



HAL
open science

Designing Resource Allocation Tools to Promote Fair Allocation: Do Visualization and Information Framing Matter?

Arnav Verma, Luiz Morais, Pierre Dragicevic, Fanny Chevalier

► To cite this version:

Arnav Verma, Luiz Morais, Pierre Dragicevic, Fanny Chevalier. Designing Resource Allocation Tools to Promote Fair Allocation: Do Visualization and Information Framing Matter?. CHI 2023 - Conference on Human Factors in Computing Systems, ACM, Apr 2023, Hamburg, Germany. 10.1145/3544548.3580739 . hal-04074164

HAL Id: hal-04074164

<https://inria.hal.science/hal-04074164v1>

Submitted on 19 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Designing Resource Allocation Tools to Promote Fair Allocation: Do Visualization and Information Framing Matter?

Arnav Verma
arnav@cs.toronto.edu
University of Toronto
Toronto, Canada

Pierre Dragicevic
pierre.dragicevic@inria.fr
Université de Bordeaux, CNRS, Inria, LaBRI
Bordeaux, France

Luiz Morais
luiz.morais@inria.fr
Inria
Bordeaux, France

Fanny Chevalier
fanny@cs.toronto.edu
University of Toronto
Toronto, Canada

ABSTRACT

Studies on human decision-making focused on humanitarian aid have found that cognitive biases can hinder the fair allocation of resources. However, few HCI and Information Visualization studies have explored ways to overcome those cognitive biases. This work investigates whether the design of interactive resource allocation tools can help to promote allocation fairness. We specifically study the effect of presentation format (using text or visualization) and a specific framing strategy (showing resources allocated to groups or individuals). In our three crowdsourced experiments, we provided different tool designs to split money between two fictional programs that benefit two distinct communities. Our main finding indicates that individual-framed visualizations and text may be able to curb unfair allocations caused by group-framed designs. This work opens new perspectives that can motivate research on how interactive tools and visualizations can be engineered to combat cognitive biases that lead to inequitable decisions.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in visualization.**

KEYWORDS

framing, cognitive bias, visualization, resource allocation, donation

ACM Reference Format:

Arnav Verma, Luiz Morais, Pierre Dragicevic, and Fanny Chevalier. 2023. Designing Resource Allocation Tools to Promote Fair Allocation: Do Visualization and Information Framing Matter? . In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3544548.3580739>

1 INTRODUCTION

When a tornado tears down a town in Florida, North American donors and charities may provide an excess of monetary support to help with the rebuilding process while leaving little financial aid to Haiti, devastated by a similar but larger-scale disaster. This situation can occur because people tend to be biased in favor of humanitarian causes which they feel socially connected to [15]. This is problematic because some communities in need of urgent care might end up with less support than they should receive. This is an example of how *cognitive biases* can negatively affect decision making: if decision makers only trust their intuition, available resources may end up poorly distributed.

There are tools available to help or motivate people to donate. Donation websites such as donorchoose.org or fundraising platforms like gofundme.com act as tools which let users search for similar causes to which allocate their resources, but comparison across these causes is not supported. Prior research has focused on encouraging people to contribute more towards a charity using tools such as conversational agents [65], web interfaces [47, 55], and visual media like photographs and visualizations to elicit affective emotions such as compassion [28, 40]. These aim to elicit *more* donations, but little is known about tools to help people allocate available charitable resources more equally.

While there are many aspects to consider when making a resource allocation (such as the relative effectiveness of different programs [32]) cognitive biases play a major role in humanitarian aid decisions [4, 6, 66]. A charitable donor might struggle to decide how to split their personal donations among charities, or decision makers working on behalf of charities or governments may be confused about which programs should be given priority, when distributing the collected donations. Many cognitive biases have been observed in prior research in cognitive and social psychology [24], especially when making a decision which requires the allocation of resources, like time or money [5, 11, 57]. We identify that in the context of humanitarian resource allocation, prevalent cognitive biases stem from reasoning about people in need of aid as groups or communities, as opposed to considering each individual separately (§2.1). However, how to curb these biases remains an open question [45].

Motivated by the promise of visualizations to mitigate bias [25, 27, 44, 59, 68], we investigate how interactive visualization tools can be designed to aid people achieve more fair and equitable resource

allocation decisions across multiple groups. More specifically, our first core research question is:

(RQ_{vis}) – Can visualization improve allocation fairness when compared to text?

Our work expands on research that has studied the use of visualization for humanitarian purposes, which looked at eliciting donations [32], or prompting emotional connection to data [53]. Our study addresses other biases.

We also explore the effect of a specific information framing strategy, i.e. showing resources allocated to individuals instead of showing resources allocated to groups. Our research is grounded in past work looking at the effects of information framing on decision making tasks [41], and nudging in computational interfaces [19, 55]. These works suggest that the way information is presented can influence people's choices, but did not explore effects on fair allocations. Our other core research questions are thus:

(RQ_{frm}) – Can presenting resource allocation information on a per-individual basis (showing the amount allocated to each person) improve fairness compared to doing so on a per-group basis (showing the amount allocated to each cause)?

(RQ_{int}) – How do presentation format and information framing interact with one another?

We contribute the results of three crowdsourced experiments with large samples ($N = 239, 244, 495$) where participants were presented with different visual and textual framing in an interactive allocation tool used to allocate resources between two hypothetical charitable programs that are similar but provide support for populations of different sizes. Our main finding indicates that individual-framed visualizations and text may be able to curb unfair allocations caused by group-framed designs. Complementary findings are described in the sections about the experiments. This work opens new perspectives that can motivate research on how interactive tools and visualizations can be engineered to combat cognitive biases that lead to inequitable decisions.

2 RELATED WORK

Our work draws from prior work studying cognitive biases, the effect of framing information in text and visualizations, and tool designs relevant to supporting resource allocation decisions.

2.1 Cognitive Biases

Efficient response to humanitarian crises requires pragmatic, unbiased decision making. But people can be prone to biases which are well documented in psychology (see comprehensive reviews [15, 20]). A bias is a type of thinking error that is not necessarily caused by lack of knowledge or information, but rather by applying a decision rule that is helpful in some situations but not in others. In the context of charitable giving, cognitive biases can be harmful and may “lead to the systematic misallocations of funds and waste of resources” [3].

