsible for about 11% of the variance in decisions. Third, Towal et al. (2013) perform an eye-tracking experiment on snacks, where they first elicited the value of snacks from participants. They calibrate the parameters of a modified drift-diffusion model (Ratcliff, 1978), where the drift rate can depend on the product’s value and/or salience, a measure constructed from the perceptual

¹Other recent papers have focused on establishing a causal effect of attention, by manipulating attention via visual salience, exposure time or other contextual, bottom-up interventions. Evidence has been presented for such attentional influences on choice in a multitude of domains (see e.g. Armel et al. (2008); Reutskaja et al. (2011); Atalay et al. (2012); Pärnamets et al. (2015); Pachur et al. (2018); Ghaffari and Fiedler (2018); Gluth et al. (2018, 2020)). In economics, Dertwinkel-Kalt et al. (2017) and Dertwinkel-Kalt and Köster (2020) have tested recent models of salience discussed above. These studies have shown that there is a causal effect of attention, although its size is often modest.

features of the products appearance. Value appears as a more important predictor than salience, with a relative weight that is about 3 times higher. Navalpakkam et al. (2010) ask their participants to choose between multiple targets that vary in value and salience, finding a significant effect of both on the decision.

Our paper adds to this literature by providing a theoretical foundation to decompose attentional variation, and assess its contribution to variance in choices. Our traditional multi-attribute utility model is adapted to allow for both top-down and bottom-up attention via their impact on decision weights on the attributes. In doing so, we elucidate the assumptions under which one can approximate bottom-up and top-down processes by within- vs. between-subject variation in attention and choice.

Apart from our methodological insights, we contribute to a literature about the role of attention in risky choice (Fiedler and Glöckner, 2012; Pachur et al., 2018). In particular, we complement findings by Pachur et al. (2018), who show that loss aversion parameters are correlated with attention, and that exogenous variations in attention cause shifts in loss aversion. Our paper adds to this evidence, and shows that loss aversion is correlated with between-subject variation in attention. This is in line with our theoretical approach, which associates between-subject variation in attention with mechanisms that are internal to the agent. Additionally, our finding that bottom-up attention plays a role in risky choice may help explain the instability of decisions in risky choice across contexts (Bordalo et al., 2012; Johnson and Schkade, 1989).

3 Disaggregating attention: A theoretical framework

In this section, we present an attention-based model that accounts for two channels of attention that jointly influence choice: top-down and bottom-up control of attention. We model the decision and attention processes simultaneously. For the decision process, we present first a seemingly independent model without attention. We use the McFadden (1980) representation of a random utility model (RUM), where the randomness in the

decision process can arise from perceptual biases, such as those introduced by attention. We then incorporate top-down attention, and show how this translates into individual differences in both attention and behavior. Finally, we introduce salience (or context-based factors) and its effect on attention and choice. We show how one can exploit the trial-wise variations in attention to identify the effects of salience and bottom-up attention on the decision.

The decision process

Consider the case of a population of agents, indexed $j = 1, 2, \dots, J$. Over a series of (experimental) decisions or trials, indexed $t = 1, 2, \dots, T$, each agent accepts or rejects a choice option x_t . In line with most experimental designs, we assume that all agents face the same set of alternatives $X = \{x_t\}_{t=1}^T$. Option x_t has real-valued attributes indexed $s = 1, 2, \dots, S$, i.e. $x_t = (x_{1,t}, x_{2,t}, \dots, x_{S,t})$. We model the decision of accepting the choice option as a random process:

$$D_{j,t} = \begin{cases} \text{Accept} & \text{if } u_{j,t} \geq \bar{u}_j \\ \text{Reject} & \text{if } u_{j,t} < \bar{u}_j. \end{cases} \quad (1)$$

Here, \bar{u}_j is the outside option associated with rejection, and

$$u_{j,t} = \sum_{s=1}^S \omega_{j,s,t} x_{s,t} \quad (2)$$

is an additively separable multi-attribute utility function that reflects the value of the alternative x_t for the agent. Thus, decisions are determined by the attributes of x , as well as the decision weight ω for attribute s in trial t . We assume that decision weights are random variables

$$\omega_{j,s,t} = \beta_{j,s} + \epsilon_{j,s,t}, \quad (3)$$

where $\beta_{j,s}$ represents the *preference* of agent j over the attribute s . We assume that preferences are stable across the set of trials T . Crucially, agents may value different attributes differently, as reflected in different decision weights. Note that while we refer to β as a “preference”, this should be interpreted as any agent-specific mental process that determines the value of the attribute to the agent. Decision weights may deviate from preferences due to a mean-zero error term $\epsilon_{j,s,t}$. Below, we will operationalize $\epsilon_{j,s,t}$ to reflect the salience of different contextual factors.

Incorporating Attention

During the decision process, the agent allocates attention to the different attributes, which we model parallel to the decision process. In doing so, we distinguish top-down from bottom-up processes of attention. Top-down attention a_s^{TD} depends on payoff valence, as well as the agent’s preferences, i.e. $a_s^{TD} = a_s^{TD}(\beta_{j,s})$. We assume that top-down attention does not directly depend on the size of the attribute, an assumption we will discuss in Section 7.

The second process, bottom-up attention (a_s^{BU}), is a random variable driven by *salience*, denoted by σ_t . Salience is a trial-specific property of an attribute that reflects how it stands out from its surroundings (i.e. $\sigma_t = (\sigma_{1,t}, \sigma_{2,t}, \dots, \sigma_{S,t})$). Salience is driven by contextual factors as color, font size and location on a computer screen. It can also be a function of the way different attributes are contrasted or presented, so that $\sigma_t = \sigma(x_t)$, as in Bordalo et al. (2012). We will be agnostic about the determinants of salience, but we note that this is an important topic for research. Below, we will expand on how salience determines bottom-up attention.

We model the two attentional processes as separable and additive, i.e.

$$a_{j,s,t} = a_s^{TD}(\beta_{j,s}) + a_s^{BU}(\sigma_t) + \nu_{j,s,t}. \quad (4)$$

Here, $\nu_{j,s,t}$ is a zero-mean error-term, uncorrelated with the attentional effects of salience

and preferences. While the separability assumption is of course mathematically and empirically convenient, there is evidence to support it. Pinto et al. (2013) measure bottom-up and top-down attention with independent tasks, and shows that performance is independent across tasks. Moreover, a survey of the literature by Orquin and Lagerkvist (2015) concludes that the effects of the top-down and bottom-up control of attention occur in separate moments of the decision, while Corbetta and Shulman (2002) present evidence that different brain networks are involved in these two attentional processes. In case that the linear separability of top-down and bottom-up assumptions were not valid, we can consider $a_{j,s,t} = \log(A_{j,s,t})$, where $A_{j,s,t}$ is the actual measure of attention. This log-linear representation allows for both types of attention to have a multiplicative effect on the total level of attention. We elaborate further in the discussion the reasons for our representation of attention.

Top-down attention and individual differences

To operationalize the top-down channel of attention empirically, we calculate the average attention (Eq. (4)) allocated by individual j to the attribute s :

$$\bar{a}_{j,s} = a_s^{TD}(\beta_{j,s}) + \bar{a}_s^{BU} + \bar{v}_{j,s}. \quad (5)$$

If all agents observe the same alternatives, as will be the case in our experiments, average salience-driven attention $\bar{a}_s^{BU} = T^{-1} \sum_{t=1}^T a_s^{BU}(\sigma_t)$ is constant across agents, as it only depends on the choice alternatives. Hence, the only source of variation for average attention are the individual differences in $\beta_{j,s}$. Using a linear approximation (instead of assuming any linear representation for a^{TD} and a^{BU}), we show in Appendix 9.1 that we can approximate the decision weight $\omega_{j,s,t}$ defined in Equation (3) as:

$$\omega_{j,s,t} \approx \pi_{0,s} + \pi_{\bar{a},s} \bar{a}_{j,s} + \tilde{\epsilon}_{j,s,t}. \quad (6)$$

The parameter $\pi_{\bar{a},s}$ will be identifiable as long as agents' preferences have an effect on attention $\partial a^{TD}/\partial \beta_{j,s} \neq 0$. Conversely, if there is no top-down process, then individual differences in average attention should not have any correlation with the decision weights. Thus, average attention reflects top-down processes of attentional control.

Our statistical approach thus uses individual differences in average attention to approximate how individual characteristics affect choice via the top-down control of attention. This captures key elements of theories of rational inattention, which describe how optimizing behavior guides attention and choice (Caplin and Dean, 2015). For instance, rational inattention theories predict that lower information costs will make decisions more reactive to the available information. Our statistical model captures this reactivity through the $\pi_{\bar{a},s}$ parameter, which measures the correlation between average attention to an attribute and its importance in the decision. Thus, exogenous increases in information costs should lead to a smaller $\pi_{\bar{a},s}$.

