

The Attraction Effect in Information Visualization

Evanthia Dimara, Anastasia Bezerianos, and Pierre Dragicevic

Abstract—The attraction effect is a well-studied cognitive bias in decision making research, where one's choice between two alternatives is influenced by the presence of an irrelevant (dominated) third alternative. We examine whether this cognitive bias, so far only tested with three alternatives and simple presentation formats such as numerical tables, text and pictures, also appears in visualizations. Since visualizations can be used to support decision making — e.g., when choosing a house to buy or an employee to hire — a systematic bias could have important implications. In a first crowdsourcing experiment, we indeed partially replicated the attraction effect with three alternatives presented as a numerical table, and observed similar effects when they were presented as a scatterplot. In a second experiment, we investigated if the effect extends to larger sets of alternatives, where the number of alternatives is too large for numerical tables to be practical. Our findings indicate that the bias persists for larger sets of alternatives presented as scatterplots. We discuss implications for future research on how to further study and possibly alleviate the attraction effect.

Index Terms—Information visualization, decision-making, decoy effect, attraction effect, asymmetric dominance effect, cognitive bias.

1 INTRODUCTION

Suppose you are voting for primary elections and need to choose between candidates Bob and Alice (Table 1). Bob has a solid education plan, but not much concern for crime control. In contrast, Alice's education plan is weak but she has an excellent strategy for crime control. If both education and safety are important to you, this can be a difficult choice. Now suppose there is a third candidate, Eve. Like Alice, Eve focuses more on crime control than education, but her crime control strategy is not as good as Alice's. O'Curry and Pitts [38] used a similar decision task in a study, and showed that adding Eve as an option shifted participants' preference towards Alice.

Table 1: Three hypothetical candidates in political elections

	Bob	Alice	(Eve)
education	★★★★	★★	★★
crime control	★★	★★★★★	★★★★

This shift in preference called the *attraction effect* (also known as the *decoy effect* and the *asymmetric dominance effect*), is a cognitive bias whereby people tend to favor the option for which there exists a similar, but slightly inferior, alternative. Like other cognitive biases, the attraction effect leads to irrational decisions and has important implications in many areas such as politics and advertising. Our goal in this article is to find out whether the attraction effect also has implications for information visualization design. In our example, voters' decision is influenced by the presence of Eve, which is inferior in all respects and therefore irrelevant to the choice. If, in the same way, if someone uses a visualization to choose among several options (e.g., when buying an apartment [55]), will the presence of inferior choices affect their decision? In other words, does the attraction effect transfer to visualizations?

The current information visualization literature does not offer much empirical data to help us answer this question. Although biases and misjudgments have been studied, the focus has been on perceptual biases such as in color perception or magnitude estimation [53]. Cognitive biases differ from perceptual biases in that they

persist even if the information has been correctly processed at a perceptual level. There is a growing interest in cognitive biases in information visualization, but studies have so far focused on probabilistic reasoning and judgment under uncertainty [34, 42].

While there has been little work studying the role of cognitive biases in information visualization, visualization systems are increasingly used to support decision making. Large companies switch to visualization solutions to improve their human strategic decisions for profitable drug trials [46], or use visualizations to choose which features of a software they should release and when [3]. In addition, many visualization tools previously introduced in research explicitly or implicitly claim to help people make decisions such as choosing a house to buy [55], finding a nursing home [57], selecting healthy cereals [58], choosing a digital camera [16, 32], finding a profitable investment [42, 13], or selecting a site for a new factory branch [2].

A visualization is generally considered effective if it helps people extract accurate information [9, 59]. Nevertheless, we know from decision making research that full access to information does not necessarily yield good decisions [28]. Generally, the more complex a decision, the more we resort to heuristics, i.e., “*simple procedures that help find adequate, though often imperfect, answers to difficult questions*” [28]. While heuristics can be very effective [18], they can also lead to cognitive biases [28]. Therefore, in order to fully understand how information visualizations can support decision making, we need to study how they interact with cognitive biases.

We focus on the attraction effect for two reasons. First, it is one of the most studied cognitive biases in fields such as psychology, consumer research and behavioral economics. Second, these studies generally employ very small sets of alternatives (typically three) and numerical presentation formats, so it is still unknown whether the bias generalizes to data visualizations. Although some visual representations have been considered, there is conflicting evidence and a heated debate on whether the effect generalizes [17, 26, 48, 56]. Some argue that the effect occurs only in numerical stimuli [17], e.g., when attributes are presented in tables. Whereas others argue that it is generic and robust, and can be observed in many contexts such as visual judgments in shapes [52], oral instructions [44], or even among animals when they choose their food [31]. This debate suggests that the attraction effect is far from being fully understood and needs to be investigated from a variety of perspectives.

We study the attraction effect from an information visualization perspective in two crowdsourcing experiments. In the first, we test and verify that the attraction effect indeed persists when alternatives are presented in a scatterplot rather than in a numerical table. We then generalize the attraction effect procedure to more than three alternatives, and verify that the effect can persist when participants are presented with more realistic scatterplot visualizations involving about 20 data points. We finally discuss our findings and conclude with implications for future research.

- E. Dimara is with Inria and Université Paris-Saclay. E-mail: evanthia.dimara@gmail.com
- A. Bezerianos is with Univ Paris-Sud & CNRS (LRI), Inria, and Université Paris-Saclay. E-mail: Anastasia.Bezerianos@lri.fr
- P. Dragicevic is with Inria and Université Paris-Saclay. E-mail: Pierre.Dragicevic@inria.fr

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx/

2 BACKGROUND

We review work on decision-making and cognitive biases in information visualization, and on the attraction effect in other fields.

2.1 Information Visualization

Compared to numerical and textual formats, it is known that data visualizations can highlight relationships in the data, facilitate the recognition of patterns, and reduce cognitive load [9, 43, 50]. As they aid data exploration and understanding, it is generally assumed that data visualizations can support better decision making [2]. Based on this intuition, several decision-support systems that rely on interactive data visualization have been developed [59].

2.1.1 Use of Visualizations in Decision Support

A range of interactive data visualization tools have been proposed to help people make decisions. Sometimes decision making is simply used to demonstrate a new visualization or interaction technique. For example, HomeFinder [55] helps people find a house to buy using scatterplot visualizations and dynamic queries. The Dust & Magnet tool [58] is illustrated with a scenario for choosing cereals by moving magnets that attract or repel cereals according to attributes such as calories. Other tools target explicitly decision support. For example, Asahi et al. [2] visualize hierarchical decision criteria using treemaps, augmented with interactions to make decisions such as whether or not to construct a dam, file a patent, or choose a factory's location. ValueCharts [5] let consumers choose a TV set or a hotel by providing a set of domain-independent visualizations.

Domain-specific decision-support visualization systems have also been proposed. For example Decision Map / Table & Box [57] helps people find an appropriate nursing home by combining several coordinated views. Stratos [3] helps software project managers select which features to include in each production stage, by simultaneously visualizing all possible software release plans. VisIDM [13] helps people choose a financial investment through uncertainty visualizations and support for personalized risk preferences.

Although most of these tools come from research, similar ones are used in industry. For example, after losing millions of dollars in late drug trial failures, a large pharmaceutical company decided to use interactive visualizations to better track and facilitate decisions of "cut or go" projects in their early stages [46].

2.1.2 Limitations of a Pure Informational Approach

Interactive data visualizations facilitate data exploration and sense-making, making data accessible and promoting informed decisions. Furthermore, the use of interactive systems, rather than automatic analysis, leaves room for human judgment, which is crucial where expert knowledge or subjective preferences cannot be fully formalized (e.g., importance of education vs. crime control). However, most visualization tools for decision support appear designed under the assumption that decisions are made by rational people who only need to be given complete information to be able to make good decisions. And thus, that good decisions should be the natural outcome of reliable data conveyed with well-designed visualizations.

It is by now widely recognized that even perfectly informed people are not perfect decision makers [28]. The imperfections of heuristics we routinely use manifest themselves as cognitive biases, like the attraction effect. Cognitive biases are far from trivial to overcome: they occur even when all relative information is available and well perceived, and they persist even when we inform or train people on how to overcome them [19]. Thus we need to investigate further if visualization designs are likely to suffer from cognitive biases, and whether we could improve our designs to alleviate these biases.