“If I look at the mass I will never act. If I look at the one, I will.” This famous quote captures a powerful phenomenon where the way people conceptualize individuals in need of help directly impacts their response to the issue. Our work is interested in this phenomenon and cognitive biases at play when allocating resources between charitable programs.

The more the people in need (“the mass”), the less we feel compelled to aid. The *compassion fade* or *compassion fatigue* effect is a cognitive bias that refers to a decrease in compassionate intent and helping behaviour as the number of people needing aid increases [8, 11, 50, 61]. For instance, one might allocate \$300 to a charity program supporting 10 children, but would not proportionally increase this amount for a program supporting 500 children.

The opposite is also true! The less the people we can help, the less we feel compelled to aid. The *drop-in-the-bucket* effect refers to when people feel that helping few people accomplishes nothing [5]. Our work is interested in whether either of these effects generalizes to resource allocations, i.e. whether people tend to allocate resources to the larger or smaller group.

It could also be the case that neither apply. The *diversification bias* effect states that people tend to allocate available resources evenly over different choices in contexts like stock investment or product consumption [58], thereby discounting differences between the different options. Our work expects to identify whether this effect happens in the context of humanitarian resource allocation with visualizations.

Recent research also suggests that the effect of *entitativity*, which shows bias relative to perception of group belonging, can be another factor influencing reasoning [42]. For instance, children who are identified as part of a family are more likely to collectively receive more aid than children who are seen as a set of independent individuals without explicit group membership, or as statistics [62]. In the context of allocating resources between programs, we identify a tension between preserving entitativity of each group, and emphasizing individuality to focus on what amount of resources each person receives regardless of which program they are aided through.

Our study allows to evaluate whether tool designs are successful at curbing bias. We manipulate the number of people in two charitable programs, as well as the way information is presented (see §2.2). This allows us to interpret results as to whether and why different resource allocation tool designs result in participants' resource allocation strategies that tend to favor the larger, smaller, or none of two groups, or whether individuals are equally served.

2.2 Framing Effects

Tools that support decision making are concerned with the effective presentation of information relevant to the decision. We explore designs that vary in the way information is framed, which research has shown can influence how people interpret information and make subsequent decisions – a cognitive bias referred to as *framing effect* [41]. Framing is a broad term, and there are numerous ways information can be framed. Framing effects can be referred to as any “small changes in the presentation of an issue or an event, such as slight modifications of phrasing, [that can] produce measurable changes of opinion” [39].

Text can be manipulated to communicate the same fact, but with a different framing. Take for instance a glass “half full” or “half empty”. The underlying data is the same, but saying one or the other may influence how people interpret the same bit of information. Work by Kahneman and Tversky [41] showed that different

framings about loss of human lives resulted in people making different decisions regarding actions they took. For instance, people are more sensitive to text saying: “Of 600 people, 200 will be saved” than they are to the statement “Of 600 people, 400 people will die”.

Hullman and Diakopoulos [39] propose a taxonomy of framings in the context of communicative visualizations, which covers changes to the data (e.g., aggregating data) and the way it is represented (e.g., using visual metaphors, visual grouping). Studies have explored the influence of visualization framing on people’s behaviors [39], interpretation of data [14, 30], and trust in information [17]. A comprehensive review is beyond the scope of this paper, as it can be argued that the whole field deals with optimizing visualization design to achieve fair representation of data.

Since we study how people split resources between two groups, we are interested in a specific framing effect: showing resources each *group* gets as a result of the decision (group framing) vs. resources each *individual* gets (individual framing). This focus stems from the nature of our research question and of the task itself: since cognitive biases covered in §2.1 can be conceptualized as putting too much weight on entitative groups and not enough on individuals, we deem it important to look at whether individual framing vs. group framing can impact allocation fairness (RQ_{frm}).

2.3 Design for Resource Allocations

Researchers have argued that computer tools can help overcome cognitive biases in decision making involving humanitarian scenarios [20, 21]. However, we still lack comprehensive guidelines to avoid cognitive biases in decision making [45, 68]. To the best of our knowledge, there is no existing review of user interfaces that funding managers use in such scenarios, and research on developing interfaces to support unbiased resource allocations is scarce.

Our study focuses on identifying if and how visualization can improve allocation fairness when compared to text (RQ_{vis}). Allocated resources can be represented with simple visual media such as images or bar charts. However, studies have also focused on the design of the iconography and its influence on potentially aiding users to grasp concepts such as scale or help induce affective responses. Concrete scales [17] exploit visceral responses to the volume occupied by familiar objects re-expressing otherwise difficult-to-grasp quantities. Such representations can help better apprehend the amount of resources and people, and is a common strategy to raise awareness on casualties in a disaster (e.g. memorials using flags or shoes to represent victims). Anthropomorphized data graphics, which portray characteristics of humans, is another commonplace visual communication strategy employed as a potential means to elicit more prosocial behaviours [9, 54]. We note that conveying data through visual data storytelling to elicit empathy has had limited success, and prior studies do not indicate strong effects [54]. Yet, these approaches may help create more relatable iconography when visualizing people in contexts where cognitive biases tend to put too much weight on entitative groups. Our approach builds on this premise, but other designs could be explored in future work.

Prior studies have also investigated presentation and framing within interfaces and visualizations that help mitigate bias. A tool allowing to remove irrelevant information locally within a visualization to make a comparative decision has been shown to mitigate

biases such as the attraction effect [25]. Visualizations which show predictions of future outcomes have been shown to encourage less optimistic estimates [44], and interaction with visualizations to present information from different perspectives has been proposed to help combat anchoring effects [59]. Our work extends this line of research, by investigating the role of visualization and information framing on a set of cognitive biases unique to allocating resources in a humanitarian aid context (§2.1).