Bottom-up attention and salience

We will now model the impact of salience: if an attribute is more salient, it will be attended more and therefore can be “over-weighted” in the decision process. The effects of salience on choice occur through bottom-up control of attention. For this, we allow the error term of the decision weight, $\epsilon_{j,s,t}$ (defined in equation 3), to depend on bottom-up attention:

$$\epsilon_{j,s,t} = \delta_{j,s} a_s^{BU}(\sigma_{s,t}) + \eta_{j,s,t}. \quad (7)$$

Here a_s^{BU} is the bottom-up control of attention of attribute s , which we assume to be an increasing function of salience and $\eta_{j,s,t}$, the left-over noise. By substituting Equation (7) into Equation (3), we can write the decision weight $\omega_{j,s,t}$ as:

$$\omega_{j,s,t} = \beta_{j,s} + \delta_{j,s} a_s^{BU}(\sigma_{s,t}) + \eta_{j,s,t} \quad (8)$$

We operationalize salience empirically by relating it to trial-by-trial deviations in attention. This *residual trial-wise attention* ($\tilde{a}_{j,s,t}$) is the difference between the allocated attention to attribute s on trial t and the average attention to that attribute. Thus, we can write

$$\begin{aligned}\tilde{a}_{j,s,t} &:= a_{j,s,t} - \bar{a}_{j,s} \\ &= a_s^{BU}(\sigma_{j,s,t}) - \bar{a}_s^{BU} + (\nu_{j,s,t} - \bar{\nu}_{j,s})\end{aligned}\tag{9}$$

where the last step uses Equations (4) and (5). If \bar{a}_s^{BU} is constant across participants, which will be the case in experiments (like ours) where all participants observe the same stimuli, the two sources of variation for $\tilde{a}_{j,s,t}$ are bottom-up attention in trial t ($a_s^{BU}(\sigma_t)$) and the zero-mean error term $\nu_{j,s,t}$. Similar to our approach with average attention, we can approximate the decision weights as:

$$\omega_{j,s,t} \approx \pi_{0,s} + \pi_{\bar{a},s}\bar{a}_{j,s} + \pi_{\tilde{a},s}\tilde{a}_{j,s,t} + \tilde{\eta}_{j,s,t},\tag{10}$$

where the parameter $\pi_{\tilde{a},s}$ is proportional to the marginal effect of the salience σ on attention a_s^{BU} (See Appendix 9.1 for a full derivation). Thus, if bottom-up control of attention is present, it is reflected in a correlation between decision weights and residual trial-wise attention. Conversely, if the decision process does not depend on bottom-up attention, then trial-wise attention should be uncorrelated with the decision weights.

Our statistical approach thus uses trial-wise deviations from average attention to a particular attribute to approximate how various contextual elements drive choice. In this way, we can incorporate predictions from models that focus on bottom-up control of attention during choice, such as the attentional drift diffusion model (aDDM; Krajbich et al., 2010) and salience-based utility models (Bordalo et al., 2012). These models predict that more salient attributes attract more visual attention. Thus, trials with highly salient attributes will generate higher trial-wise deviations from the average attention to those specific at-

tributes. If this salience-driven attention affects the decision process, as posited by theories of salience, our model will capture this effect via an increase in the estimated $\pi_{\bar{a},s}$ parameter.

Finally, it is worth noting a parallel between our approach and aDDM, which also captures bottom-up attentional processes. More specifically, Webb (2018) and Smith et al. (2019) show that the relative effect of attention on the diffusion process can be estimated by using a logistic RUM with the value of the attributes moderated by relative attention. Our model extends this insight to include also processes of top-down attention.

Disaggregating variation in attention

To understand the relative importance of top-down and bottom-up attention, we can ask how much they contribute to the variance in decision making. To quantify this, we can use the expression for the decision weights, Equation (10). By construction, residual trial-wise and average attention are orthogonal to each other. If we assume they are also independent from the residual error $\tilde{\eta}_{j,s,t}$, we can write the variance of the decision weights as:

$$Var(\omega_{j,s,t}) = \pi_{\bar{a},s}^2 Var(\bar{a}_{j,s}) + \pi_{\tilde{a},s}^2 Var(\tilde{a}_{j,s,t}) + Var(\tilde{\eta}_{j,s,t}) \quad (11)$$

Ideally, we would like to compare how much the average and residual trial-wise attention contribute to the variance in the decision weights. This measure would suggest how ‘important’ one process is relative to the other. Since we do not observe $(\tilde{\eta}_{j,s,t})$, we construct the ratio of the contributions to the variance:

$$R_s = \frac{\pi_{\bar{a},s}^2 Var(\bar{a}_{j,s})}{\pi_{\tilde{a},s}^2 Var(\tilde{a}_{j,s,t})} \quad (12)$$

Theoretical and empirical evidence suggests that the bottom-up process becomes more relevant when the decision is made under pressure, the stakes are low or the participants do

not have a clear idea of what to do or how to compare their options (Fehr and Rangel, 2011; Gottlieb et al., 2013). By comparing the relative contributions of the different sources of variance, our framework provides a way to evaluate these claims by comparing R_s across decision-making contexts.

4 Experimental Design

4.1 Participants

In total 99 participants took part in two experiments ($n_1 = 53$, $n_2 = 46$), which were identical except for small details (more on that below). Data from 8 participants were excluded, because of technical problems that occurred during data collection (5 in Exp.1 and 3 in Exp.2) due to wearing eye-tracker incompatible glasses or contact lenses ($n = 5$) and problems with recording the behavioural data ($n = 3$). One participant made the same decision in all trials, therefore their data was excluded. Partial data for one of two sessions was included for 3 more subjects (2 in Exp.1 and 1 in Exp.2), due to incomplete measurement of the visual data in one of the sessions (data loss of more than 75% due to calibration difficulties). The final data used for analysis therefore contains 91 participants (59 females, average age is 23.5 years).

Participants in both experiments were students from the University of Amsterdam, with no impaired or corrected vision. The recruitment was done via the website of the Behavioral Science Lab that houses the eye-trackers used in the current experiment (<https://www.lab.uva.nl/lab>). The participants signed an informed consent (available in the Appendix) and the experiments were approved by the FMG Ethics Committee of the University of Amsterdam.

4.2 Experimental Procedures

On the day of the experiment, participants performed the main task in a darkened testing room. This was done to reduce the effects of ambient light changes on pupil dilation. Jointly, the instructions, practice session and calibration procedures provided ample time to adjust to the background light in the experiment room. Eye movements made throughout the experiment were recorded using an EyeLink 1000 desk-mounted eye-tracker with a sampling rate of 500 Hz. To improve the accuracy of eye-tracking data collection, participants were asked to rest their heads on a chinrest to stabilize the head position and maintain a constant distance from the screen throughout the experiment. The stimuli were presented on a 22-inch screen with the resolution set to 1920×1080 pixels and a refresh rate of 60 Hz. At the start of the experiment and at the half-way point (after 80 trials) a 9-point calibration was performed to ensure proper calibration of the eye-tracker throughout the experiment.

4.3 Main Task

The main task in both experiments consisted of a series of 160 individual decisions involving risk. In each trial, participants were asked to accept or reject a mixed gamble with two equally likely outcomes. The outcomes were always a positive (“gains”) and a negative one (“losses”). Figure 1 shows the sequence of an example trial. At the beginning of the trial, participants were asked to focus on a fixation cross presented in the middle of the screen for a jittered period of time (300-1100ms). This ensured that in each decision period eye fixations started from the same central position and that attention was not biased towards a single location. Then the two potential outcomes appeared at each side of the screen, with the left stimulus located at $(x = 480\text{px}, y = 580\text{px})$, and the right one at $(x = 1420\text{px}, y = 580\text{px})$. This wide separation between lottery options along the x -dimension (of approximately 2.5° of visual angle) ensured that eye movement patterns can be well separated during the analysis stage (see Figure S1). The location of gains and

losses was counterbalanced, such that they had an equal chance of appearing on the left or right in each trial.

The participants were asked to press the Up-Key on the keyboard to accept the gamble or the Down-Key to reject it. Subjects were given a period of 5 seconds to make the decision. If the subject did not respond within those 5 seconds, a message appeared on the screen reminding participants to ‘Respond Faster’. In total, 47 of the 14,372 analysed trials exceeded the time limit; these ‘miss’ trials were excluded from the analysis. Participants were aware that if they did not respond within the 5-second period, they would receive the loss outcome of that trial in case it was selected at random at the end of the experiment. In experiment 2, the trial continued with a question of how confident the subject was about their decision, which was the only difference between the two experiments.

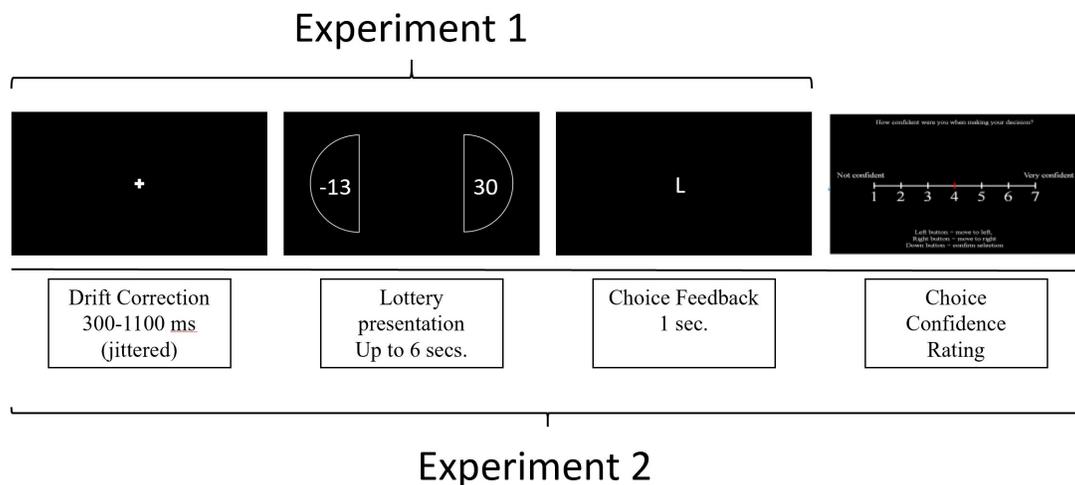


Figure 1: Example of Experimental Trial

Initially, a white fixation cross is shown for a random duration that is jittered between 300ms and 1100ms. The prospect is then presented. Participants then communicated their decision by pressing the up or down keys of the keyboard to accept or reject respectively. Feedback informed participants what option they had chosen before the next trial began in experiment 1. Experiment 2 differed only in that participants were asked to rate their confidence before the next trial.

The attributes presented on the left and right were pseudo-randomized, such that the

subject would never observe a loss or a gain more than three consecutive times on one side. The values of the Gains and Losses varied across trials. The gains fell between 20 to 38 ECU (experimental currency units) in steps of two units (10 cases). The losses ranged from -13 to -27 ECU in steps of two (8 cases). Gains and losses were independent from each other, and participants observed all possible combinations between gains and losses twice (80 trials per session in 2 sessions).