2.1.3 Cognitive Biases and Visualizations

Information visualization has studied perceptual biases [53, 60], but cognitive biases have comparatively received little attention.

Zuk and Carpendale [61] discuss cases where visualizations may aid to remediate uncertainty biases. Researchers have studied how visualizations, such as Euler diagrams and frequency grids, can reduce the base rate bias in probabilistic reasoning [34, 29]. FinVis

[42] is a tool that shows investment options using tables and visualizations to help investors overcome the uncertainty aversion and diversification bias. Miller et al. [35] used scatterplots and histograms to help fantasy baseball experts overcome regression bias in their predictions. Although many of these previous studies try to examine how visualizations can help overcome cognitive biases, some studies found that visualization-based remediation can be challenging [34, 30], or that cognitive biases can co-occur with visualizations. For example, Zhang et al. [59] showed that startup companies presented with tabular visualizations were subject to conservatism and loss aversion biases in their probability judgments. Some biases, like the within-the-bar bias, only appear with visualizations [11].

Most of these previous studies focused on judgment under uncertainty. Although reasoning based on uncertain information is hard and pervades our everyday lives, uncertainty is not the only cause of irrationality in decision making. In the attraction effect, irrationality stems instead from the fact that decisions are influenced by irrelevant information (the presence of a decoy). Our work is thus significantly different from previous work about reasoning under uncertain information.

2.2 The Attraction Effect

We next define the attraction effect and the terminology used in this article. We present theories on why the effect exists, and discuss recent studies investigating the effect on visual stimuli.

2.2.1 Terminology

A **decision task** involves choosing one among several **alternatives** (i.e., Alice, Bob or Eve in our example). Alternatives are characterized by **attributes** (e.g. their support for education and crime control), which take values that are unambiguously ordered in terms of preference (e.g. more crime control or education is better than less).

An alternative **dominates** another if it is strictly superior in one attribute and superior or equal in all others. An alternative is **dominated** within a set of alternatives if there is at least one alternative that dominates it. In our example Eve is dominated by Alice, because she is equal in education and worst in crime control. In this decision task Eve would be formally a "wrong" answer.

An alternative is **asymmetrically dominated** within a set of alternatives if it is dominated by at least one alternative, but is not dominated by at least one other [25]. Eve is asymmetrically dominated because she is dominated by Alice but not Bob, since Eve offers better crime control than Bob. We call two alternatives **formally uncomparable** if neither dominates the other, as is the case for Alice and Bob. The best candidate is a matter of personal choice.

A typical attraction effect experiment involves a decision task with three alternatives, two that are formally uncomparable, and one that is asymmetrically dominated. They are referred to as: the **decoy**, the asymmetrically dominated alternative (Eve); the **target**, the alternative that dominates the decoy (Alice); the **competitor**, the alternative that does not dominate the decoy (Bob). This decision task is typically compared with a task where the decoy is absent, i.e. that involves only the two formally uncomparable alternatives.

The **attraction effect** is a cognitive bias where the addition of a decoy (Eve) in a set of two formally uncomparable alternatives increases people's preference for the target (Alice) [25, 27]. In experimental settings this preference switch is observed not for any single individual but between groups, where a higher percentage of people generally choose the target when the decoy is present. This switch in preference is irrational because it violates a basic axiom of rational choice theory, the principle of *regularity*, according to which the preference for an alternative cannot be increased by adding a new alternative to the choice set [25]. Attraction effect experiments assume that decision makers behave rationally in all other respects, and that they are able to perceive dominance relations. As a consequence, they are expected to never choose the decoy.

Later on, we will generalize the attraction effect to more than three alternatives. For now, we discuss previous work on the attraction effect, which always involves two alternatives plus a decoy.

2.2.2 Why does the Attraction Effect Occur?

Two types of explanatory theories have been offered for the attraction effect: strategic ones and perceptual ones [33].

Strategic Explanations: According to strategic theories, people use the dominance over the decoy as a heuristic to simplify an otherwise difficult decision. Choosing the target is also easier to justify to others [47] — in our example, someone who chooses Alice could argue that she is at least better than Eve. Neuroimaging studies have additionally shown that the presence of a decoy tends to reduce negative emotions associated with the decision task [21].

Perceptual Explanations: So-called “perceptual” theories assume that the addition of a decoy changes how people perceive the relative importance of the attributes involved, giving more weight to the attribute on which the target is strong [1, 22]. By analogy with perceptual contrast effects (e.g., an object appears larger when surrounded by small objects), the target appears more attractive when surrounded by unattractive alternatives [49]. In our example, if Eve is present, crime control may appear more important as two candidates perform relatively well on this criterion. Since this is the strength of Alice, it may raise her perceived value compared to Bob.

All explanations agree that for the attraction effect to occur, a perceptible dominance relation between the target and the decoy is key.

2.2.3 Can the Attraction Effect Occur with Visualizations?

Studies suggest the attraction effect is quite general and robust, e.g., it occurs when people choose consumer products like beers, cars, or films [25], when they gamble [54], select candidates to hire [24], decide which suspect committed a crime [51], or vote [38]. Even animals like hummingbirds [4], bees [45], and amoebae [31] appear to be subject to the same bias when selecting their food.

The attraction effect has been observed under a variety of experimental conditions, the majority of which present decision tasks as numerical tables. A few studies have shown that the effect generalizes to non-tabular representations, such as pictures of consumer products [49], verbal instructions [44], and physical objects (i.e., people choosing between cash and a pen, or between tissues and towels) [49]. Studies have further suggested that the effect occurs when carrying out visual judgment tasks, such as finding the largest rectangle [52] or finding similarities in circle and line pairs [10].

Nevertheless, several authors [17, 56] have recently argued that the attraction effect only occurs when attributes are presented in numerical format, and reported failures to replicate the previous studies involving the representations mentioned above. Others subsequently questioned the validity of these replications [26, 48]. This debate on whether the effect generalizes to non-numerical presentations opposes (i) *numeric* displays of *quantitative* information with (ii) displays of *qualitative* information such as photos, verbal descriptions, or physical objects. As most data visualizations are *pictorial* displays of *quantitative* information, the debate does not provide evidence on whether the effect occurs in visualizations.

Frederick et al. [17] however studied a gambling task with two or three bets presented either as a table, or as a diagram. Each bet had a prize in dollars and a probability to win. In the diagram condition, the probability of each ticket was shown as a “probability wheel” (analogous to a pie chart), and the prize was shown underneath, as a number. When gambles were presented as numeric tables, the decoy nearly doubled the share of the target, but when pie charts were used, the effect disappeared. To the best of our knowledge, this is the study that comes closest to a test of the attraction effect on visualizations. Nevertheless the diagram design was very domain-specific, and only one of the two attributes (probability, but not price) was encoded visually. We address this by using 2D scatterplots.

Although why the attraction effect occurs is still not fully understood, the possibility that it persists in visualizations is consistent with both the strategic and the perceptual explanatory theories. Both assume that the effect requires the ability to make attribute-to-attribute comparisons and to recognize the dominance relation between target and decoy. If anything, visualizations could make these tasks easier and could perhaps even amplify the effect.

3 GYM EXPERIMENT: TABLE/SCATTERPLOT, 3 CHOICES

The purpose of this first experiment is to replicate the design of a standard attraction effect experiment (two alternatives plus a decoy presented in a numerical table), and then to test if the effect persists when alternatives are shown using a scatterplot visualization.

Similar to Frederick et al. [17] who successfully replicated the attraction effect with tables but not with non-numerical formats, our study was conducted using crowdsourcing. Crowdsourced experiments are now commonly used in information visualization [23], including in studies involving judgment and decision making [34, 29]. We used Crowdfunder¹ as the crowdsourcing platform.

3.1 Design Rationale

Although the attraction effect is thought to be robust, a replication can fail if not enough attention is paid to the details of the experimental design [26, 48]. We therefore based our design choices on lessons and recommendations from the attraction effect literature.

3.1.1 Scenario and Attribute Values

By *scenario* we refer to the semantic and narrative context of the decision task. In our introduction example, alternatives are candidates, attributes are support for education and crime control, and the decision consists of voting for a candidate.