Also relevant are tools that allow to explore the result of alternative scenarios, through direct manipulation and visualization. Their design can serve as inspiration for features which allow for the user to visually compare the effects of different allocations. For instance, the ‘What-if Tool’ lets machine learning practitioners evaluate the effect of manipulating input data on machine learning models performance, such as ML fairness metrics [69]. Through interactive widgets, users can change aspects of the datasets and model parameters, and see the results of such hypothetical situations. Similarly, tools like the climate change calculator [10] allow viewers to visually explore the effect of varying government action scenarios on the future climate. Our tool borrows from these past works, by allowing users to dynamically adjust their allocation decision, and get immediate feedback on its result through visualization.

3 A TOOL FOR RESOURCE ALLOCATION: PROBLEM AND DESIGN CONSIDERATIONS

Imagine that you decide to contribute to humanitarian causes – how should you split your money if you had two different causes to aid? Or think about how organizations like the United Nations should allocate their funds between the programs they support. Those are examples of the **resource allocation problem**, which happens when two or more groups should be benefited by the allocation of resources. The question is: how to make effective or fair allocations based on the information that the decision-maker has available?

The problem is complex because in real-world scenarios groups can have completely different characteristics. Those factors can play a role in how people make judgements to allocate money. If people rely on their own experience or intuition, they might be influenced by cognitive biases and their attitudes. If decision makers decide to allocate resources based on data but without a proper tool, they can be overwhelmed by the amount of information.

This paper is a first attempt to explore the design of resource allocation tools to promote fair allocations. We consider resources to be *fairly allocated* when each beneficiary receives the same amount of resources as any other beneficiary, irrespective of the group they belong to. For the sake of simplicity and reproducibility, we focus on a task in which all individuals from two groups have comparable needs and contexts, with only the group sizes changing. This way, we make sure that the same amount of resources would benefit different individuals equally. This decision helps breaking the resource allocation problem down into a simpler problem and controls for confounding factors in the experiments we conducted. Possible limitations for this approach are addressed in §8.2.

We are specifically interested in the effect of two design dimensions on resource allocation. **Presentation format** corresponds to how the resource allocation information is represented – for example, we can use *text* to describe the resource split across groups, or a

visualization. **Information framing** refers to how precisely the allocation information is conveyed to the user – for example, a *group* framing describes the amount donated to each group, whereas an *individual* framing focuses on how much would be allocated to each individual. We examined how presentation format (RQ_{vis}) and information framing (RQ_{frm}) affect resource allocation by designing and testing interactive tools that vary in terms of those dimensions. In the following, we describe examples of designs that combine different presentation formats and information framings, and introduce the interactive tools used in our experiments.

3.1 Textual Information

Information needed to make resource allocation can often be composed of a textual description of a problem accompanied by other presentation formats such as a picture or a visualization. The way the text is framed can affect people's judgements [16]. We explore two ways of framing the textual information as follows:

When **textual information is framed by group**, the message is expressed as a total donation to each group (e.g., "give \$1,000 to aid the victims of Florida's hurricane and \$300 to help victims of Haiti's disaster"). This framing allows readers to think holistically about the allocation given to the respective groups. This can induce the effect of entitativity [62] as it can encourage the donor to associate the group of people in each program to singular group units, possibly resulting in compassion fade [61] or drop-in-the-bucket [5]. Group framing can also result in the naive diversification bias [6], when people prefer to allocate resources evenly across groups, even when these have different characteristics.

When **textual information is framed by individuals**, the message is expressed in terms of the amount received per person (e.g., "give \$15 to each victim of Florida's hurricane and \$12 to each victim of Haiti's disaster"). This framing allows readers to reason in terms of what individuals in the respective communities each receive. This might alleviate the entitativity and diversification biases, since individuality is directly expressed, as well as the effect of compassion fade, because large numbers become more palatable when expressed as what that means for individuals.

3.2 Visual Information

For the display of information using a data visualization, we opt for an iconic representation [9] of victims (i.e., icons of a human), and resources (i.e., coins). Like for Isotypes [56], the number of icons directly encodes quantities, using partial rendering of icons when necessary (see Figure 1). To vary framing, we manipulate the organization of icons representing quantities as follows:

Visual information can be framed by group using Gestalt psychology [43] to create a perception of unity for each group. Icons representing people from the same groups can be located together so that they are perceptually perceived as a unit. Similarly, icons representing resources can be grouped together spatially to be interpreted as an overall amount of resources allocated to the group. We chose spatial grouping (see an example in Figure 2-c), but other perceptual groupings could be considered in future designs.

Visual information can be framed by individuals using the same Gestalt principles, combining visual representations that correspond to each beneficiary. An individual framing can be created

by placing the exact amount of resources allocated to each individual, besides each individual (see an example in Figure 2-d). This way, viewers can easily perceive the result of splitting one group's allocation across individuals in that group.

We used an iconic representation to show individual people because it makes it easier for viewers to understand the true scale of a community. Although it is possible to use aggregated visual encodings such as bar charts to convey population sizes, such an aggregated representation would not work well in an individual framing, where the goal is to emphasize the share of resources each individual gets. In contrast, a unit chart representation where each person is represented by a symbol works well with both group and individual framings. Unit charts are also common when conveying information about populations, such as in data journalism [53].

3.3 An interactive tool for resource allocation

We created a resource allocation tool that combines different presentation formats and framings. Figure 1 shows a screenshot of the interface we used in our study. The tool encompasses: (a) a legend, which informs the number of people and resources encoded by each mark in the visualization; (b) a display, showing textual and/or visual information about the groups and corresponding allocations; and (c) a slider, which allows the user to change the amount of resources to each group. Since we only consider 2-group allocations, the display is divided in two parts; one for each program/group. The display dynamically updates when the user moves the slider to the left (to benefit group Alpha) or to the right (to benefit group Zefa). That template was used for all of our experiments.

To study how the two dimensions affect resource allocation, we manipulated the tool design in terms of presentation format and information framing. We tested designs composed of only text, visualization, or both. Also text and visualization had group or individual framings. In the example shown in Figure 1, the display encompasses both textual and visual information. The text follows an individual framing because it contains the exact amount to be donated for each individual (in blue). The visualization also follows an individual framing because the coins encoding the values allocated to each beneficiary are aligned with that representing individuals. As both text and visualization have the same framing, we consider that the tool has a **congruent framing**. Conversely, when the text and visualization have opposite framings, the tool has an **incongruent framing**. There are many other aspects to the design of a resource allocation tool that could be considered, which we cover in the discussion section.