4.4 Incentives and Payment

Participants filled out a 30-minute online questionnaire consisting of a number of established Personality Questionnaires (e.g., ERQ, STAI, BIS-11) up to 1 day before the main experiment. The participants received €10 as a payment for completing the questionnaires. This amount served as an endowment for the main task to avoid the house money effect (Thaler and Johnson, 1990). Participants were informed that one of the 160 trials would be chosen at random. If the gamble was accepted on that trial, then the lottery would be resolved via a virtual coin flip. The outcome would be added to the initial endowment if it was a gain, or subtracted from the initial endowment if it was a loss. The ECUs were converted to € at a rate of $\frac{1}{5.4}$. In case the gamble was rejected, participants would receive the initial payment only. On average participants earned €10.80 and €10.94 in Experiment 1 and 2 respectively.

4.5 Eye-tracking Data Acquisition and Pre-processing

Fixation points were carefully calibrated using a 9-point calibration at two time points in the experiment (before the start of the experiment and after 80/160 completed trials). Furthermore, throughout the experiment, gain and loss attributes were clearly separated by presenting one attribute on the left and another on the right of the center. This clear separation of lottery attributes on the screen allowed us to define well-separated regions of interest and thereby to improve the identification of fixations. Next, using k-means,

we clustered the fixations along the horizontal axis representing fixation areas for left and right gamble attributes, and central fixation, which occurred only at the beginning of each trial. We ignore the vertical position for clustering, since all the stimuli were positioned at the same vertical location. This allowed us to discriminate between fixations for each outcome (left and right ROI) and central fixations (see Supplementary Figure B1). Finally, K -means clustering was performed for each session separately, as separate calibrations were performed for each session.

Table 1 shows the number fixations for each region of interest by their order of occurrence. A large majority of the first fixations are on the centre (90%), indicating that subjects followed task-instructions to focus on the fixation cross between trials. Most subsequent fixations go to the left first (68.9%), reflecting a commonly observed upper-left location bias (Orquin and Mueller Loose, 2013).

Fixation	Left	Right	Total
1	10,463	3,195	13,658
2	2,859	9,780	12,639
3	5,265	2,057	7,322
4	922	1,906	2,828
> 5	922	876	1,798
Total	20,431	17,814	38,245

Table 1: Number of fixations by order of Fixation and Region of Interest

We focus our analyses of the eye-tracking data on the dwell times, defined as the period participants fixate on a lottery attribute throughout one trial. We do this, because dwell times are the dominant measure of attention in the literature. Another measure, the number of saccades or switches of gaze between options, are less informative for our purpose. As shown in Table 1, the majority of trials do not contain more than three fixations, hence this number has little variation across trials and participants. We consider additional measures of attention in Section 7 which show similar results.

5 Hypotheses

Our experimental setup provides a natural application for the framework outlined in Section 3 with $J = 91$ and $T = 160$. Since all outcomes are equiprobable, the choice options in the experiment (gambles) have only two attributes, gains and losses, i.e. $s \in \{G, L\}$. When it comes to attention, we assume the agents' attention to be fully captured by their gaze patterns.² Hence, if an agent pays more attention to an attribute, this outcome should be observed for longer. From here onwards, we will sometimes refer to total dwell time (i.e. total time spent fixating on a stimulus) as attention.

We can now apply our results to generate testable hypotheses. Equation 6 implies that differences in preferences translate into differences in average attention, via the top-down channel.

Hypothesis 1 (Top-down attention). *A higher average attention of participant to an attribute (gains or losses) is associated with a higher decision weight for that attribute.*

Similarly, Equation 10 implies that differences in attribute salience translate into differences in residual trial-wise attention, via the bottom-up channel.

Hypothesis 2 (Bottom-up attention). *A higher deviation from average attention to an attribute (gain or loss) in trial t increases the decision weight for this attribute in trial t .*

Thus, our framework allows us to assess the relative importance of top-down and bottom-up attention by testing these two hypotheses. To do so, our main empirical exercise is to estimate the decision weights on the different attributes (gains and losses, $\pi_{\bar{a},s}$ and $\pi_{\bar{a},s}$ in our model) and to test whether there is an interaction with the two different types of attention.

²Under the “eye-mind assumption” the current focus of attention is what is being processed (Just and Carpenter, 1980). However, it is also possible in some contexts that attention diverges from the visual focus (Egeth and Yantis, 1997).

6 Results

Our main results are presented in Table 2³. The table presents logit regression models, in which we regress the binary acceptance decision on the lottery attributes (gains x_G and losses x_L), as well as interactions with average individual attention (\bar{a}_s^{TD}) and trial-wise deviations in attention (\tilde{a}_s^{BU}). Each model includes individual fixed effects to account for differences in the value of the outside option (\bar{u}_j).⁴

The fourth column of Table 2 presents the simplest model including only the coefficients for the two attributes x_G and x_L . These reflect the parameters $\pi_{0,G}$ and $\pi_{0,L}$ that contribute to the decision weights (see equation 10). Both are statistically significant at $p < 0.001$, and the intercept for losses is larger than that for gains ($\pi_{0,G} - |\pi_{0,L}| = -.083$, $p\text{-value} < 0.001$). This result suggests that participants are loss averse on average, which we will analyze in more detail below.

To test our model predictions, we focus on the full model shown in column 1, which includes interaction terms with both types of attention. We find that the interactions of these coefficients with the attentional measures are statistically significant for both attributes and for both types of attention. When we evaluate the differences of the attentional impact across gain and loss attributes with a Wald test, we find that differences (in absolute value) are not significant for average attention ($\pi_{\bar{a},G} - |\pi_{\bar{a},L}| = .039$, $p\text{-value} = 0.504$), but we find significant differences for trial-wise attention ($\pi_{\tilde{a},G} - |\pi_{\tilde{a},L}| = -.024$, $p\text{-value} < 0.001$). This indicates that decisions weights are affected significantly more by salience in the domain of losses compared to gains.

To evaluate the importance of the attentional channels in predicting decisions, we compare the model fit of the full model with various benchmark models in columns 2-4 of Table 2 using standard criteria of model fit shown at the bottom of the table (i.e. the

³Note that we report the results from the combined dataset for simplicity, because results are highly similar when estimating each model separately for experiments 1 and 2. These results are reported in Supplementary Table A2

⁴We use the package for fixed effect logits from Cruz-Gonzalez et al. (2017) to analytically correct for the incidental parameter problem in discrete choices (Katz, 2001; Coupé, 2005; Arellano and Hahn, 2006).

Table 2: Estimations for Decisions

	(1)	(2)	(3)	(4)
	Full Model	Avg. Attention	Res. Attention	Constant Weights
$x_{G,t}$	0.190*** (0.022)	0.189*** (0.022)	0.352*** (0.008)	0.352*** (0.008)
$x_{G,t} \times \bar{a}_{j,G}$	0.353*** (0.048)	0.356*** (0.047)		
$x_{G,t} \times \tilde{a}_{j,G,t}$	0.008*** (0.003)		0.008*** (0.003)	
$ x_{L,t} $	-0.301*** (0.028)	-0.306*** (0.028)	-0.433*** (0.010)	-0.435*** (0.010)
$ x_{L,t} \times \bar{a}_{j,L}$	-0.315*** (0.065)	-0.308*** (0.065)		
$ x_{L,t} \times \tilde{a}_{j,L,t}$	-0.031*** (0.005)		-0.030*** (0.005)	
N	13057	13057	13057	13057
AIC	7293.469	7334.256	7373.802	7415.865
BIC	7338.332	7364.164	7403.710	7430.819
LL	-3640.735	-3663.128	-3682.901	-3705.932

Results of the logistic estimations of binary decisions on lottery attributes and attention. All models include individual fixed effects (91 individuals). Note that we excluded observations with only one fixation (8.8% of all trials) as this indicates that participants did not fully consider all choice options on a given trial. The loss amounts were entered as absolute values for easier interpretation of the weights. The error terms are estimated with jackknife resampling (in parentheses). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC) and the Log Likelihood (LL)). In column 2, we leave out the interaction with trial-wise attention. In line with the small interaction effect in column 1, removing trial-wise attention does not have a large impact on the other coefficients, although it does worsen the model fit. By contrast, when we drop the interaction with average attention from the model in column 3, we find that the coefficients of the attributes almost double for gains, and rise by about 40% for losses. This indicates that a large part of the variation in decision weights can be attributed to individual differences in attention. Finally, column 4 shows the model without attention, which is the worst performer of the four models in terms of model fit, further underlining the value of incorporating attention when predicting decisions.

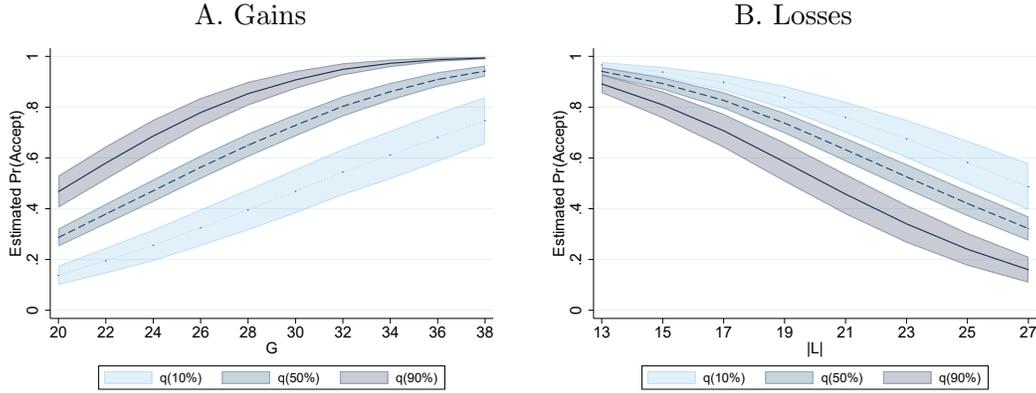


Figure 2: Acceptance probability by attention quantiles and gain/loss magnitude. The estimated probability of accepting the lottery (vertical axis) conditional on the outcome values (Left panel for gains, right panel for losses) and average attention to the same outcomes. The lines reflect different levels of average attention, which are the sample quantiles 10% (dotted), 50% (dashed) and 90% (solid). The predictions are presented with their 95% confidence intervals.