Many different scenarios and attribute values have been employed since the original studies of the attraction effect [25, 27]. We reasoned that a recent study is more likely to employ an optimal design, since it has more accumulated knowledge to build on. We therefore chose to replicate the scenario from the first experiment of recent work by Malkoc et al. [33], that involved choosing a fitness club (or gym), and found a clear attraction effect.

In Malkoc et al.’s study, each gym was defined by its variety and its cleanliness, both rated from -10 to +10. A positive rating meant better than average, and a negative rating meant worse. The study investigated whether undesirable options (all negative ratings) eliminate the attraction effect. But as the effect was strong for their control condition (all positive ratings), we chose it for our replication.

The experiment employed four gyms g_C, g_V, g_C^*, g_V^* , where g_C was cleaner, g_V had more variety, and g_C^* and g_V^* were slightly less attractive than g_C and g_V respectively. The attribute values were $g_C(\text{variety} = 1, \text{cleanliness} = 4)$, $g_V(4, 1)$, $g_C^*(0, 4)$, and $g_V^*(4, 0)$. Three decision tasks were tested: $\{g_C, g_V\}$ (no decoy), $\{g_C, g_V, g_C^*\}$ (decoy on g_C), and $\{g_C, g_V, g_V^*\}$ (decoy on g_V). These attribute values however cause the data points g_C^* and g_V^* to overlap with scatterplot axes, possibly creating visual anchoring effects that could affect participant responses. Since such effects were outside the scope of our study, we incremented all values by one. Thus we used as attribute values $g_C(2, 5)$, $g_V(5, 2)$, $g_C^*(1, 5)$, and $g_V^*(5, 1)$. These values preserve all dominance and similarity relationships between alternatives.

3.1.2 Stimuli: Tables and Scatterplots

We used a numerical table as a control condition, to test our experiment design and compare our results with previous studies. Figure 1a shows the 2×2 table representation for the decision task $\{g_C, g_V\}$, and Figure 1b shows the 3×2 table for the decision task $\{g_C, g_V, g_V^*\}$. Attributes were presented in rows and alternatives in columns, as in Malkoc et al. [33]. Alternatives were labeled A, B or A, B, C from left to right. The ordering of rows and columns in the table will be discussed in the next subsection.

In the visualization condition, alternatives were conveyed with scatterplots (see Figure 1c,d) and similarly labeled A, B or A, B, C from left to right and from top to bottom.

There are four main reasons behind the choice of scatterplots for the visualization condition. First, 2D scatterplots are a standard information visualization technique [16, 37]. Second, they are suited for visualizing any tabular dataset with two quantitative dimensions, which captures the decision tasks used here and most decision tasks used in previous studies on the attraction effect. Third, a

¹<http://www.crowdfunder.com/>

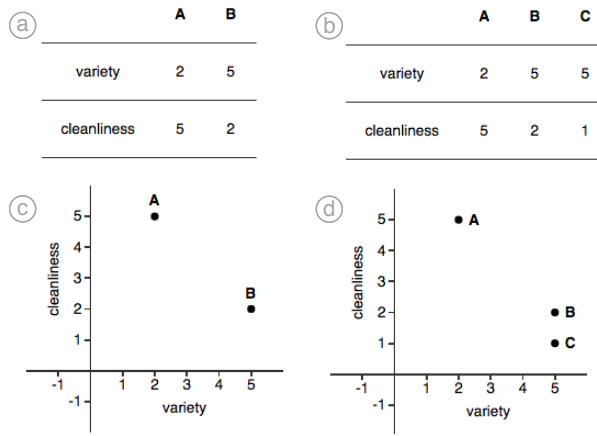


Fig. 1: Examples of experimental stimuli for the table (a,b) and the scatterplot (c,d) conditions. The left decision task (a,c) has no decoy, while the right decision task (b,d) has a decoy on B.

scatterplot shows all data cases within the same frame of reference, thus providing a rapid overview of all alternatives. A unified frame of reference also likely supports comparisons better than side-by-side views such as Frederick et al.'s [17] pie charts discussed in the background section. In fact, scatterplots are used as figures in most articles on the attraction effect for conveying the alternatives used in the experiments [4, 20, 24, 25, 27, 31, 33, 36, 38, 40, 45, 47]. Finally, scatterplots scale up to more than three items, which is an important requirement for our follow-up experiment.

The appearance of tables and scatterplots was kept as similar as possible to avoid experimental confounds due to choices in visual design. Both presentation formats took approximately the same amount of screen real estate, and graphical attributes (colors, line thickness and font sizes) were kept consistent. In both conditions, participants indicated their choice through separate radio buttons.

3.1.3 Ordering of Alternatives and Attributes

Although Malkoc et al. [33] used a fixed order of presentation for attributes and alternatives, the choice of ordering may affect participant responses, in particular in our experiment where different presentation formats are used. For example, participants may give more weight to variety if it is shown first on a table, but on a scatterplot, it is not clear whether the choice of horizontal vs. vertical axis would have a similar effect. In addition, alternatives can be presented in any order within a table, while on a scatterplot the way alternatives are laid out is dictated by attribute values.

To balance out any possible order effect, we thus randomized the order of presentation of attributes and alternatives across participants. In the scatterplot condition, axes can be flipped, leading to 2 different attribute orderings (variety on x and cleanliness on y , or vice versa). In a 2×2 table, there are 2 ways to order rows and 2 ways to order columns, yielding 4 different tables. Similarly, a 2×3 table can be presented in 12 different ways. Since the decoy is typically placed next to the target in attraction effect experiments (e.g., [17, 20, 21, 22]), we removed cases where the target was not next to the decoy (4 tables out of 12). Since the decoy cannot appear between the target and the competitor in the scatterplot, we also removed cases where the decoy was in the middle (4 tables out of 12). In summary, we used 4 different table stimuli and 2 different scatterplot stimuli for each of the three decision tasks $\{g_C, g_V\}$, $\{g_C, g_V, g_C^*\}$ and $\{g_C, g_V, g_V^*\}$, for a total of 18 different experimental stimuli.

3.1.4 Crowdsourcing Quality Control

Quality control is important in any crowdsourcing experiment [23], and in attraction effect studies in particular [48]. Quality was ensured by recruiting highly-rated crowdsourcing contributors (level 3 on the Crowdfunder platform), by including test questions, and by

devising a job assessment scheme prior to running the experiment. Four criteria were used for job assessment:

Completion time. A job completion time of less than 1 minute or more than 30 minutes was considered abnormal. Our pilots indicated an average task completion time of 6 minutes.

Justification. Participants had to provide a textual justification for their choice. Justifications were classified by one investigator as either proper or improper, depending on whether it made a reference – direct or indirect – to either cleanliness or variety. Participants were informed in advance that they would have to justify their choice, as this has been linked to a stronger attraction effect [47].

Prior preferences. After the experimental task, participants were asked if they suffered from an abnormal fear of dirt (or bacteriophobia), with “no”, “yes”, or “unsure” as answers. This identified participants with a strong prior preference for cleanliness, as strong prior preferences are known to reduce the effect [36, 41].

Table and scatterplot tests. After carrying out the task, participants were subjected to two screening tests: a *numerical table test*, and a *scatterplot test*, irrespective of the condition they saw. Both tests involved choosing between three laptops based on their RAM and CPU, with one laptop clearly dominating the other two (i.e., had both higher RAM and higher CPU). The tests were designed to be trivial, with a single correct answer, using a presentation format similar to the experimental task (see Figure 1). The purpose of the table test was to screen for contributors who did not pay attention to the tasks. The purpose of the scatterplot test was to control for visualization literacy [8], and make sure that participants were able to read scatterplots and to perceive dominance relations [26].

We classified jobs in three categories: the *Red*, where the job is rejected (and the contributor not paid); the *Orange*, where the job is accepted but the data discarded from our analysis; and the *Green*, where the job is accepted and the data kept in our analysis. Due to limitations in the Crowdfunder platform we had to pay all contributors, but we report here on the three categories nonetheless.