4 EXPERIMENTAL DESIGN: OVERVIEW

We conducted three experiments to gain a better understanding of how people approach the resource allocation problem in a fictional scenario. Each experiment builds on findings and lessons from the previous experiment(s). We describe the general context and procedures for our study. Elements specific to a particular experiment are discussed in the corresponding sections. All experimental materials and scripts were pre-registered for each experiment.¹

¹All pre-registrations can be found at <https://osf.io/zrbej/>; Experiment 1: <https://osf.io/xqser/>; Experiment 2: <https://osf.io/7ebsm/>; Experiment 3: <https://osf.io/egv85/>



Figure 1: Screenshot of our instance of an interactive resource allocation tool featuring two charitable programs, used in Experiment 1 of our study. With the help of the legend (a), a donor can directly assess the result of resource allocation presented in the information display (b). Here both text and visualization representations are displayed, each presenting allocations with an individual framing. The donor can dynamically adjust how the resources are split between the two programs, using the slider (c), which automatically updates the information presented in the information display.

4.1 Task

We created a fictional scenario for the resource allocation task. Participants were provided with a brief about a charitable organization which is helping setting two separate programs for impoverished children. The scenario brief specifies that the charity’s mission is to aid children, and that the money to allocate will be of significant value to accomplish this goal, followed by instructions on the task:

“Your task is help the organization decide how to distribute money between two of their programs.”

The two programs, Alpha and Zefa, differ in size: Alpha aids less children than Zefa does (sizes are reported for each experiment). The brief specifies how the money will be utilized – “the money will go toward food, clothes, and education”, but it includes little information about individuals, as we wanted to measure the presentation of summarized information rather than the presentation of identifiable victims (which comes with biases outside our scope [31, 40]).

Our scenario was inspired by the scenarios found in related psychological studies [33]. We chose money because it is the main resource that decision makers deal with in charitable giving contexts, and monetary donations is common in similar studies from psychology [66], economics [60], and HCI [9, 54]. We chose to keep population sizes relatively small, i.e not in the thousands or millions, since a smaller number can be easier for participants to reason about both when shown in text and when encoded in a visualization. We believe this is a reasonable choice as similar studies

showed that compassion fade stems from the proportional difference between populations [50], meaning that effects observed on small populations would theoretically hold for population sizes tens or hundreds times larger [33].

4.2 Procedure

We deployed our experiment online. Upon sign up, participants were directed to a self-hosted website where they were first prompted for consent. Participants who consented were presented with the scenario and task (see §4.1). When proceeding to the next screen, they were presented with the interactive tool to make their allocation (see §3.3). Participants were not instructed anything more than allocating resources between the two programs as they saw fit. No indication of performance goals was provided, i.e. participants were not hinted at what would be a reasonable or desirable split of resources, neither were they given time or speed constraints.

After submitting their chosen allocation, participants were asked to explain their rationale in a text box. We did not inform participants that they were going to be asked to justify their decision until after they decided how to split the resources between the programs, and participants were not allowed to return to the tool to change their allocation. We did so to minimize potential social desirability bias [29] which may arise when one feels they will be held accountable for their decision, which is also why we opted against a donation intention question for our task. Participants

were also asked to rate confidence in their choice, on a five-point Likert scale (“not at all confident” to “very confident”).

Finally, participants were posed a multiple-choice question about the purpose of the charity, which we used as an attention check to exclude potentially inattentive participants or bots, whose responses would introduce noise.

4.3 Measurements

Our primary measurement is obtained from the slider (Figure 1-c), which records how the full amount of available resource was split between the two programs. We developed metrics to gauge the extent to which allocations were **fair** (or **unfair**), described in each Experiment’s section. We also characterized what different allocation values mean in terms of **allocation strategy** employed, as follows:

- The **fair allocation** strategy treats all children equally, i.e. money is allocated such that individual people in the two programs get the same amount of money. This allocation favors the larger program (Zefa), because more money is allocated to it compared to Alpha, but this strategy treats individuals equally, which can be conceptualized as the allocation strategy that is fair towards individuals.
- The **unfair allocation** strategy encompasses any other allocation, where children in different programs are aided a different amount of resources.

We further break down the **unfair allocation** strategy into four subcategories as follows:

- The **minority-skewed allocation** strategy favors the program with less people, i.e., more money is allocated to the program with the smaller number of children. The allocation favors both the smaller program (Alpha) and the individual people in this program.
- The **naive allocation** strategy treats the programs equally, i.e., the same amount of money is allocated to the two charity programs. The allocation treats charity programs equally but favors individual people in the smaller program (Alpha). This response is indicative of diversification bias [6].
- The **mixed allocation** strategy balances causes and beneficiaries, i.e., more money is allocated to the program with the larger number of children (Zefa), but individual people from the smaller program (Alpha) still get more money. This type of response is indicative of compassion fade effect [11].
- The **majority-skewed allocation** strategy favors the program with more children, i.e., more money is allocated to individual people in the charity program with the larger number of children (Zefa). The allocation favors both the larger program and the individual people in this program, and indicates a drop-in-the-bucket bias [5].

4.4 Analyses

We follow an estimation approach to statistical reporting, which focuses on reporting effect sizes and interval estimates instead of p -values, and encourages non-directional research questions [22, 26]. We draw our primary inferences from graphically-reported point and interval estimates [23]. Effects are estimated using the percentile bootstrap method to obtain 95% confidence intervals

(CIs). We distinguish between planned (preregistered) and post-hoc analyses. Instead of adjusting for multiplicity, we distinguish between primary and secondary research questions [7, 48]. We also performed a qualitative analysis of the content of the free-text responses to the question asking participants to justify their choice. Methods used are described for each Experiment.