The value of explicitly modeling attention is further illustrated in Figure 2, which shows the impact of average attention on decision weights, by graphing the changes in the probability of accepting the lottery for the 10%, 50% and 90% quantiles based on the average attention distribution. While greater attention to gains is associated with a

greater likelihood and greater attention to losses with a decreased likelihood of accepting the gamble, the extent of this effect depends on stake size. In the domain of gains (left panel), the difference between the 10% and 90% quantiles decreases with increasing gain amounts, starting from a difference of 32.9% for small gain amounts ($x_G=20$) and reaching a difference of 24.6% for large gain amounts ($x_G=38$). In the domain of losses (right panel), the difference between the 10% and 90% quantiles increases from 7.4% for small loss amounts ($|x_L| = 13$) to 32.7% for large amounts ($|x_L| = 27$). The analogous plot for differences in trial-wise attention is given in Supplementary Figure (B2). Although statistically significant, the differences are small and not clearly visible in the plot.

Finally, Table 2 allows us to compare the relative importance of the two types of attention in decision making. We calculate the ratios of their contribution to the variance for both gains and losses (as described in Equation 11). We use the sample variances for the average and residual attention and the estimates from column 1 in Table 2 to construct the ratio of the contributions to the variance R_s (see Equation 12). Relative to trial-wise attention, the contribution of average attention to the variance of the weights is significantly larger for both losses and gains ($\ln(R_l) = 3.121, p\text{-value} < 0.001$ and $\ln(R_g) = 6.159, p\text{-value} < 0.001$). This result confirms that in the context of our experiment, the main driver of changes in decision weights is average attention, reflected by individual differences in attentional patterns, and not trial-wise residual attention that reflects attribute salience.

Result 1. *Both average attention and residual trial-wise attention correlate significantly with the decision weights for the different attributes. The effect size is much larger for average attention, indicating that in the context of our task, top-down attention contributes more to the decision process.*

Attention and Loss Aversion

To more specifically characterize the role of top-down attention in decision biases, we investigate the relationship between attention and loss aversion on an individual level. Loss

aversion refers to a preference for avoiding losses rather than obtaining gains (Kahneman and Tversky, 1979). Since loss aversion is an agent-specific characteristic, the model predicts that it should be correlated with average individual attention to losses relative to gains. We define and compare individual levels of loss aversion in our sample and investigate whether this is driven by the participants’ average attention, and by the decision weights predicted by the attention of participants.

First, to obtain simple measure of loss aversion, we estimate our benchmark model without attention (column 4 in Table 2) for each individual. This results in an estimate of two decision weights ($\hat{\omega}_{j,s}$) for each individual, based on the behavioral data only.⁵ We then use these weights to calculate each individual’s level of loss aversion λ as

$$\lambda_j := \left| \frac{\hat{\omega}_{j,L}}{\hat{\omega}_{j,G}} \right|. \quad (13)$$

Second, we compute an agent-specific measure reflecting the relative allocation of attention to losses compared to gains as $\Delta\bar{a}_j = \bar{a}_{j,L}/\bar{a}_{j,G}$.

The left panel of Figure 3 shows the resulting relationship between loss aversion λ_j and $\Delta\bar{a}_j$. The model reported at the top of panel A includes an intercept and shows that $\Delta\bar{a}_j$ is significantly associated with λ_j ($p < 0.001$). This underlines that loss aversion (as measured by attribute weights) is indeed correlated with relative attention to losses, as one would expect if loss aversion is an individual preference.

To further probe the impact of attention within the context of our structural model, we correlate λ_j with an alternative measure of loss aversion that is predicted by the attentional patterns. To calculate the latter, we use the attentional data to predict the weights $\omega_{j,G,t}$ and $\omega_{j,L,t}$ for every trial and every individual. We do so using Equation (10) from our model, where the π parameters are based on the interaction terms of our full model (Table 2 column (1)). This allows us to have trial and individual specific measures of attention-

⁵We exclude three cases for which λ could not be estimated at the agent level: two cases with decision weights that had opposite signs, and an additional case with an insufficient number of observations.

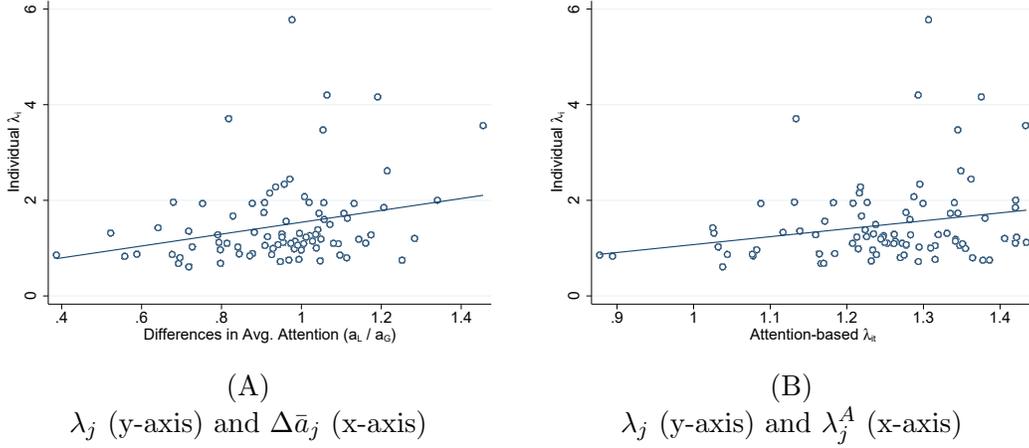


Figure 3: Correlation between loss aversion and attention measures. The figures present the relationship between the differences in average attention to losses and gains, the agent-specific and attention-based levels of Loss Aversion. Agent-specific loss aversion, shown on the y-axis in (A) and (B), is the ratio between the individual weights ($\lambda_j = -\omega_{j,L}/\omega_{j,G}$). The difference in average attention, shown in the x-axis in (A) is defined as $\Delta \bar{a}_j = \bar{a}_{j,L} - \bar{a}_{j,G}$. λ_{jt}^A , shown on the x-axis in (B), is the median of the ratio of the attention-based decision weights ($\lambda_{j,t}^A = -\omega_{j,L,t}/\omega_{j,G,t}$). The correlations in (A) and (B) are significant at 0.1%.

based loss aversion $\lambda_{j,t}^A := \left| \frac{\omega_{j,L,t}}{\omega_{j,G,t}} \right|$. These individual loss aversion measures fluctuates due to the variability of attention over trials, so we take the median over trials for each individual to obtain individual weights:

$$\lambda_j^A := \text{median} \{ \lambda_{j,t}^A \}_{t=1}^T \quad (14)$$

The right panel of Figure 3 shows the significant ($p < 0.001$) correlation between the individual level estimates of loss aversion (λ_j) and those predicted from attention only (λ_j^A). This shows that individual differences in loss aversion are partially captured by attention-based proxies.

Result 2. *We find a modest but statistically significant correlation between a behavioral measure of individual loss aversion and attention-based proxies of loss aversion.*

These results complement those in Pachur et al. (2018), who also find a relationship between attentional patterns and loss aversion parameters.

7 Discussion

In this section we report additional results concerning the robustness of our model, the causal effect of attention, the validity of our model assumptions, and the applicability of our method to other experiments. Briefly, our model yields robust results when using alternative measures of attention, assumptions are not violated and our model can be flexibly applied across different experimental contexts and requires a relatively low trial number.

7.1 Robustness to other measures of attention

While dwelling times are a typical measure for attention, our theoretical framework is agnostic about which measure best captures attention. Therefore, Table A3 (see Appendix)

repeats our analysis using another set of measures, namely relative dwell times to the two attributes (column 2), the number of fixations to each attribute (column 3), and a logarithmic transformation of dwell time (column 4). For convenience column 1 repeats the results of the absolute dwell time model in Table 2.

As the estimates in Table A3 demonstrate, the results are qualitatively and quantitatively robust to the use of different measures reflecting attention to specific attributes. First, all the interactions between the attributes and the attentional measures remain statistically significant, although the relative dwell time and fixation measures show a slight reduction in p-values for the trial-wise attention measures in the gain domain. Furthermore, in all specifications, trial-wise attention has a much smaller impact than average individual attention. These robustness checks further underline that attention matters, and that bottom-up attention has a weaker impact in our experiment than top-down attention in this context.

7.2 Validity of assumptions

Our interpretation of residual trial-wise variations as reflecting bottom-up attention and average attention measuring top-down attention rests on a number of assumptions. Perhaps the most fundamental assumption underlying our framework is that the influences of top-down and bottom-up attention on the decision process are additively separable. Above, we cited experimental evidence in support of this assumption (e.g., Pinto et al., 2013). This assumption would be violated by interactive effects between bottom-up and top-down attention. These may occur if top-down attention to an attribute raises the impact of salience variations in that attribute.

To test for this type of violation, we estimate the interaction effect of average and residual trial-wise attention. Table A4 in the appendix shows the results of regressions including these interactions in both the loss and gain domain (column 2), the gain domain only (column 3) and the loss domain only (column 4). Column 1 reproduces the original

estimates from Table 2 for ease of comparison. Results show a low but significant interaction effect for losses. This interaction effect, however, is not robust to other model specifications (see column 4). More importantly, the inclusion of both interaction terms does not cause the estimates on either type of attention to change substantially or become insignificant. We conclude that any interaction effects between average and trial-wise attention, if they exist, are relatively small and do not change the estimates of the individual effects of the two attention channels.