A total of 437 jobs were submitted, after removing invalid completion codes and duplicate worker IDs. A job was marked *Red* if: the completion time was abnormal (1 % of all submitted jobs), the gym choice was not properly justified (14%), or the contributor failed the table test (12%). A job was marked *Orange* if: the response to the bacteriophobia question was “yes” (12% of all submitted jobs), or the contributor failed the scatterplot test (13%). In total, 16% of all submitted jobs were marked *Red* and 14% were marked *Orange*. These jobs were discarded from all our analyses.

3.2 Experiment Design

The experiment followed a 3×2 between-subjects design. The first independent variable was the decision task, which involved three different datasets: $\{g_C, g_V\}$, referred to as the *no decoy* condition; $\{g_C, g_V, g_C^*\}$, referred to as *decoy on cleanliness*; and $\{g_C, g_V, g_V^*\}$, referred to as *decoy on variety*. The second independent variable was the presentation format, with two conditions: *table* and *scatterplot*.

3.2.1 Procedure

We conducted a first pilot study to ensure the clarity of the instructions, and we then uploaded the experiment as a Crowdfunder job.

Participants had to open an external 8-page Web form. They were told they would have to choose a fitness club based on two attributes: variety of the machines and cleanliness of the club. They had to assume that they had done some preliminary research, and had narrowed down their choices to two (in the no-decoy condition) or three (in the decoy conditions) clubs. They were then shown the gyms as a table or a scatterplot (Figure 1) and asked to choose one.

Once finished, participants rated their confidence on a 7-point scale and provided an open text justification for their choice. They also rated their enthusiasm towards fitness clubs on a 7-point scale and reported on whether they suffered from bacteriophobia. Finally, they were given the table and scatterplot tests (Section 3.1.4), and filled a short questionnaire with demographic information.

At the end of the experiment, participants copied the provided completion code and pasted it in the crowdflower platform to receive payment. The entire job took on average 6 minutes to complete, and participants were paid \$0.60 upon completion.

3.2.2 Participants

Our population sample consisted of 305 crowdsource contributors who submitted valid responses, i.e., jobs classified as *Green* (Section 3.1.4). Job assignments were left on the crowdsourcing server until the planned sample size of $n=50$ per condition was approximately reached. We obtained $n=54$, 51, 50 for the *table* decision tasks, and $n=47$, 53, 50 for the *scatterplot* tasks.

A summary of our participants' self-reported demographics is shown in Figure 2 (map and bar charts labeled "Gyms"). As can be seen, participants tended to be educated young male adults.




Fig. 2: Participant demographics for both experiments.

3.2.3 Hypotheses

Our statistical hypotheses were:

- H1** A larger proportion of participants will choose the target in the *table* \times *decoy on cleanliness* and the *table* \times *decoy on variety* conditions than in the *table* \times *no decoy* condition.
- H2** A larger proportion will choose the target in the *scatterplot* \times *decoy on cleanliness* and the *scatterplot* \times *decoy on variety* conditions than in the *scatterplot* \times *no decoy* condition.

3.3 Results

We analyze, report and interpret all our inferential statistics using interval estimation [15]. The experimental stimuli, data and analysis scripts are available at <http://www.aviz.fr/decoy>.

3.3.1 Planned Analyses

All analyses reported in this section were planned before data was collected. One planned analysis (an analysis of differences between attraction effects) was not conducted because it required equal sample sizes across all conditions.

Only one participant out of 306 chose a decoy, which is low compared to previous studies, where decoy selection rates can be as high as 13% [17]. This shows that participants carried out the tasks seriously and could perceive dominance relationships. The decoy choice is removed from the rest of this analysis.

Participant choices are shown in the top of Figure 3 marked "Gyms" ("Bets" refers to our second experiment). The top three bars are for the *table* format, in the conditions *no decoy*, *decoy on cleanliness* and *decoy on variety*. Adding a decoy is expected to increase the proportion of choices of the target, in the direction indicated by the arrow. This was indeed the case for the *decoy on variety* condition (a 20% increase), but not for *decoy on cleanliness* (a 6% decrease). The next three bars refer to the *scatterplot* format. Here the expected increase was observed for both *decoy on cleanliness* (a 18% increase) and *decoy on variety* (a 3% increase). We now turn to inferential statistics to determine to what extent these effects are reliable.

The previously reported effects are shown in Figure 4 — the four black dots under the category "Gyms". Effects are expressed in percentage points, where a positive value (i.e., to the right of the vertical dashed line) indicates an attraction effect. Dots are sample




Fig. 3: Proportions of participant choices in both experiments.




Fig. 4: Point estimates and 95% confidence intervals for the attraction effects in Malkoc et al. [33], and in our two experiments.

statistics, while error bars are 95% confidence intervals indicating the range of plausible population effects [12]. Confidence intervals were computed using score intervals for difference of proportions and independent samples.

Figure 4 shows that the unexpected reversal observed in the *table* \times *decoy on cleanliness* is too unreliable for any conclusion to be drawn. The same is true for the small effect found for *scatterplot* \times *decoy on variety*. However, we have good evidence for an attraction effect in the other two conditions. The magnitude of the effect is comparable to Malkoc et al. [33], shown on the top of Figure 4.

Thus, our results partially confirm **H1** and **H2**, but are less "clean" than in Malkoc et al.'s [33] original study.

3.3.2 Additional Analyses

Participants reported similar confidence in their answers across all conditions (Figure 5). They were overall highly confident, with a mean rating of 5.9 to 6.1 on a 7-point Likert scale, depending on the condition. Participants' reported familiarity with fitness clubs varied, but they were overall rather familiar (Figure 5).

We computed *combined attraction effects*, shown as purple dots and error bars in Figure 4. A combined attraction effect is the sum of the attraction effects obtained in both decoy conditions, or equivalently, the difference in choice proportions between these two conditions (i.e., the differences between the bars marked "decoy" in Figure 3). This combined measure generally yields more statistical power and facilitates comparisons of results since some experiments (e.g., [54] and our next experiment) do not include a no-decoy condition and thus only report combined attraction effects.

The two purple error bars in Figure 4-Gyms show that the data overall speaks in favour of an attraction effect, both for the *table* and the *scatterplot*. To better quantify the strength of evidence, we conducted a Bayesian analysis using the Jeffreys prior for proportions [7]. Ignoring previous studies and considering our data only, the presence of a combined attraction effect in the *table* condition is 34 times more likely than a practically null effect (set to $\pm 1\%$), and 11 times more likely than a "repulsion" effect. In the *scatterplot*, a combined attraction effect is 150 times more likely than a practically null effect, and 66 times more likely than a repulsion effect.




Fig. 5: Self-reported confidence and familiarity in both experiments.

3.3.3 Discussion

We found evidence for an attraction effect on *table* for the *decoy on variety* condition, but not for the *decoy on cleanliness* condition, where the effect may be smaller or even possibly negative (see Figure 4). We do not have an explanation for this asymmetry, but the wide confidence intervals and their large overlap suggests that the difference may be due to a large extent to statistical noise [12].

Based on the combined attraction effect which is a more holistic measure with more statistical power, we replicated the attraction effect on tables (H1) but the results are less strong than in the initial study [33] (i.e., about half of the original study, as shown by the purple CIs in Figure 4). It is common for a replication to yield smaller effect sizes [39], but the differences in results could also be due to modifications we made to the original experiment design.

We produced four different stimuli for each decision task in order to eliminate possible presentation order effects for alternatives and attributes, whereas Malkoc et al. [33] used a unique table. The use of different stimuli could have yielded a higher variability in responses.

Our study was also a crowdsourcing experiment, whereas Malkoc et al. conducted theirs with students in a lab, where participants are less diverse and generally more focused [34]. Perhaps the feeling of being evaluated was also stronger for students, which we know can amplify the attraction effect [47]. Our rejection criteria (e.g., textual justification for the answer, table and scatterplot test, attention test) could have also filtered subsets of the population that are more vulnerable to the effect. Finally, our participants were on average rather familiar with gyms (Figure 5), and 11% were unsure if they suffered from bacteriophobia, and we know that familiarity with the subject matter and strong prior preferences can reduce the effect [36, 41]. Malkoc et al. [33] do not report on familiarity and prior preferences.