4.5 Recruitment, Inclusion & Exclusion Criteria

Participants were recruited through the Prolific crowdsourcing platform, which was chosen for its larger set of positive actors compared to other similar platforms [37]. In a pre-screening process, participants were required to speak English fluently, to have a 99% approval rate on Prolific, and to have not participated in any of our pilots to be eligible. We also recruited only from majority English speaking countries (UK, USA, Canada). We excluded data from participants who did not complete the study, or failed at the attention test (see §4.2).

4.6 Implementation

We deployed our study tool as a standalone web-based page compatible with most browsers. Our tool is built in Typescript with React.js, D3.js, and Tailwind.css libraries for the frontend, and Node.js for the backend. Besides practical considerations for a smooth crowdsourced study experience, we also chose a web-based implementation as future extensions of our tool may be integrated into larger websites such as donorschoose.org or gofundme.com.

5 EXPERIMENT 1: INITIAL INVESTIGATION OF AVERAGE RESOURCE ALLOCATION

Since no prior studies, to our knowledge, have looked at visualization framing in the context of charitable resource allocations, we designed this first experiment to investigate which donation strategies people follow when presented with text and visualizations with different framings. We also set out to verify whether effects caused by cognitive biases would manifest within our study design. More importantly, we explored the effect of presentation format and information framing on resource allocation.

In our pre-registrations, we distinguish between our primary (or main) research questions, and secondary (or auxiliary) research questions. Each primary research question further elaborates on one of our core research questions outlined in the introduction, and is indicated in parenthesis.

- (RQ_{vis}) **RQ1.1** To what extent does a visual representation of information affect donation allocations compared to a design with the same information presented as text only?
- (RQ_{frm}) **RQ1.2** To what extent does framing the donation amount per individual affect donation allocations when compared to framing the donation amount per program?
- (RQ_{int}) **RQ1.3** To what extent do the factors of information representation and information framing interact with one another regarding donation allocations?
- RQ1.4** What are the most common stated reasons for the donation allocations?

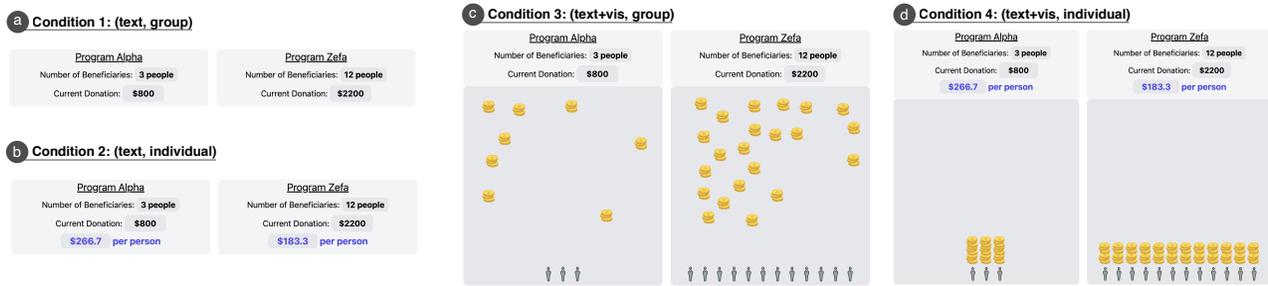


Figure 2: Experiment 1: Stimulus used in the different conditions. See a demo of the tool at <https://resallocdesign.github.io/>

5.1 Experimental Design

Experiment 1 followed the procedure described in section 4. The Alpha and Zefa programs aided 3 and 12 children respectively. The amount to allocate was USD\$3,000.

Our between-subject design manipulated two independent variables. The PRESENTATION-FORMAT determined whether the same information was either conveyed using text only (*text*), or using both text and visualization (*text+vis*). The FRAMING variable determined whether the presentation formats (§3.3) present information with either a focus on splitting resources between causes (i.e. *group* framing), or a focus on splitting resources between individuals in the causes (i.e. *individual* framing). Our full factorial design encompassed four conditions. The content presented in the display component of the tool (see Figure 1-b) for each condition is depicted in Figure 2.

5.2 Participants

Our target sample size was 256 participants (~64 per condition; participants were randomly assigned a condition at sign up). This gave us a 0.8 power to detect a “medium” Cohen *d*’s effect size of 0.5 between any two conditions (per the G*Power software) for differences between independent means. We ran this study on August 7, 2021 until target sample size was reached (within hours), after 2 pilots with 5 and 10 participants respectively. Out of the 256 participants who consented, 17 were excluded for not completing the study or failing the attention check. The 239 remaining responses were divided as follows: 63 were assigned the (*text+vis, individual*) condition, 63 completed the (*text, individual*) condition, 56 completed the (*text+vis, group*) condition, and 60 completed the (*text, group*) condition.

5.3 Measurements & Analyses

We followed methods described in §4.3 & §4.4. We used the allocation to the larger group (Zefa) to measure allocation strategy: a value of \$2,400 corresponds to a fair allocation strategy, where all individuals are given the same amount. With this measure, we used a difference in means to generate the bootstrapped confidence intervals used to answer our research questions.

For the qualitative coding of justification comments (method not preregistered), we followed an emergent coding approach [63]: four coders independently assigned labels to participant responses. One coder did so across all responses while the three other coders each covered a separate one-third of the responses. Amount allocated to

each program was hidden during this process. All coders discussed respective codes, refining, combining, and merging codes where appropriate, converging on a common consolidated coding scheme comprising of 5 categories, each of which comprising several codes (see OSF repository). One coder then re-coded all shuffled responses using the new scheme.

Our approach to data collection and methods for analysis were refined through unregistered piloting of our study. This helped us refine our wording and layout of the story, adding an attentional check, and refining analysis scripts.

5.4 Results

Figure 3 shows a summary of results. The stacked bar chart (a) shows a break down of responses per allocation strategy (see §4.4), and the error bars (b) show 95% CIs of mean allocation to Zefa, compared to a fair allocation (\$2,400).