For residual trial-wise attention to reflect bottom-up attention, we assume that the size of the attributes does not affect top-down attention. If this is violated, our empirical model would then be attributing a percentage of the variability in the top-down process due to attribute size to residual trial-wise attention. Since we observe greater coefficients for average attention compared to residual trial-wise attention, we would be over-weighting the bottom-up effect within our model (and underestimate the impact of top-down attention). To test this assumption, we assess whether more attention is paid to larger gain and loss attributes. Appendix Table A5 shows the estimations for the determinants of attention. The results show that attention is significantly correlated with gain size. Even though these effects are significant, their impact is relatively small; increasing the gains by one unit leads to an estimated increase in dwell times by approximately 2ms out of an average 448ms for gains and 3ms out of 417ms for losses. Therefore, even if these effects reflect top-down attention to gain and loss magnitude, the effects on our estimations would be negligible. Moreover, these effects are only significant in the gain domain, but not in the loss domain.

Finally, we make the assumption that preferences are stable, and do not change over trials. This is a natural assumption in our experiment, given that no feedback about decision outcomes is provided until the end of the experiment, thereby preventing learning. However, we can test this assumption by assessing the evolution of attention over trials. Table A5 shows that dwell times for both gains and losses decrease slightly, but significantly,

throughout the experiment. This effect likely reflects increased familiarity with the general task setup as the experiment progresses. This is supported by the results reported in column 3 in Supplementary Table A5, which show that participants also become slightly faster at the task, and this reduction in average response times per trial will consequently reduce the dwell times for both attributes throughout the experiment. In Table A6 we control for potential changes in the decision weights over the course of the experiment. The results show that none of our original predictions of attention change. Nonetheless, we observe a small, but significant effect of trials on the decision weight for gains, but not for losses, indicating that the weighting of gains increased throughout the experiment.

The assumptions discussed in this section indicate the knowledge frontier in attention research, and are subject to wider discussion in the psychology and neuroscience literature. For instance, Awh et al. (2012) argue that some contextual elements, like rewards, may trigger top-down attention, because people have built up mental associations with them. Future research that clarifies the impact of such interactions between attentional and reward processes will inform the foundation of our model.

7.3 Applicability of the method

Like any method, ours has a number of strengths and limitations. One strength is that it is quite general: it does not depend on a specific number of attributes, specific measures of salience or particular measures of attention. It is thus potentially applicable to a large number of experimental approaches and datasets.

One limitation is that the model requires a sufficient number of trials in order to produce a reliable estimate for the impact of trial-wise variance. To get a sense of how many trials are needed for a stable estimate, Supplemental Figure B3 in the appendix, shows how estimates and the model log-likelihood change when progressively more trials are included in the estimation. As is apparent from these figures, it appears that about 40 trials are sufficient to obtain relatively stable estimates, a number that is below the typical number

of trials in the attention literature in psychology and neuroscience. The model proposed here, may therefore yield relatively robust estimates with a relatively low number of trials (given that a sufficiently wide range of attributes is included within these 40 trials).

When it comes to experimental manipulations, our method can be easily applied to experiments that incorporate within-subject manipulations of salience: these manipulations lead to variation in trial-wise attention, which the model correctly attributes to bottom-up processes. The framework thus provides a useful method to verify if the salience manipulations did indeed have an impact on the importance of bottom-up attention.

7.4 Causality

As we discussed in Section 2, several papers have tried to measure the causal impact of attention. Our framework shows that the causal effect of attention is not as straightforward as it seems at first sight, as there may be different causal pathways. On the one hand, bottom-up attention can exert a causal influence on decision making by directing attention to salient features of the choice context. On the other hand, top-down attention exerts its impact on attention by focusing on features that reflect the agent’s preferences, thereby allowing the agent to translate his or her preferences and beliefs into relevant decisions.

In the context of decision-making it appears that top-down attention is harder to manipulate than bottom-up attention, because it is driven by personal characteristics, goals and preferences, rather than contextual variables that are under direct control of the experimenter. Thus, the majority of papers we discuss in the literature section have explored the causal effects of bottom-up attention through the manipulation of salience. Causal manipulations of top-down attention are rare in the context of judgement and decision making research (e.g., Ghaffari and Fiedler, 2018; Hausfeld et al., 2021). Other common experimental manipulations, such as the use of time quotas for observing different attributes, are likely to affect both bottom-up and top-down processes simultaneously. Future research can build on the correlational approach developed here, by independently manipulating

top-down and bottom up attention, and measuring the relative impact of both.

8 Conclusions

For good reasons, there is an increasing focus on attention in economic theory. One of the most fundamental open questions in this regard is to what extent attention is driven by agents' characteristics and to what extent it is determined by the context. In this paper, we have provided a basic framework to evaluate this question. Our experiments show that both bottom-up and top-down attention contribute to choices under risk. Aggregate differences in attention also correlate with an individual loss aversion parameter, underlining their relation with the agent's specific goals.

Among economists, there is some expectation that attention can be a “unifying” variable that ties together hitherto separate phenomena (Gabaix, 2019). We show here that this may indeed be the case as attention explains parts of the variance underlying loss aversion. Additional studies investigating the role of attention in other decision biases are important to further delineate the role of attention in decision biases. Similarly, the potential of attention and eye-tracking are attracting scholars from a wide range of research fields, such as management and organization (Meißner and Oll, 2019). The framework we propose here can be flexibly applied to different experimental contexts and can help answer a number of questions that are crucial to fulfill this promise of attention research. For instance, how does the influence of bottom-up vs. top-down attention vary across environments? How do various aspects of salience affect bottom-up attention and the occurrence of behavioral biases? How do individual differences in attention correlate with personal characteristics and decision parameters? Answering these questions will be valuable to both theorists and policy makers alike. More generally, our approach demonstrates the fruitful interaction between cognitive (neuro-)science and economic analysis.

References

- Amasino, D. R., Sullivan, N. J., Kranton, R. E., and Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature Human Behaviour*, 3(4):383–392.
- Anderson, B. A., Laurent, P. A., and Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences*, 108(25):10367–10371.
- Arellano, M. and Hahn, J. (2006). A likelihood-based approximate solution to the incidental parameter problem in dynamic nonlinear models with multiple effects. *Unpublished manuscript*.
- Armel, K. C., Beaumel, A., and Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision making*, 3(5):396–403.
- Atalay, A. S., Bodur, H. O., and Rasolofoarison, D. (2012). Shining in the center: Central gaze cascade effect on product choice. *Journal of Consumer Research*, 39(4):848–866.
- Awh, E., Belopolsky, A. V., and Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in cognitive sciences*, 16(8):437–443.
- Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6):1437–75.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

- Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.
- Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201–215.
- Coupé, T. (2005). Bias in conditional and unconditional fixed effects logit estimation: A correction. *Political Analysis*, pages 292–295.
- Cruz-Gonzalez, M., Fernandez-Val, I., and Weidner, M. (2017). LOGITFE: Stata module to compute analytical and jackknife bias corrections for fixed effects estimators of panel logit models with individual and time effects.
- Dertwinkel-Kalt, M., Köhler, K., Lange, M. R., and Wenzel, T. (2017). Demand shifts due to salience effects: Experimental evidence. *Journal of the European Economic Association*, 15(3):626–653.
- Dertwinkel-Kalt, M. and Köster, M. (2020). Salience and skewness preferences. *Journal of the European Economic Association*, 18(5):2057–2107.
- Dessein, W., Galeotti, A., and Santos, T. (2016). Rational inattention and organizational focus. *American Economic Review*, 106(6):1522–36.
- Egeth, H. E. and Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual review of psychology*, 48(1):269–297.
- Fehr, E. and Rangel, A. (2011). Neuroeconomic foundations of economic Choice—Recent advances. *Journal of Economic Perspectives*, 25(4):3–30.
- Fiedler, S. and Glöckner, A. (2012). The Dynamics of Decision Making in Risky Choice: An Eye-Tracking Analysis. *Frontiers in Psychology*, 3.
- Fisher, G. (2021). Intertemporal choices are causally influenced by fluctuations in visual attention. *Management Science*.

- Fosgerau, M., Sethi, R., and Weibull, J. W. (2020). Categorical Screening with Rational Inattention. *SSRN Electronic Journal*.
- Fudenberg, D., Strack, P., and Strzalecki, T. (2018). Speed, Accuracy, and the Optimal Timing of Choices. *American Economic Review*, 108(12):3651–3684.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.
- Gabaix, X. (2019). Chapter 4 - Behavioral inattention. In Bernheim, B. D., DellaVigna, S., and Laibson, D., editors, *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2 of *Handbook of Behavioral Economics - Foundations and Applications 2*, pages 261–343. North-Holland.
- Ghaffari, M. and Fiedler, S. (2018). The power of attention: Using eye gaze to predict other-regarding and moral choices. *Psychological science*, 29(11):1878–1889.
- Gluth, S., Kern, N., Kortmann, M., and Vitali, C. L. (2020). Value-based attention but not divisive normalization influences decisions with multiple alternatives. *Nature Human Behaviour*, pages 1–12.
- Gluth, S., Spektor, M. S., and Rieskamp, J. (2018). Value-based attentional capture affects multi-alternative decision making. *Elife*, 7:e39659.
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., and Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11):585–593.
- Hausfeld, J., von Hesler, K., and Goldlücke, S. (2021). Strategic gaze: an interactive eye-tracking study. *Experimental economics*, 24(1):177–205.
- Johnson, E. J. and Schkade, D. A. (1989). Bias in utility assessments: Further evidence and explanations. *Management science*, 35(4):406–424.

- Just, M. A. and Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological review*, 87(4):329.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-291. *Econometrica*.
- Katz, E. (2001). Bias in conditional and unconditional fixed effects logit estimation. *Political Analysis*, pages 379–384.
- Kőszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly journal of economics*, 128(1):53–104.
- 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.
- Krajbich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, 3:193.
- Lim, S.-L., O’Doherty, J. P., and Rangel, A. (2011). The Decision Value Computations in the vmPFC and Striatum Use a Relative Value Code That is Guided by Visual Attention. *Journal of Neuroscience*, 31(37):13214–13223.
- Mackowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.
- Maćkowiak, B. and Wiederholt, M. (2015). Business cycle dynamics under rational inattention. *The Review of Economic Studies*, 82(4):1502–1532.
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.
- McFadden, D. (1980). Econometric models for probabilistic choice among products. *Journal of Business*, pages S13–S29.