Despite mixed results for the table condition, we obtained good evidence for an attraction effect in the *scatterplot* condition. There still appears to be an asymmetry between the two decoy conditions (this time, in the opposite direction), but CIs show no evidence for a difference. The combined attraction effect provides compelling evidence that the attraction effect can generalize to scatterplots (H2). This observed shift in preference after adding an irrelevant option to a two-point scatterplot gives credence to the idea that people may make irrational decisions even when they use visualizations as decision making aids. Thus we decided to explore the effect further, using scatterplots with larger sets of alternatives.

4 EXTENDING THE ATTRACTION EFFECT

Our gym experiment confirmed that the attraction effect can extend to scatterplot formats. However, we have so far only considered three data points, which does not capture most real-world decision tasks where visualizations would be used.

Previous work has focused on only three alternatives because in numeric tables, it is hard to perform rapid attribute-to-attribute comparisons and recognize dominance relationships between many points. Bettman et al. [6] point out that the attraction effect requires asymmetric dominance relationships to be “perceptual in nature” and “easy to access”. They expect that the bias will be eliminated with multiple alternatives, as the number of pairwise comparisons increases and these relationships become harder to understand. This may be true for numerical tables, but not necessarily for visualizations such as scatterplots, that are designed to aid viewers read and understand complex data, and support comparison of many data points at once [37]. It is thus plausible that visualizations of many alternatives can also elicit attraction effects.

4.1 Ways of Adding More Alternatives

There are three ways the classical attraction effect procedure can be extended to include more than three alternatives:

1. By adding more non-dominated options. In our introduction example, the only non-dominated alternatives were Bob and Alice. We could add more candidates that neither dominate nor are dominated by Bob and Alice. The set of formally uncomparable or non-dominated alternatives is also called the *Pareto front*.
2. By adding more decoys. In our example the only decoy is Eve. We could however add more decoys similar to Eve.
3. By adding “distractors”, i.e., irrelevant options that play neither the role of target, of competitor, or of decoy. An example would be a dominated candidate that appears both in the baseline condition and in the decoy condition.

The first approach is problematic in at least two respects. One is that since it breaks the dichotomy between target and competitor, it would require a major change in the way the attraction effect is measured in experiments. A second problem is that it would cause the attraction effect to interfere with other cognitive biases. For example, the *compromise effect* is a bias by which if presented with several formally uncomparable alternatives, people tend to avoid extremes and choose options in the middle [47]. Even though it could be informative to study how the two effects may combine, we decided here to focus on the attraction effect only.

Adding an arbitrary number of distractors (option 3) is however possible. With many distractors a single decoy is unlikely to produce a measurable effect, but more decoys can be added (option 2). The Pareto front however still needs to consist of only two alternatives – a target and a competitor. We present an extension of the attraction effect procedure using this approach.

4.2 Extended Procedure




Fig. 6: A baseline decision task T_0 and two possible test decision tasks: T_A , where A is the target, and T_B , where B is the target.

The procedure consists of starting with a **baseline decision task** T_0 (see middle of Figure 6 for an example). This baseline decision task has two non-dominated alternatives, A and B . All other alternatives are dominated by A and/or B , and are called **distractors**.

For convenience, we divide the space of all possible alternatives into three dominance regions, shown in Figure 6. If d_A is the region dominated by A (blue hatches in the Figure) and d_B is the region dominated by B (red hatches), then $R_A = d_A \setminus d_B$ (region dominated by A but not by B), $R_B = d_B \setminus d_A$ (region dominated by B but not by A), and $R_{AB} = d_A \cap d_B$ (region dominated by both). In the figure, the baseline decision task contains two distractors per region.

From the baseline decision task one can derive two types of test decision tasks, labeled T_A and T_B in the figure. The decision task T_A is created by adding extra alternatives to the region R_A . Thus T_A only differs from T_0 in that it contains more alternatives that are dominated by A but not by B (twice as many, in this example). These *extra* asymmetrically dominated alternatives are referred to as **decoys**, while A is called the **target** and B the **competitor**. Similarly, the task T_B is created by adding extra decoys to the region R_B , and this time B is the target and A is the competitor.

In case no distractor is included in T_0 and a single decoy is added to T_A and to T_B , we obtain a classical attraction effect experiment. Thus our new definitions for decoy, target and competitor are consistent with the definitions from Section 2.2.1 and generalize them to more complex cases. However, decoys, targets and competitors are always defined with respect to a baseline decision task.

5 BET EXPERIMENT: SCATTERPLOT, MANY CHOICES

We expand our study of the attraction effect to situations where participants are presented with scatterplots with multiple alternatives. We conducted another experiment prior to the one reported here, with a different design and inconclusive results. This inconclusive experiment is reported in a separate research report [14].

5.1 Design Rationale

Here we describe and motivate the design of this new experiment, highlighting the differences with the first (gym) experiment.

5.1.1 Replicated Study

Most attraction effect studies (including our previous experiment) follow a between-subjects design. However, these designs typically suffer from low statistical power. The width of confidence intervals in our gym experiment indicates this was the case there.

We therefore decided to adopt a within-subjects design. Wedell [54] was able to measure a clear attraction effect with numerical tables using a within-subjects procedure, where participants were given multiple decision tasks. He further tried to increase statistical power by *i*) excluding no-decoy conditions and only measuring the combined decoy effect, and *ii*) choosing a scenario with which people were less familiar in an attempt to amplify the effect [36, 41]. We therefore decided to replicate Wedell's design.

5.1.2 Scenario and Attribute Values

Wedell's scenario involved choosing among three lottery tickets, each defined by two attributes: the probability of winning (*probability*), and the amount that can be won (*prize*). Participants were presented with twenty decision tasks in sequence. Each time, three lottery tickets were presented and participants had to choose one. Wedell thought that the abstract nature of the task and of the attributes would reduce possible carry-over effects, such as participants building up strategies based on past choices.

Table 2: The non-dominated alternatives used in our tasks.

	A	B	C	D	E	(F)
probability	0.83	0.67	0.5	0.4	0.3	0.25
prize	\$12	\$15	\$20	\$25	\$33	\$40

The non-dominated alternatives (targets and competitors) used in all Wedell's tasks were taken from a pool of five alternatives (*A* to *E* in Table 2). All had the same expected value of ~\$10. Thus, though a rational decision maker would only need to compare alternatives along a single dimension (expected value), the decision tasks had the same dominance structure as tasks involving two independent attributes such as in the previous gym experiment.

For each possible pair of alternatives in (*A,B,C,D,E*) Wedell generated two decision tasks, one with a decoy on probability, and one with a decoy on prize. We use the notation *XY* to refer to a task where *X* is the target and *Y* is the competitor, and refer to the two decision tasks *XY* and *YX* as *matched*. For example, the pair of alternatives (*A,C*) yields the two matched tasks *AC* (where the decoy is on *A*) and *CA* (where the decoy is on *C*). Wedell's design resulted in 10 pairs of matched decisions tasks (20 tasks in total).

Although we planned to reuse the same targets and competitors, it appeared that the distance between the target and the competitor was visually very small in some scatterplots compared to others. Thus we added an alternative with the same expected value (*F* in Table 2) and excluded all tasks that involved adjacent target/competitor pairs (e.g., *AB*, or *DE*). This new design also resulted in 10 pairs of matched decisions tasks, and 20 tasks in total.

5.1.3 Adding Distractors and Decoys

While Wedell only added one decoy to each of the decision tasks, our goal was to present many alternatives as explained in the previous section. For each pair of matched decision tasks, the procedure consisted of two steps. We explain the procedure for *AC* and *CA* (see results in Figure 7), but it is the same for all other pairs:




Fig. 7: Experimental stimuli for the two matched decision tasks *AC* and *CA* (black-and-white background images), and explanatory annotations (box overlays). See Section 4.2 for the full details.

Step 1. A baseline decision task analogous to T_0 in Figure 6 was created by adding distractors dominated by *A* and/or *C*. One or two distractors (number randomly drawn) were added in each of the regions R_A , R_C and R_{AC} , following a uniform spatial distribution.