Overall, most participants followed a fair allocation strategy. Figure 3-b shows mean allocations to Zefa nearing or covering the \$2,400 value. Figure 3-a shows that the condition with fewer participants choosing a fair allocation (in light blue) was the one with textual information and group framing (40% of participants), while showing text+visualization with individual framing led to about 80% fair allocations. Also, a notable amount of participants chose a mixed allocation strategy in *group* framing conditions (30% for *text*, 34% for *text+vis*), which shifted the mean allocations of those conditions farther from the fair allocation. Below, we discuss our results organized by research question.

(RQ_{vis}) **RQ1.1: Effect of Adding a Visualization Support.** We were interested in learning the extent to which different presentation formats affect resource allocation. Figure 4-a shows the effect of supplementing text with a visualization on our metric, by framing. We could not find evidence that the addition of a visualization affects the fairness of the mean allocation compared to a design with the same information presented as text only.

(RQ_{frm}) **RQ1.2: Effect of Framing.** We also explored the extent to which different framings affect resource allocation. Figure 4-b shows the effect of varying framing on our metric, for each presentation format. Overall, we have some evidence suggesting that an individual framing leads to a fairer mean resource allocation than a group framing, for the condition where both textual and visual information are available.

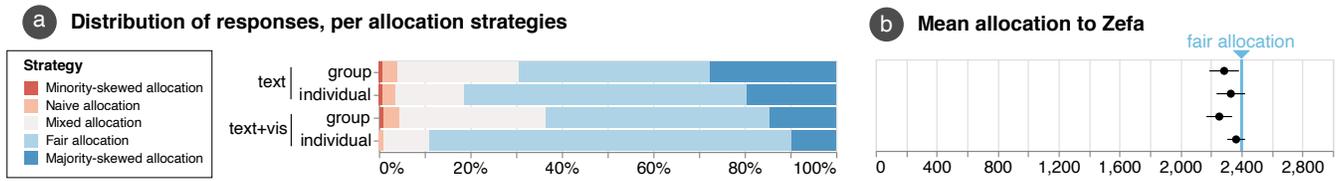


Figure 3: Results of Experiment 1: (a) Distribution of allocation strategy chosen by participants, by PRESENTATION FORMAT and FRAMING. (b) Mean allocation granted to Zefa by condition. Error bars are 95% CIs.

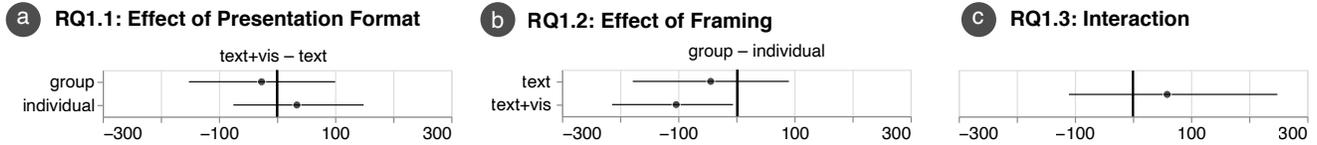


Figure 4: Results of Experiment 1: Effect of PRESENTATION FORMAT (a) and FRAMING (b) on the mean allocation (here in \$), and interaction between our different factors (c). Error bars are 95% CIs.

(RQ_{int}) **RQ1.3: Interaction.** We investigated the interaction between the two design factors (i.e. presentation format and information framing). Figure 4-c shows a quantification of this interaction. We could not find evidence that our factors interact with respect to mean resource allocation.

RQ1.4: Stated Rationale. We also looked for the most common stated reasons participants provide for their donation allocations. From coding the justifications (see §5.3), our main insights are:

(i) Inconsistencies between the strategy used and the justification provided. The *sense of equality* code determined if participants responses indicated that they tried to allocate equitably across charitable programs. We coded 70 responses as such (out of 239), however 12 of those participants had not used a fair strategy.

(ii) Mental calculations (e.g. dividing a total amount by number of individuals) are not necessarily performed to achieve a fair allocation, and are prone to errors. The *math* code category indicated whether the response included some explicit calculation. Of the 44 participants who did include such mention, eight had not used the fair allocation strategy.

(iii) Some participants made assumptions about the programs, which have likely impacted their response. The *contextual influence* code category determined whether the response indicated that the participant’s allocation was chosen based on personal interpretations of the scenario, e.g. presuming that program Zefa solves a more significant problem, whereas our scenario did not state so. This finding motivated us to reword the scenarios in subsequent experiments, to leave less room to such extrapolations.

(iv) The two other coding categories did not provide additional insight, and are thus omitted in interest of space.

5.5 Discussion

Contrary to our expectations, presentation format does not seem to greatly affect resource allocations (RQ_{vis}). We could not find evidence that supplementing text with a visualization produces fairer mean resource allocation compared to only presenting text. This is

inconsistent with our initial assumption that showing a visual representation would help participants to evenly distribute resources across beneficiaries. Note that we did not find that visualization negatively affects allocation fairness either.

Our results also suggest that the way the information is framed could affect how people allocate resources (RQ_{frm}). We specifically found evidence that individual framing influenced participants to perform fairer allocations compared to a group framing, in presence of a visualization. We suspect that presenting the exact amount that each individual would receive influenced participants to choose a value that would help everybody equally. This finding reinforces previous evidence that information framing can play a role on decision making [19, 35].

We later realized that the weak evidence for the effect of framing could be attributed to the aggregated metric that we used. Mean allocation and mean differences are sensitive to extreme values; and the majority-skewed responses could have pushed the mean closer to 2400, even in conditions where a smaller proportion of participants made fair allocations. We decided to re-analyze our data using the proportion of fair allocations – the ratio of fair allocations to the number of responses in each condition. This new metric better aligns with our research aim of exploring fair allocations. We did find evidence that individual framing causes more participants to use a fair allocation strategy (which can also be seen in the descriptive statistics of Figure 6-a). Since this unplanned analysis was not preregistered, we re-ran a new experiment (Experiment 2) to confirm our tentative findings with this alternative metric.

6 EXPERIMENT 2: A CLOSER LOOK AT FAIR VS. UNFAIR ALLOCATIONS

This second experiment is a replication of Experiment 1 with a more robust and meaningful metric, and improvements to the scenario brief based on limitations observed in Experiment 1. Our research questions were:



Figure 5: Results of Experiment 2: Effect of PRESENTATION FORMAT (a) and FRAMING (b) on overall proportion of fair allocations made. Error bars are 95% CIs.