- Meißner, M. and Oll, J. (2019). The promise of eye-tracking methodology in organizational research: A taxonomy, review, and future avenues. *Organizational Research Methods*, 22(2):590–617.
- Navalpakkam, V., Koch, C., Rangel, A., and Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences*, 107(11):5232–5237.
- Orquin, J. L. and Lagerkvist, C. J. (2015). Effects of salience are both short- and long-lived. *Acta Psychologica*, 160:69–76.
- Orquin, J. L. and Mueller Loose, S. (2013). Attention and choice: A review on eye movements in decision making. *Acta Psychologica*, 144(1):190–206.
- Pachur, T., Schulte-Mecklenbeck, M., Murphy, R. O., and Hertwig, R. (2018). Prospect theory reflects selective allocation of attention. *Journal of Experimental Psychology: General*, 147(2):147–169.
- Pärnamets, P., Johansson, P., Hall, L., Balkenius, C., Spivey, M. J., and Richardson, D. C. (2015). Biasing moral decisions by exploiting the dynamics of eye gaze. *Proceedings of the National Academy of Sciences*, 112(13):4170–4175.
- Peng, L. and Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80(3):563–602.
- Pinto, Y., van der Leij, A. R., Sligte, I. G., Lamme, V. A., and Scholte, H. S. (2013). Bottom-up and top-down attention are independent. *Journal of vision*, 13(3):16–16.
- Polonio, L., Di Guida, S., and Coricelli, G. (2015). Strategic sophistication and attention in games: An eye-tracking study. *Games and Economic Behavior*, 94:80–96.
- Posner, M. I. (2011). *Cognitive neuroscience of attention*. Guilford Press.

- Posner, M. I., Snyder, C. R., and Davidson, B. J. (1980). Attention and the detection of signals. *Journal of experimental psychology: General*, 109(2):160.
- Ratcliff, R. (1978). A Theory of Memory Retrieval. *American Psychological Association*, 85(2):59.
- Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics*, 53(8):1761–1800.
- Reutskaja, E., Nagel, R., Camerer, C. F., and Rangel, A. (2011). Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study. *American Economic Review*, 101(2):900–926.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, page 26.
- Sims, C. A. (2010). Rational inattention and monetary economics. In *Handbook of Monetary Economics*, volume 3, pages 155–181. Elsevier.
- Smith, S. M., Krajbich, I., and Webb, R. (2019). Estimating the dynamic role of attention via random utility. *Journal of the Economic Science Association*, 5(1):97–111.
- Thaler, R. H. and Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, 36(6):643–660.
- 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.
- Ungerleider, L. G. and Kastner, S. (2000). Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, 23(1):315–341.
- Webb, R. (2018). The (neural) dynamics of stochastic choice. *Management Science*, 65(1):230–255.

9 Appendix

9.1 Reformulation of Attention-based utility

In the most general form (including the effect of salience), the equations for attention ($a_{j,s,t}$, defined in equation 4) and the decision weights ($\omega_{j,s,t}$, defined in equation 8) were given by:

$$\begin{aligned} a_{j,s,t} &= a_s^{TD}(\beta_{j,s}) + a_s^{BU}(\vec{\sigma}_t) + \nu_{j,s,t} \\ \omega_{j,s,t} &= \beta_{j,s} + g(a_s^{BU}(\vec{\sigma}_t)) + \eta_{j,s,t} \end{aligned}$$

We decomposed attention into two measures, average attention ($\bar{a}_{j,s}$, defined in equation 5) and residual trial-wise attention ($\tilde{a}_{j,s,t}$, defined in equation 9). Average attention, by aggregating over all trials, is not affected asymptotically by the trial-wise effects of salience. On the other hand, the residual trial-wise attention captures only the differences in bottom-up control of attention.

$$\begin{aligned} \bar{a}_{j,s} &= a_s^{TD}(\beta_{j,s}) + \bar{a}_s^{BU} + \bar{\nu}_{j,s} \\ \tilde{a}_{j,s,t} &= a_s^{BU}(\sigma_{j,s,t}) - \bar{a}_s^{BU} + (\nu_{j,s,t} - \bar{\nu}_{j,s}) \end{aligned}$$

Now, we proceed to partially differentiate $\omega_{j,s,t}$, $\bar{a}_{j,s}$ and $a'_{j,s,t}$:

$$d\omega_{j,s,t} = d\beta_{j,s} + \frac{\partial g}{\partial a^{BU}} \nabla_{\sigma} a^{BU} d\vec{\sigma}_t + d\eta_{j,s,t} \quad (15)$$

$$d\bar{a}_{j,s} = \frac{\partial a^{TD}}{\partial \beta} d\beta_{j,s} + o_P(1) \quad (16)$$

$$da'_{j,s,t} = \frac{\partial g}{\partial a^{BU}} \nabla_{\sigma} a^{BU} d\vec{\sigma}_t + d\nu_{j,s,t} + o_P(1) \quad (17)$$

Where $o_P(1)$ is a residual term that converges in probability to 0 at a rate of order

$1/T$. If $\frac{\partial a^{TD}}{\partial \beta} \neq 0$, we can replace $d\beta_{j,s}$ from equation (16) to (15). Similarly, we can take the overall effect of salience $\frac{\partial g}{\partial a^{BU}} \nabla_{\sigma} a^{BU} d\vec{\sigma}_t$ from equation (17) to (15). Then we rewrite the partial differential and do a linear approximation of $\omega_{j,s,t}$:

$$d\omega_{j,s,t} = \left(\frac{\partial a^{TD}}{\partial \beta} \right)^{-1} d\bar{a}_{j,s} + \frac{\partial g}{\partial a^{BU}} d\tilde{a}_{j,s,t} + o_P(T) - d\nu_{j,s,t} + d\eta_{j,s,t}$$

$$\omega_{j,s,t} \approx \pi_0 + \pi_{\bar{a},s} d\bar{a}_{j,s} + \pi_{\tilde{a},s} d\tilde{a}_{j,s,t} + \tilde{\eta}_{j,s,t}$$

Where $\pi_{\bar{a},s} = \left(\frac{\partial a^{TD}}{\partial \beta} \right)^{-1}$ and $\pi_{\tilde{a},s} = \frac{\partial g}{\partial a^{BU}}$. The error term $\tilde{\eta}_{j,s,t} = \eta_{j,s,t} - \nu_{j,s,t} + o_P(1)$ expected value converges to 0 when T increases.

9.2 Robustness Check - Estimations by Experiment

The estimations in Table A2 assess whether decision weights change or remain consistent across the two separate experiments. Column 1 reflects our original estimation reported in Table 2 that contains the data from both experiments (91 participants). Columns 2 and 3 reflect estimations based on the data recorded for experiments 1 (39 participants) and 2 (52 participants) respectively. The estimations include individual fixed effects and the error terms were estimated with jackknife resampling (in parentheses).

Table A2: Estimations for Decisions (by Experiment)

	(1)	(2)	(3)
	Both	Exp. 1	Exp. 2
$x_{G,t}$	0.190*** (0.022)	0.035 (0.041)	0.163*** (0.031)
$x_{G,t} \times \bar{a}_{j,G}$	0.353*** (0.048)	0.664*** (0.107)	0.360*** (0.059)
$x_{G,t} \times \tilde{a}_{j,G,t}$	0.008*** (0.003)	0.010* (0.005)	0.004 (0.003)
$ x_{L,t} $	-0.301*** (0.028)	-0.174*** (0.043)	-0.260*** (0.041)
$ x_{L,t} \times \bar{a}_{j,L}$	-0.315*** (0.065)	-0.620*** (0.111)	-0.338*** (0.088)
$ x_{L,t} \times \tilde{a}_{j,L,t}$	-0.031*** (0.005)	-0.016* (0.009)	-0.034*** (0.006)
N	13057.000	5426.000	7631.000
AIC	7293.469	3014.678	4152.059
BIC	7338.332	3054.272	4193.699
LL	-3640.735	-1501.339	-2070.029

Standard errors in parentheses

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

Table A3: Estimations for Decisions

	(1)	(2)	(3)	(4)
	Dwell-Time	Relative Time	Fixations	ln(D. Time)
$x_{G,t}$	0.190*** (0.022)	0.005 (0.040)	0.042 (0.040)	0.547*** (0.024)
$x_{G,t} \times \bar{a}_{j,G}$	0.353*** (0.048)	1.118*** (0.127)	0.228*** (0.030)	0.183*** (0.020)
$x_{G,t} \times \tilde{a}_{j,G,t}$	0.008*** (0.003)	0.015* (0.009)	0.003* (0.002)	0.006*** (0.002)
$ x_{L,t} $	-0.301*** (0.028)	-0.175*** (0.048)	-0.176*** (0.051)	-0.618*** (0.032)
$ x_{L,t} \times \bar{a}_{j,L}$	-0.315*** (0.065)	-0.897*** (0.162)	-0.191*** (0.038)	-0.166*** (0.026)
$ x_{L,t} \times \tilde{a}_{j,L,t}$	-0.031*** (0.005)	-0.056*** (0.014)	-0.010*** (0.003)	-0.014*** (0.003)
N	13057	13057	13057	13057
AIC	7293.469	7267.744	7329.232	7247.709
BIC	7338.332	7312.607	7374.095	7292.572
LL	-3640.735	-3627.872	-3658.616	-3617.855

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

9.3 Robustness Check - Alternative attention variables

Table A3 shows our model estimations with different measures for attention. The dependent variable is the participants' decisions. Column 1 shows our chosen variable, total dwell time (Estimation is identical to column 1 in Table 2). The model represented in Column 2 replaces absolute dwell time by relative dwell time that depends on the response time on a given trial (DT/RT). Finally, the model in column 3 uses the number of fixations as attentional proxy, while the model in column 4 uses the logarithmic transform of total dwell time. These estimations use fixed effects and the error terms were estimated with jackknife resampling (in parentheses).