Step 2. Two separate decision tasks *AC* and *CA* were then created by adding decoys as shown in Figure 6. For the task *AC* (decoys on *A*), 10 to 20 decoys (number randomly drawn) were added to the region R_A following a bivariate half-normal probability distribution. On each axis, the mode of the half-normal was *A*'s value on this axis, and the mean was this value multiplied by 0.7. The use of half-normals yielded decoys that tend to cluster near *A*, but whose density smoothly decreases with distance to *A* for a more natural look. The same was done for the decision task *CA*.

In both steps, overlaps were eliminated by *i*) defining overlap between two alternatives as a distance less than 0.025 in normalized coordinates (prize divided by 40, probability left unchanged) and *ii*) whenever a new alternative is randomly drawn, iterating until there is no overlap. The reason why the number of alternatives to draw was randomized (i.e., 1–2 for each region in Step 1 and 10–20 in Step 2) was to create more variation across scatterplots and make it more difficult for participants to infer patterns in the experiment.

5.1.4 Ordering of Decision Tasks

Our presentation order for the 20 decision tasks was inspired from Wedell [54], but modified to account for our different set of tasks and for the fact that we present each task on a separate Web page, while Wedell used a four-page paper-and-pencil test.

We created a task ordering such that *i*) a decision task and its matched task (e.g., *AC* and *CA*) are always at least 5 pages apart; and *ii*) the role of an alternative alternates over time. For example, if *D* appears as a target in a task, it will be a competitor the next time it appears. To reduce further possible ordering effects, we created a second ordering where each task is replaced with its matched task. Participants were randomly assigned to each ordering.

To make it more difficult for participants to infer patterns in the sequence of decision tasks, we additionally inserted seven irrelevant decision tasks at various positions, which were not used in our analyses. These tasks differed in that they had either one or three non-dominated alternatives (instead of two), and they did not exhibit an imbalance in the number of asymmetrically dominated alternatives.

5.1.5 Stimuli: Interactive Scatterplots

In this experiment, we added minimal interaction to the scatterplot visualizations. In the first experiment, the scatterplots were static and each data point was labeled with a letter (Figure 1), so that participants could specify their choice through separate radio buttons. As we are now dealing with more data points, labels were removed to prevent clutter (Figure 7), and participants were asked to specify their choice by selecting the data point. Points were highlighted when hovered. Hovering a point also displayed horizontal and vertical projection lines, and the data point's *X* and *Y* values were overlaid on the axes. Such interactions help examine the data and are not uncommon in scatterplot visualizations. After a point

was clicked, its color changed and the participant was asked to confirm her choice by clicking on a button at the bottom of the page.

We added a short flicker during task transitions in order to elicit change blindness and prevent participants from easily detecting similarities and differences between two successive scatterplots.

5.1.6 Crowdsourcing Quality Control

We made two major modifications to the previous procedure: *i*) we added a preliminary tutorial, *ii*) we used a real decision making task where choices affected subsequent monetary gains.

The tutorial simultaneously explained the scenario (the lottery tickets, and what their probability and prize meant), and how to read scatterplots. Although Wedell [54] did not provide similar training, crowdsourcing contributors do not necessarily have the same qualifications as university students, and the notion of probability in particular is known to be challenging [34]. In order to prime participants to use their intuition rather than doing calculations, probability was explained qualitatively rather than quantitatively.

After the tutorial, participants were given a test question consisting of choosing one among 13 lottery tickets presented as a scatterplot. Three tickets were non-dominated (and thus formally uncomparable), and the remaining 10 were considered wrong answers.

In order to better approximate real-life decisions and motivate our participants, we informed them that a computer will run the lottery after the experiment is completed, and for every winning ticket they picked, they will be paid a bonus proportional to the ticket's prize. The use of a real decision task with consequences is common in behavioral economics and is occasionally used when studying the attraction effect (e.g., choosing between objects or money [49]).

Similarly to our previous experiment, we defined our rejection criteria in advance and categorized jobs as *Red* (rejected and not paid), *Orange* (paid but not analyzed) and *Green* (analyzed).

A total of 120 jobs were submitted with a valid completion code. A job was marked as *Red* (12%) if its completion time was abnormal (0.8%), if the contributor failed the tutorial test (11%), or if during the experimental trials, the contributor selected a dominated option more than half of the time (12%). A job was marked as *Orange* (27%) if the contributor always chose the highest probability (27%) or the highest prize (0%). These contributors had a too strong prior preference (in this case, risk aversion) to be sensitive to the attraction effect. The remaining 61% ($N=73$) were marked as *Green*.

5.2 Experiment Design

The design consisted of two within-subjects factors: *task pair* (10 pairs of matched tasks), and *decoy position* (on *probability* or *prize*).

5.2.1 Procedure

We first briefed our crowdsourcing contributors that they will have to choose lottery tickets and will receive a bonus for each winning ticket, for a total of \$0.60 on average. They then opened an external link to the 10-page tutorial. Contributors who chose a valid ticket on the test were told that the ticket won, and that they would get a \$0.10 bonus for the ticket if they proceed and complete the job.

Participants then opened a second external link to the main study, a 31-page form, where they saw the twenty decision tasks, mixed with the seven distractor tasks. After completing all decision tasks, participants rated their overall confidence, their perceived difficulty of the job, their familiarity with gambling games, and whether they knew of the notion of "expected value" in probability. They then filled a short demographic questionnaire.

Finally, participants were presented again with one of the study's scatterplots with the target and competitor labeled *A* and *B*, and were asked whether the higher number of tickets near *A* affected their choices, why they thought there were more tickets, and whether they had this explanation in mind during the study.

All participants received a baseline payment of \$0.20, while *Orange* and *Green* received a bonus of \$0.10 plus a lottery bonus. The expected lottery bonus was \$0.50 if no dominated alternative was chosen, based on a conversion rate of 0.0025 between the scenario's

"virtual dollars" and USD. After the experiment was over, we determined each lottery bonus by *i*) running Bernoulli random draws to determine the winning status of each chosen ticket, *ii*) summing up the prizes of winning tickets *iii*) multiplying by the conversion rate.

5.2.2 Participants

Our participants were 73 crowdflower contributors whose job was marked *Green*. Their demographics, shown in Figure 2, were similar to the first experiment.

5.2.3 Hypotheses

Our statistical hypothesis was **H3**: the mean *attraction score* (as defined in the next section) will be strictly positive.

5.3 Results

One planned analysis for assessing the consistency of participants' responses within matched tasks is not reported for space reasons. We report on all other planned analyses.

5.3.1 Planned Analyses

We first report descriptive statistics of participant choices in a similar way to Wedell [54]. We recorded a total of 1460 choices (73 participants \times 20 decision tasks). We pair choices according to matched tasks (e.g., tasks *AC* and *CA* in Figure 7), yielding $73 \times 10 = 730$ choice pairs. Of all these choice pairs, only 24 (3.3%) included a dominated alternative. Wedell reports similar results (2%), even though his tasks only involved a single dominated alternative.

Wedell '91				Bets			
Decoys on prob				Decoys on prob			
	chose prob	chose prize	total		chose prob	chose prize	total
Decoys on prize							
chose prob	54%	6%	60%	chose prob	59%	8%	67%
chose prize	21%	20%	40%	chose prize	23%	10%	33%
total	75%	25%	100%	total	82%	18%	100%

Fig. 8: Contingency tables showing choice pairs for all matched tasks

Figure 8 (right) summarizes the remaining 706 choice pairs as a contingency table, shown next to Wedell's (on the left). Choice pairs fall into four categories. One is choosing the ticket with highest probability in both tasks (i.e., ticket *A* in Figure 7). This represents 59% of all choice pairs, and is reported in the top-left cell in Figure 8. A second possibility is choosing the ticket with highest prize twice, which represents 10% of all cases. The remaining two possibilities, shown in bold cells, consist in always choosing the target (23%), or always choosing the competitor (8%).

The patterns in our contingency table follow Wedell's closely [54]: participants favoured higher probability overall (reflecting again risk aversion), but when their choice was inconsistent across two matched tasks, they chose the targets more often than they chose the competitors. We now turn to inferential statistics.