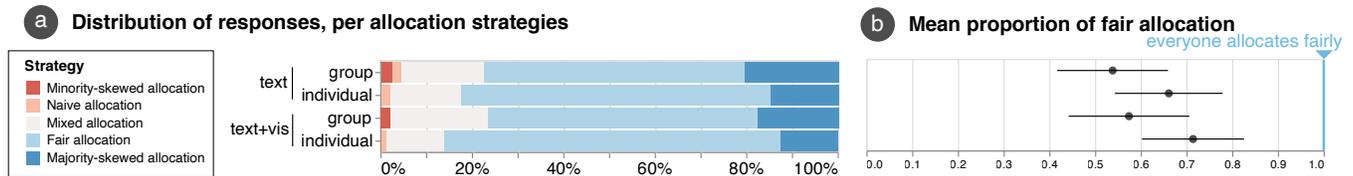


Figure 6: Results of Experiment 2: (a) Distribution of allocation strategy chosen by participants, by PRESENTATION FORMAT and FRAMING (see §3.3). (b) Mean proportion of people making a fair allocation by condition. Error bars are 95% CIs.

- (RQ_{vis}) **RQ2.1** To what extent does a visual representation of information influence the proportion of fair allocations compared to a design with the same information presented as text only?
- (RQ_{frm}) **RQ2.2** To what extent does framing the donation amount per individual influence the proportion of fair allocations when compared to framing the donation amount per program?

6.1 Experimental Design

We followed the same procedure and design as Experiment 1. The only difference was in the scenario brief. To prevent participants from speculating on the programs' effectiveness, we added: (i) a description of needs (i.e. "Children from both programs are in a very similar situation and need money equally. They will not benefit from any other donation than yours."), and (ii) a clarification of how the money will be allocated within the programs (i.e. "All the money you donate to a program will be split evenly between the children in the program."). See a demo at <https://resallocdesign.github.io/>. With these clarifications about the underlying assumptions, there is less ambiguity about which allocation is the most fair.

6.2 Participants

Like for Experiment 1, we targeted a sample size of 256 (~64 per condition). We ran this study on March 1, 2022 until target sample size was reached. Twelve (12) participants were excluded. The 244 remaining responses were divided as follows: 62 in the (*text+vis, individual*) condition, 63 in the (*text, individual*) condition, 54 in the (*text+vis, group*) condition, and 65 in the (*text, group*) condition.

6.3 Measurements & Analyses

Unlike in Experiment 1, here we used the relative proportion of people who chose to use a fair allocation strategy as a means to assess overall allocation fairness. While we report descriptive statistics showing the breakdown of responses per allocation strategies using all five strategies from §4.3, our inferential statistics only contrast

fair vs. unfair responses. With this measure, we used a difference in means to generate the bootstrapped confidence intervals.

We validated our analysis methodology through unregistered experimentation: we explored the data from Experiment 1 to arrive at our measure used for analysis in Experiment 2. All analyses in Experiment 2 were pre-registered.

6.4 Results

Figure 6 shows a summary of our results. The stacked chart (a) shows responses per allocation strategy, by condition. The plot (b) shows the proportion of fair allocations.

Both plots suggest that most participants performed a fair allocation. We observe that individual-framing conditions tend to result in slightly higher proportions of fair allocations compared to group-framing ones. Actual evidence for these possible effects is further discussed under our research questions below.

(RQ_{vis}) **RQ2.1: Effect of Adding a Visualization Support.** We were interested in the effect of presentation format on the proportion of fair allocations. Figure 5-a shows the difference in proportions between the *text* format and the *text+vis* format for each of the two framings. The effect of format seems to be negligible, independently of framing.

(RQ_{frm}) **RQ2.2: Effect of Framing.** We were also interested in the effect of information framing. Figure 5-b shows the difference in proportion of fair allocations between *individual* and *group* framing, for each presentation format. There is evidence that individual framing may motivate fairer allocations than group framing: taken individually each of the two CIs provides very weak evidence, but taken together they provide converging evidence.

6.5 Discussion

The findings from this experiment are in line with the results from Experiment 1. This means for our core question (RQ_{vis}), results suggest that supplementing text with a visualization produces a comparable proportion of fair allocations as using stand-alone text. For