Table A4: Estimations with interacting effects

	(1)	(2)	(3)	(4)
$x_{G,t}$	0.190*** (0.022)	0.188*** (0.022)	0.188*** (0.022)	0.190*** (0.022)
$\bar{a}_{j,G} \times x_{G,t}$	0.353*** (0.048)	0.358*** (0.048)	0.357*** (0.048)	0.354*** (0.048)
$\tilde{a}_{j,G,t} \times x_{G,t}$	0.008*** (0.003)	0.020** (0.010)	0.019* (0.010)	0.007*** (0.003)
$\bar{a}_{j,G} \times \tilde{a}_{j,G,t} \times x_{G,t}$		-0.018 (0.018)	-0.015 (0.018)	
$ x_{L,t} $	-0.301*** (0.028)	-0.299*** (0.028)	-0.300*** (0.028)	-0.300*** (0.028)
$\bar{a}_{j,L} \times x_{L,t} $	-0.315*** (0.065)	-0.316*** (0.065)	-0.316*** (0.065)	-0.314*** (0.065)
$\tilde{a}_{j,L,t} \times x_{L,t} $	-0.031*** (0.005)	-0.065*** (0.019)	-0.031*** (0.005)	-0.061*** (0.019)
$\bar{a}_{j,L} \times \tilde{a}_{j,L,t} \times x_{L,t} $		0.064* (0.035)		0.056 (0.034)
N	13057	13057	13057	13057
AIC	7293.469	7287.432	7290.206	7291.850
BIC	7338.332	7347.248	7342.545	7344.189
LL	-3640.735	-3635.716	-3638.103	-3638.925

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

9.4 Robustness Check - Interaction effects between average attention and residual trial-wise attention

The estimations in Table A4 test for possible non-linear interactions between average and residual trial-wise attention. Column 1 shows our specification reported in the main text, column 2 includes the non-linear moderation of average and residual attention for both the gains and losses. Columns 3 and 4 include the non-linear moderation only for the gains and for the losses respectively. The estimations include individual fixed effects and the error terms were estimated with jackknife resampling (in parentheses).

Table A5: Estimations for Attention

	(1)	(2)	(3)
	DT Gains	DT Losses	RT
$x_{G,t}$	0.003*** (0.001)	0.002** (0.001)	0.006*** (0.002)
$ x_{L,t} $	-0.002 (0.001)	0.001 (0.001)	-0.002 (0.002)
L. Left	0.063*** (0.018)	-0.062*** (0.017)	0.015 (0.026)
L. First	0.018 (0.016)	-0.022 (0.015)	-0.006 (0.024)
L. Left \times LossFirst=1	0.021 (0.020)	0.039* (0.021)	0.000 (0.041)
L. Last	-0.167*** (0.011)	0.180*** (0.010)	-0.054*** (0.013)
t	-0.001*** (0.000)	-0.000*** (0.000)	-0.002*** (0.000)
Constant	0.018 (0.023)	-0.084*** (0.024)	1.430*** (0.052)
N	14372	14372	14372
AIC	7007.764	3747.179	26082.792
BIC	7060.775	3800.190	26135.803
LL	-3496.882	-1866.589	-13034.396

Standard errors in parentheses

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

9.5 Determinants of Attention

Table A5 shows the estimations for the determinants of attention. We estimated separate linear regressions with fixed effects for the dwell times for gains (column 1), losses (column 2) and total reaction time (column 3). The error terms are clustered at the level of participant. The variable L. Left corresponds to a dummy variable that takes the value of 1 if in that trial the losses were presented on the left. Similarly, L. First and L. Last take the value of 1 if the first and last fixation was on the losses respectively.

Table A6: Estimations with trial variation

	(1)	(2)
	Decision	Decision
$x_{G,t}$	0.190*** (0.022)	0.160*** (0.022)
$\bar{a}_{j,G} \times x_{G,t}$	0.353*** (0.048)	0.359*** (0.048)
$\tilde{a}_{j,G,t} \times x_{G,t}$	0.008*** (0.003)	0.006** (0.003)
$\tilde{\tau} \times x_{G,t}$		0.042** (0.021)
$ x_{L,t} $	-0.301*** (0.028)	-0.286*** (0.028)
$\bar{a}_{j,L} \times x_{L,t} $	-0.315*** (0.065)	-0.314*** (0.065)
$\tilde{a}_{j,L,t} \times x_{L,t} $	-0.031*** (0.005)	-0.030*** (0.005)
$\tilde{\tau} \times x_{L,t} $		-0.026 (0.026)
N	13057	13057
AIC	7293.469	7296.216
BIC	7338.332	7356.032
LL	-3640.735	-3640.108

Standard errors in parentheses

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

9.6 Robustness Check - Changes in decision weights over time (trials)

The estimations in Table A6 assess whether decision weights change or remain consistent as the experiment progresses. Column 1 replicates our main design, while column 2 adds trials normalized to $\tilde{\tau} = (\tau - 80)/160$ as an additional factor that can affect the decision weights. The estimations include individual fixed effects and the error terms were estimated with jackknife resampling (in parentheses).

10 Supplementary Figures and Materials

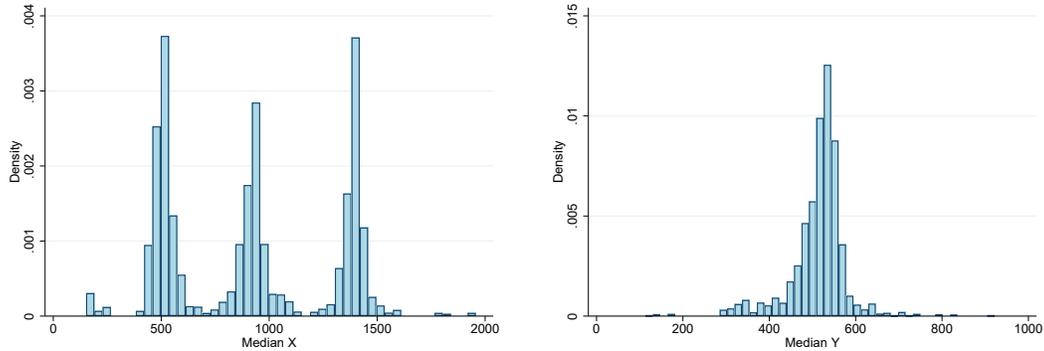


Figure B1: Horizontal and Vertical Clusters of visual fixations

The figures above describe the center of the individual clusters on the x- and y-axis of the screen. Left Panel: three main clusters for the horizontal axis, consistent with the regions of interest (left, middle and right). Right panel: on the vertical axis, there is only one concentration point since all regions of interest are aligned at the same height.

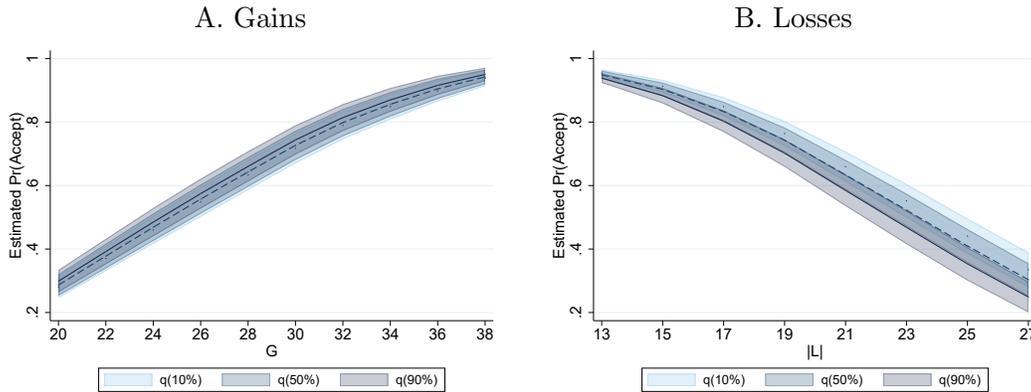


Figure B2: Acceptance probability conditional on residual attention

The figures above show the estimated probability of accepting the lottery (vertical axis) conditional on the outcome values (Left panel for gains, right panel for losses) and the residual attention to the same outcomes. Lines reflect different levels of residual attention, which are the sample quantiles 10% (dotted), 50% (dashed) and 90% (solid). Results reflect the significant interaction (although not visually different) between the gamble's outcomes and average attention present in model 2 (Table 2). The predictions are presented with their 95% confidence intervals.

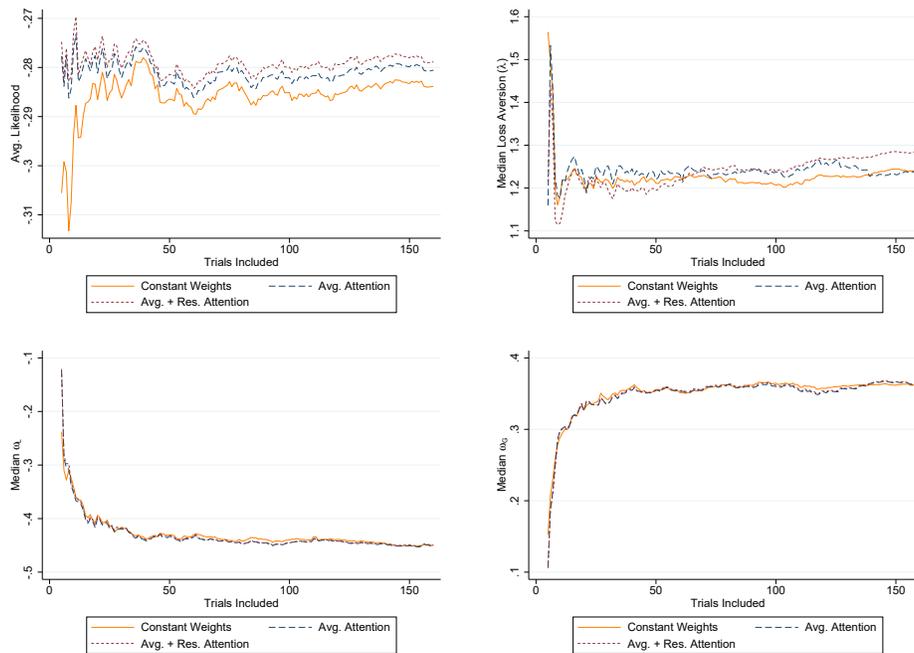


Figure B3: Stability of estimations

The figures above show the results of our estimations by the number of trials included (horizontal axis). The upper-left figure shows the average log-likelihood of the included observations. The upper-right panel shows the median degree of loss-aversion estimated as the ratio of the decision weights of losses over gains. The lower panels show the median decision weights for the gains and the losses respectively.