Similarly to Wedell, we used as dependent variable an *attraction score*, calculated on a per-participant basis as follows. Each of the 20 decision tasks was assigned a score of 1 when the ticket with highest probability was chosen, a score of 0 when the ticket with highest prize was chosen, and a score of 0.5 when another (dominated) ticket was chosen. Then, we averaged all scores for the 10 decision tasks where the decoys were on probability (yielding a score S_{prob}) and did the same for the 10 tasks where the decoys were on prize (yielding a score S_{prize}). The difference between the two scores $S = S_{\text{prob}} - S_{\text{prize}}$ was the *attraction score*.

A participant who is not subject to the attraction effect should exhibit the same preference for high probability irrespective of the position of the decoys, thus her attraction score should be close to zero. We multiplied the attraction score by 100 to obtain a percentage analogous to the combined decoy effect reported in the gym experiment. The difference here is that the percent difference is com-

puted within-subjects instead of between-subjects, and it incorporates choices of dominated options as “neutral” observations.

The mean attraction score was 15%, with a 95% bootstrap confidence interval of [10%, 22%] (see Figure 4). Thus we have very solid evidence for H3, even if the effect is smaller than in Malkoc's gym study [33]. We cannot directly compare our effects with Wedell's [54] due to the use of different statistical methods, but Figure 8 suggests the effect sizes are comparable.

5.3.2 Additional Analyses

As shown in Figure 5, participants reported various levels of familiarity with gambling and were confident in their choice overall, although slightly less than in the gym experiment. Data on participants' knowledge of expected values was missing due to a bug.

Concerning the final questionnaire on how participants interpreted the presence of decoys (see Section 5.2.1), 8 participants reported not being able to see the scatterplot image, leaving data from 65 participants. When asked whether the uneven distribution of tickets affected their choices, 41% replied “never” or “rarely”, 46% replied “sometimes”, 12% replied “often”, and none replied “always”. When asked why they thought there were more tickets in one region than the other, most (86%) gave responses that were irrelevant or unintelligible based on an informal content analysis of open text responses. Out of the 9 remaining responses, 5 referred to a strategy employed by the lottery organizer (e.g., “*To tempt people to choose tickets of high prize but with low probability, increasing the profitability of lottery owner*”; “*To distract from choosing the higher chances of winning*”), and 4 referred to tickets as past choices from other players (e.g., “*Customers want to win a higher prize*”; “*Maybe more people played the same*”). Only 4 participants (quoted here) reported that they had their explanation in mind while performing the task, while the other 5 reported that it was prompted by our question. Thus there is little evidence that participants' preference for the target was motivated by deliberate, reasoned strategies.

6 GENERAL DISCUSSION AND CONCLUSIONS

Taken together, our two experiments suggest that the attraction effect generalizes to data visualizations. While the first experiment focuses on a traditional procedure with only two or three alternatives, the second experiment shows that the effect can persist with more alternatives. Bettman et al [6] expected that the effect would disappear as more alternatives are added, since pairwise comparisons and dominance recognition becomes hard if numerical tables are used. Our findings suggest that this may not be the case when using visualizations, as visualizations such as scatterplots support fast comparisons and dominance recognition. Overall, our study indicates that when people visualize choice alternatives using scatterplots, the number and position of irrelevant (dominated) alternatives may influence their choice. This shift in preferences violates basic axioms of rational choice theory [25]. In addition to being the first infovis study on the attraction effect, our work contributes to the ongoing debate in decision-making research on whether the effect generalizes to non-numerical formats [17, 26, 48, 56].

6.1 Implications for Design

On a general level, our study indicates that cognitive biases can affect decisions even if the data is well visualized and fully understood, thus traditional visualization design rules may not apply when the goal is to support decision making. This article has not considered debiasing techniques for the attraction effect, but a simple way to eliminate the bias would be to only show the Pareto front, i.e. to hide all dominated options. However, this approach assumes that the system has full knowledge of the user's choice criteria, which may not be the case in practice. In addition, dominated options can help understand dataset trends, and may in some cases provide useful context when making decisions. Thus, debiasing techniques should only be available as options, and activated on demand. Alternatively, one could consider techniques such as de-emphasizing dominated options or highlighting the Pareto front, but the effectiveness of such techniques remains to be experimentally tested.

6.2 Limitations

There are several potential limitations to our study. One stems from a general criticism of cognitive bias research, namely, that heuristics that appear irrational may not be so upon deeper examination [18]. Concerning the attraction effect, the way dominated alternatives are distributed could in some cases provide relevant information. For example, a real estate investor may infer from a region with many dominated alternatives that a certain type of house is more common, and therefore represents a larger market. At the same time, situations also exist where the number and position of dominated alternatives is clearly irrelevant and where a preference for the target would be irrational. This was the case for our experiment involving real bets, and our data does indicate that the vast majority of our participants were unable to rationalize their choices based on where the dominated alternatives were located.

Although we have observed attraction effects, we did not investigate why they occur. In particular, we do not know how much of the effect has cognitive vs. perceptual causes. Since in the bet experiment regions with many decoys were visually more salient, it is possible that they drew participants' attention towards the target, or similarly, that participants sometimes failed to see the competitor because it was an isolated point. This possibility does not invalidate the existence of an attraction effect (as defined in Sections 2.2.1 and 4.2), but it does raise the possibility that part of the effect with scatterplots (but not with numerical tables) has perceptual origins.

Finally, we tested very specific datasets, i.e., synthetically-generated datasets with only two non-dominated options and a large number of decoys. More realistic datasets need to be tested, although our inconclusive results with real datasets suggest that the effects may be small and hard to measure [14].

6.3 Future work

While our work is a first step in investigating the attraction effect in visualizations, much more work is needed. More realistic datasets and decision making situations remain to be tested. We also focused on scatterplots, but clearly other commonly used visualizations need to be evaluated to assess whether the effect persists across visual encodings. Other cognitive biases [19] remain to be studied, both in isolation (as we did here), and in combination. How cognitive biases interact with visual perception is also an important and difficult question that has remained largely unexplored.

ACKNOWLEDGMENTS

We thank S.A. Malkoc for sharing their material and D.H. Wedell for their study advice, our colleagues J.-D. Fekete, F. Vernier, T. Isenberg, P. Goffin and L. Besancon for feedback on the paper, and our reviewers for their insightful comments and suggestions.

REFERENCES

- [1] D. Ariely and T. S. Wallsten. Seeking subjective dominance in multi-dimensional space: An explanation of the asymmetric dominance effect. *Organizational Behavior and Human Decision Processes*, 63(3):223–232, 1995.
- [2] T. Asahi, D. Turo, and B. Shneiderman. Using treemaps to visualize the analytic hierarchy process. *Information Systems Research*, 6(4):357–375, 1995.
- [3] B. A. Aseniero, T. Wun, D. Ledo, G. Ruhe, A. Tang, and S. Carpendale. Stratos: Using visualization to support decisions in strategic software release planning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 1479–1488. ACM, 2015.
- [4] M. Bateson, S. D. Healy, and H. Hurly. Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, 63(3):587–596, Mar. 2002.
- [5] J. Bautista and G. Carenini. An empirical evaluation of interactive visualizations for preferential choice. In *Proceedings of the working conference on Advanced visual interfaces*, pages 207–214. ACM, 2008.
- [6] J. R. Bettman, M. F. Luce, and J. W. Payne. Constructive Consumer Choice Processes. *J. of Consumer Research*, 25(3):187–217, Dec. 1998.
- [7] W. M. Bolstad. *Introduction to Bayesian statistics*. J. Wiley & Sons, 2013.
- [8] J. Boy, R. Rensink, E. Bertini, and J.-D. Fekete. A principled way of assessing visualization literacy. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):1963–1972, Dec 2014.