Information Brochure for Decision-Making Study

Dear participant ,

Thank you for participating in this experiment. Before you start the experiment, it is important that you are aware of the procedures followed in this study. Please read the following text carefully and do not hesitate to ask your experimenter if you have any questions.

Aim of the study

The goal of our experiment is to investigate how people make financial decisions under risk. The experiment will take about 1 hour to complete and we will track your eye movements throughout the experiment.

Experiment procedure

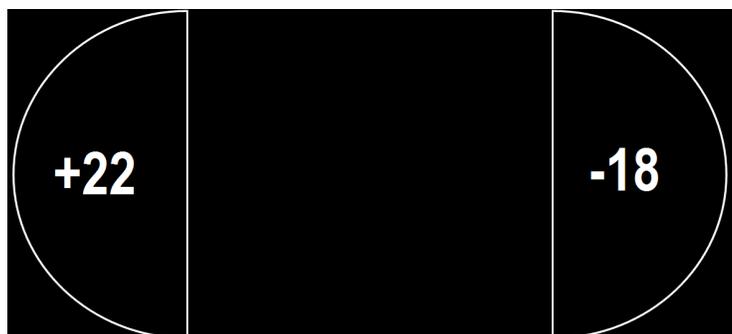
You will receive an initial payment of 10 Euros for filling out a number of questionnaires. Based on your decisions throughout the experiment, you have the chance to earn additional money, as well as to lose money from your endowment. This is because one trial will be randomly selected at the end of the experiment - the payout relevant trial. The decision you made on this trial will be realized as explained in detail below. All values shown in the experiment are in *monetary units* (MU), which have an exchange rate of 1 MU = 0.1851852 Euros.

You are not allowed to write anything down or make notes during the experiment. Moreover, it is very important that you look at the screen throughout the experiment, unless there is a break and we ask you to relax your eyes.

a. Detailed description of the choice scenarios

The experiment consists of a total of 160 decisions. Your task is to make a decision about which of two options you prefer: (A) receiving a certain payout, which leads to no change to your endowment of 10 Euros, or (B) playing a lottery, which can lead to additional earnings with a 50% probability, but also losses with a 50% probability. Choosing the lottery means that you could win, or lose, one of the amounts displayed on the screen with equal (50%) probability. Note that the values offered by the lottery will change on every trial, so please make sure that you pay attention to the amounts on every trial before you make a decision. The certain payout, on the other hand, will remain the same throughout the experiment, such that when you choose this option there will be no additional earnings added to or losses subtracted from your initial endowment.

To make this even clearer, consider the following example: On every trial, the values of the lottery will be displayed on the screen as shown in the figure on the right. In this example trial, gains are shown on the left side (gains are signified by "+") and losses are on the right (losses are signified by "-"). This means that if you decide to accept the lottery on this trial, you will have a 50% chance of



endowment (your payment of 10€ for the questionnaires). Whether you receive the gain or loss will be decided upon via a virtual coin flip, if you selected the lottery on the payout relevant trial selected at the end of the experiment. Note that the locations of gains and losses are not set and can also be reversed on some trials, with losses on the left and gains on the right. Choosing the safe option always leads to no change from your initial endowment, that means you do not receive any additional gains, nor will you incur any additional losses.

Once you have decided which option you prefer, you can communicate your choice by pressing one of two buttons:

- Press the **up arrow key** to choose the lottery.
- Press the **down arrow key** to choose the safe option.

You will receive a brief feedback after you made the decision (for ca. 1 second), which indicates what option you have chosen, such that the letter **L** appears in the center of the screen, when you chose the lottery, and the letter **C** appears in the center of the screen when you chose the certain payout. After a short break, the next lottery will be displayed.

b. Details on payout determination

After you have made your choice for all 160 lotteries, you will select the payout relevant trial by rolling three 10-sided dice. The die rolls will reflect a number between 1 and 160 (the number of all decisions that you have made) as follows: the first die that you roll indicates whether your payout relevant trial is smaller than 100 (die shows a number <5), or greater than 100 (die shows a number ≥ 5). The second and third die rolls then determine the exact trial number. If the trial number is greater than 160, you will roll dice 2 and 3 until a number ≤ 160 is generated. You will then enter the chosen trial number into the computer, which will recall the exact decision that you have made on that trial.

If you chose to play the lottery, a computer algorithm equivalent to an even coin flip will determine whether the gain amount on this trial will be added to your endowment, or whether the loss amount will be deducted from your endowment. Please remember that the monetary units will first be converted to euros using the exchange rate of 1 MU = 0.1851852 Euros. If you chose the certain option, you will receive your endowment of 10 Euros. The amounts on the randomly selected payout relevant trial, your decision and your additional wins or losses will be displayed to you on the screen. Your final payment will be calculated as follows:

If the outcome was a gain: 10 Euro (endowment) + gain amount * 0.1851852

If the outcome was a loss: 10 Euro (endowment) - loss amount * 0.1851852

If you chose the certain outcome: 10 Euro (endowment).

c. Subparts of the experiment

1. At the beginning of the experiment you will fill out questionnaires for about 30 minutes. For your work, you will receive a payment of 10 Euros for use in the following part of the experiment.
2. After the questionnaires, you will be given the chance to familiarize yourself with the experiment in 10 practice trials. These 10 decisions will not affect your final payout and will be made solely for the purpose of giving you experience with the choice scenarios and the speed of the experiment.
3. The main experiment begins after all your questions have been answered and we are certain that you have understood all aspects of the experiment. We will now set up the eye tracker, which monitors where you are looking throughout the remainder of the experiment. To this end, we will ask you to place your head on a chin rest. From this point on, it is very important that you move your head as little as possible and fixate on the screen.
4. At the end of the experiment, we ask you to fill out a final questionnaire, which will take an additional 10 minutes.
5. Finally, you will receive your payment, which will be determined as outlined in detail above.

Confidentiality

All research data will remain completely confidential. In case of either using these results in scientific publications or making these results public in any other way, this will happen anonymously. Personal data will not be seen by others without explicit approval.

VOLUNTARY

Your participation in this study is voluntary. You are free to choose whether to participate in this study. You may also choose to withdraw from the study or to decline to answer any questions at any time. You will not be penalized or lose any benefits to which you are otherwise entitled if you choose not to participate or choose to withdraw.

INSURANCE

Participation in this study involves making simple choices which is routinely used and will do no harm to your health or safety. Because this study poses no risks to your health or safety, the conditions of the regular liability insurance of the University of Amsterdam are applied.

FURTHER INFORMATION

If you have questions about this research beforehand or afterwards, please contact the responsible researcher dr. Jan Engelmänn (e-mail j.b.engelmann@gmail.com). In case of complaints about this study, you can contact Dr. Wery van den Wildenberg, member of the ethical committee of the Psychology Department of the University of Amsterdam (Fmg-UvA, REC-G1.10, Nieuwe Achtergracht 129 B, 1018 WS Amsterdam, 020-5256686, w.p.m.vandenwildenberg@uva.nl).

AGREEMENT

When you sign this document containing a written explanation of the experiment that you are participating in, you declare that you have read and understood the instructions and that all your questions have been answered by the experimenter. Moreover, with your signature you agree to participate in the procedures outlined in the instruction above.

If you have further questions about this experiment, please contact the responsible researcher dr. Jan Engelmann (e-mail j.b.engelmann@gmail.com). In case of complaints about this study, you can contact Dr. Wery van den Wildenberg, member of the ethical committee of the Psychology Department of the University of Amsterdam (Fmg-UvA, REC-G1.10, Nieuwe Achtergracht 129 B, 1018 WS Amsterdam, 020-5256686, w.p.m.vandenwildenberg@uva.nl).

[Participant]

“I have read and understood the information above and agree to participate in the current experiment and grant the experimenters permission to use my data. I reserve the right to withdraw from this agreement without giving any explanation, as well as to withdraw from participation in this experiment at any time.”

Date:

.....
Participant name

.....
Signature

[Experimenter]

“I have explained the experiment to the participant. I will answer any further questions to my best knowledge.”

Date:

.....
Researcher name

.....
Signature

Exit Questionnaire.

Thank you again for participating in our experiment. Because we are always concerned with improving the experiment and the instructions, we have just a few questions for you. Please rate how much you agree with the following statements using the scale below.

0	1	2	3	4
Strongly disagree	Disagree	Undecided	Agree	Strongly agree

Statement	Your Evaluation
During the experiment, I <i>never considered</i> that I would not receive the amount that I selected via dice rolls at the end the end of the experiment.	
During the experiment, I <i>considered</i> that the experiment was programed in such a way that would make me lose money.	
During the experiment, I <i>never thought</i> that I was being deceived by the experimenters about the additional money I could win or lose?	
During the experiment, I <i>fully understood</i> that the values shown would be converted to Euros via an exchange rate.	

- Have you ever participated in an experiment in which you were deceived? Please circle your answer.

Yes

No

Cannot tell /
do not remember

- To what extent do you think that previous experiences with deception influenced your behavior in the current experiment? Please circle one answer.
Previous experiences with deception influenced me in this experiment ...

0	1	2	3	4
Not at all	Slightly	Somewhat	Moderately	Extremely

- Please use the space below if you have any other comments or questions about the experiment.