- [9] S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [10] J. M. Choplin and J. E. Hummel. Comparison-induced decoy effects. *Memory & Cognition*, 33(2):332–343, 2005.
- [11] M. Correll and M. Gleicher. Error bars considered harmful: Exploring alternate encodings for mean and error. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2142–2151, 2014.
- [12] G. Cumming and S. Finch. Inference by eye: confidence intervals and how to read pictures of data. *American Psychologist*, 60(2):170, 2005.
- [13] M. Daradkeh, C. Churcher, and A. McKinnon. Supporting informed decision-making under uncertainty and risk through interactive visualisation. In *Proc. Fourteenth Australasian User Interface Conference-Vol 139*, pages 23–32. Australian Computer Society., 2013.
- [14] E. Dimara, A. Bezerianos, and P. Dragicevic. Testing the Attraction Effect on Two Information Visualization Datasets. Research Report RR-8895, Inria, Mar. 2016. <http://hal.inria.fr/hal-01295624>.
- [15] P. Dragicevic. Fair statistical communication in HCI. In *Modern Statistical Methods for HCI*, pages 291–330. Springer, 2016.
- [16] N. Elmqvist, P. Dragicevic, and J.-D. Fekete. Rolling the Dice: Multidimensional Visual Exploration using Scatterplot Matrix Navigation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1141–1148, Nov. 2008. Best Paper Award.
- [17] S. Frederick, L. Lee, and E. Baskin. The limits of attraction. *Journal of Marketing Research*, 51(4):487–507, 2014.
- [18] G. Gigerenzer. Why heuristics work. *Perspectives on psychological science*, 3(1):20–29, 2008.
- [19] M. L. Graber, S. Kissam, V. L. Payne, A. N. Meyer, A. Sorensen, N. Lenfestey, E. Tant, K. Henriksen, K. LaBresh, and H. Singh. Cognitive interventions to reduce diagnostic error: a narrative review. *BMJ Quality & Safety*, pages bmjqs–2011, 2012.
- [20] T. B. Heath and S. Chatterjee. Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, 22(3):268–284, 1995.
- [21] W. Hedgcock and A. R. Rao. Trade-off aversion as an explanation for the attraction effect: A functional magnetic resonance imaging study. *Journal of Marketing Research*, 46(1):1–13, 2009.
- [22] W. Hedgcock, A. R. Rao, and H. A. Chen. Could ralph nader's entrance and exit have helped al gore? the impact of decoy dynamics on consumer choice. *Journal of Marketing Research*, 46(3):330–343, 2009.
- [23] J. Heer and M. Bostock. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proceedings of CHI 2010*, pages 203–212. ACM, 2010.
- [24] S. Highhouse. Context-dependent selection: The effects of decoy and phantom job candidates. *Organizational Behavior and Human Decision Processes*, 65(1):68 – 76, 1996.
- [25] J. Huber, J. W. Payne, and C. Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9(1):90–98, 1982.
- [26] J. Huber, J. W. Payne, and C. P. Puto. Let's be honest about the attraction effect. *Journal of Marketing Research*, 51(4):520–525, 2014.
- [27] J. Huber and C. Puto. Market boundaries and product choice: Illustrating attraction and substitution effects. *Journal of Consumer Research*, pages 31–44, 1983.
- [28] D. Kahneman. *Thinking, fast and slow*. Macmillan, 2011.
- [29] A. Khan, S. Breslav, M. Glueck, and K. Hornbæk. Benefits of visualization in the mammography problem. *International Journal of Human-Computer Studies*, 83:94–113, 2015.
- [30] U. Khan, M. Zhu, and A. Kalra. When trade-offs matter: The effect of choice construal on context effects. *Journal of Marketing Research*, 48(1):62–71, 2011.
- [31] T. Latty and M. Beekman. Irrational decision-making in an amoeboid organism: transitivity and context-dependent preferences. *Proceedings of the Royal Society B: Biological Sciences*, 278(1703):307, 2011.
- [32] S. Lelis and A. Howes. Informing decisions: how people use online rating information to make choices. In *Proceedings of CHI 2011*, pages 2285–2294. ACM, 2011.
- [33] S. A. Malkoc, W. Hedgcock, and S. Hoeffler. Between a rock and a hard place: The failure of the attraction effect among unattractive alternatives. *Journal of Consumer Psychology*, 23(3):317 – 329, 2013.
- [34] L. Micallef, P. Dragicevic, and J.-D. Fekete. Assessing the effect of visualizations on bayesian reasoning through crowdsourcing. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2536–2545, 2012.
- [35] S. Miller, A. Kirlik, A. Kosorukoff, and J. Tsai. Supporting joint human-computer judgment under uncertainty. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 52, pages 408–412. SAGE Publications, 2008.
- [36] S. Mishra, U. Umesh, and D. E. Stem Jr. Antecedents of the attraction effect: An information-processing approach. *Journal of Marketing Research*, pages 331–349, 1993.
- [37] T. Munzner. *Visualization Analysis and Design*. CRC Press, 2014.
- [38] Y. P. S. O'Curry and R. Pitts. The attraction effect and political choice in two elections. *Journal of Consumer Psychology*, 4(1):85–101, 1995.
- [39] Open Science Collaboration et al. Estimating the reproducibility of psychological science. *Science*, 349(6251), 2015.
- [40] J. C. Pettibone. Testing the effect of time pressure on asymmetric dominance and compromise decoys in choice. *Judgment and Decision Making*, 7(4):513–523, 2012.
- [41] S. Ratneshwar, A. D. Shocker, and D. W. Stewart. Toward understanding the attraction effect: The implications of product stimulus meaningfulness and familiarity. *J. of Consumer Research*, pages 520–533, 1987.
- [42] S. Rudolph, A. Savikhin, and D. S. Ebert. Finvis: Applied visual analytics for personal financial planning. In *IEEE Symposium on Visual Analytics Science and Technology*, pages 195–202. IEEE, 2009.
- [43] A. Savikhin, R. Maciejewski, and D. S. Ebert. Applied visual analytics for economic decision-making. In *Visual Analytics Science and Technology, VAST'08. IEEE Symposium on*, pages 107–114. IEEE, 2008.
- [44] S. Sen. Knowledge, information mode, and the attraction effect. *Journal of Consumer Research*, 25(1):64–77, 1998.
- [45] S. Shafir, T. A. Waite, and B. H. Smith. Context-Dependent Violations of Rational Choice in Honeybees (*Apis mellifera*) and Gray Jays (*Perisoreus canadensis*). *Behavioral Ecology and Sociobiology*, 51(2):180–187, 2002.
- [46] A. Shen-Hsieh and M. Schindl. Data visualization for strategic decision making. In *Case Studies of the CHI2002*, pages 1–17. ACM, 2002.
- [47] I. Simonson. Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, 16(2):158–174, 1989.
- [48] I. Simonson. Vices and virtues of misguided replications: The case of asymmetric dominance. *J. of Marketing Research*, 51(4):514–519, 2014.
- [49] I. Simonson and A. Tversky. Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29:281–295, 1992.
- [50] J. J. Thomas and K. A. Cook. *Illuminating the path: the research and development agenda for visual analytics*. IEEE Computer Society, 2005.
- [51] J. Trueblood. Multialternative context effects obtained using an inference task. *Psychonomic Bulletin and Review*, 19(5):962–968, 2012.
- [52] J. S. Trueblood, S. D. Brown, A. Heathcote, and J. R. Busemeyer. Not just for consumers context effects are fundamental to decision making. *Psychological science*, 24(6):901–908, 2013.
- [53] C. Ware. *Information visualization: perception for design*. Elsevier, 2012.
- [54] D. H. Wedell. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(4):767, 1991.
- [55] C. Williamson and B. Shneiderman. The dynamic homefinder: Evaluating dynamic queries in a real-estate information exploration system. In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '92*, pages 338–346. New York, NY, USA, 1992. ACM.
- [56] S. Yang and M. Lynn. More evidence challenging the robustness and usefulness of the attraction effect. *Journal of Marketing Research*, 51(4):508–513, 2014.
- [57] J. S. Yi. *Visualized decision making: development and application of information visualization techniques to improve decision quality of nursing home choice*. PhD thesis, Georgia Institute of Technology, 2008.
- [58] J. S. Yi, R. Melton, J. Stasko, and J. A. Jacko. Dust & magnet: multivariate information visualization using a magnet metaphor. *Information Visualization*, 4(4):239–256, 2005.
- [59] Y. Zhang, R. K. Bellamy, and W. A. Kellogg. Designing information for remediating cognitive biases in decision-making. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2211–2220. ACM, 2015.
- [60] C. Ziemkiewicz and R. Kosara. Laws of attraction: From perceptual forces to conceptual similarity. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1009–1016, 2010.
- [61] T. Zuk and S. Carpendale. Visualization of uncertainty and reasoning. In *Smart Graphics*, pages 164–177. Springer, 2007.