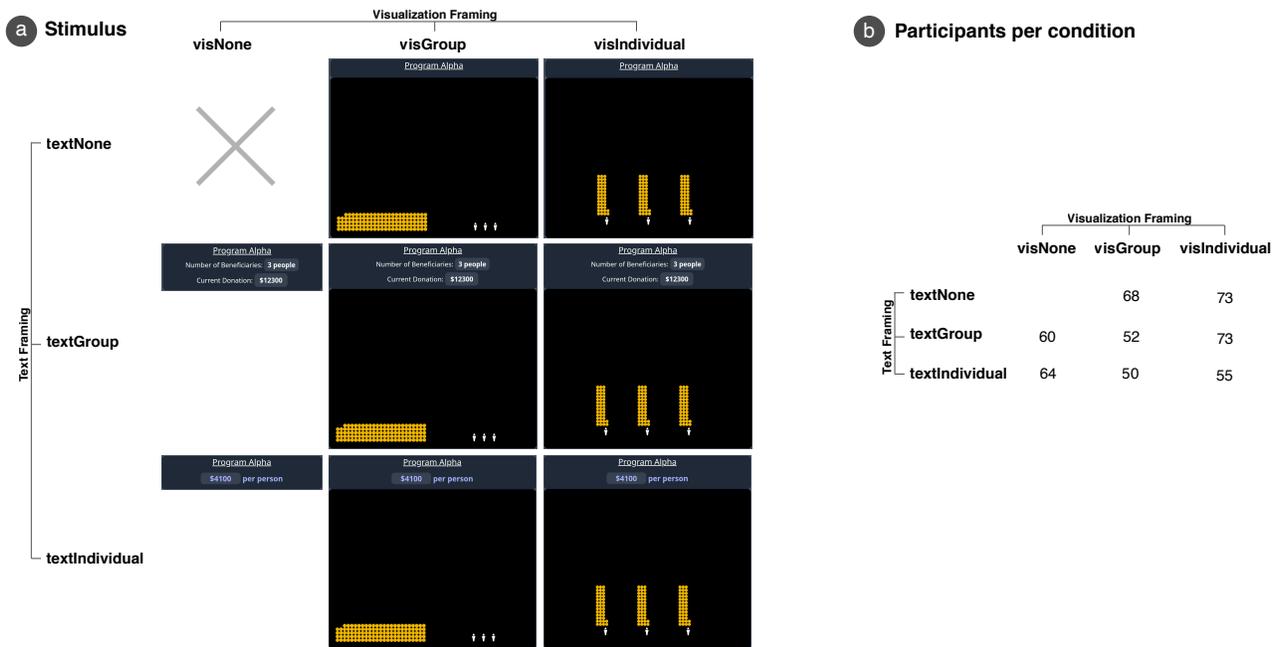


Figure 7: Experiment 3: (a) Stimulus used in the different conditions, (b) Participant count per condition.

our core question (RQ_{frm}), we found some evidence that individual-focused framing is more effective for motivating fair allocations than group framing (with weak evidence for each of the two presentation formats). Those results strengthen the assumption that information framing plays a role on resource allocation.

The high variation of fair allocations across conditions in the previous experiment was not reproduced in this experiment (compare Figure 6-a and Figure 3-a). As mentioned, maybe some instructions in Experiment 1 led participants to believe that a specific program was more important than the other, which caused a greater variance in responses. The instruction improvements in this experiment may have reduced room for personal assumptions and made the fair allocation strategy more unambiguously fair, explaining the lower variation. Findings from prior studies on cognitive biases §2.1 do not seem to apply significantly in our context (deciding how to split resources between groups with a 1:4 population ratio), as the vast majority of participants allocated resources fairly.

Experiments 1 & 2 allowed to examine the effect of framing on two presentation designs, one with text only, and one where text is supplemented with a visualization, with a congruent framing. In a follow-up experiment, we examine the role of each presentation component in isolation, as well as effects of congruent vs. incongruent framings.

7 EXPERIMENT 3: A THOROUGH LOOK AT INFORMATION FRAMING

The main purpose of the previous experiments was to investigate whether *supplementing* text with a visualization influences fair allocations. In this experiment, we take a closer look at the role of

information framing, exploring each presentation format independently (i.e text only, visualization only, and text plus visualization). This allows us to isolate the role of presenting information through visualization only. Prior studies found little influence of visualizations in inducing prosocial behaviors (§2); in Experiment 3, we examine whether visualizations can substitute text, while promoting a similar propensity for fair allocations in our context. This also allows us to examine conditions where text and visualization have incongruent framings (e.g. text focuses on individuals and visualization focuses on group, and vice-versa, see §3.3). Experiment 1 & 2 only included congruent framings (with additional caveat that text presentation included both group and individual framings in the individual-focused framing design).

Our main research question focused on the effect of visualization framing as follows²:

(RQ_{int}) **RQ3.1** To what extent does visualization framing influence the unfairness of fund allocations across all text framing conditions?

Secondary research questions aimed at uncovering the effect of visualization framing were:

RQ3.2 To what extent does visualization framing influence the fairness of allocations when (a) no text information is present; (b) text information is shown with a group framing; and when (c) text information is shown with an individual framing?

RQ3.3 Does the effect of visualization framing depend on text framing?

²Note that our pre-registered material referred to ‘text’ presentation format as ‘numerical’, and ‘individual’ framing as ‘person’ framing.

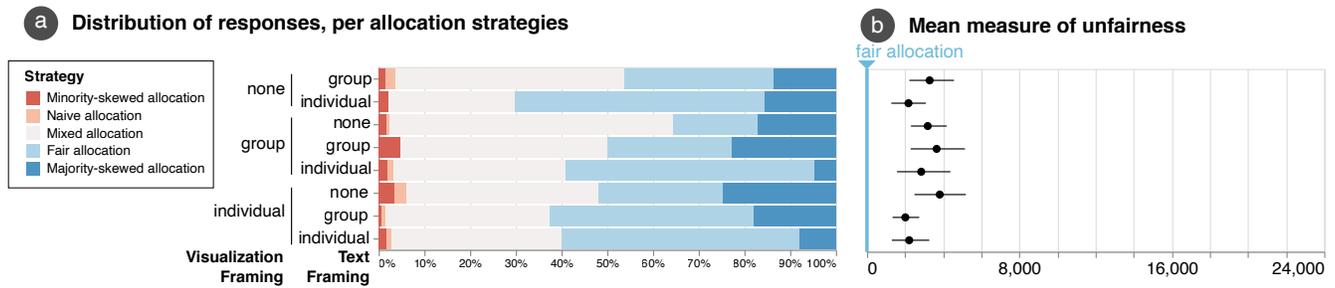


Figure 8: Results of Experiment 3: (a) Distribution of allocation strategy chosen by participants, by TEXT FRAMING and VISUALIZATION FRAMING, and (b) Mean measure of unfairness per condition. Error bars are 95% CIs.

An additional secondary research question focused on examining the effect of presenting a visualization:

- RQ3.4** To what extent does the presence of visualization influence the fairness of allocations when (a) the text information and visualization have a congruent group framing; (b) the text information and visualization have a congruent individual framing; and when (c) the text information and visualization have an incongruent framing?

Finally, our two auxiliary research questions were:

- RQ3.5** What are the most common stated reasons for the donation allocations?
- RQ3.6** How does each condition affect the duration taken to complete an allocation choice?

7.1 Experimental Design

We followed the procedure described in §4.2. Program Alpha still had 3 children, but we changed the size of program Zefa to now include 120 children, resulting in a 1:40 ratio (instead of 1:4 in the other Experiments). Since a larger difference between the programs might be more prone to compassion fade, this will allow us to further examine the robustness of our previous findings. We also changed the amount of the resource to be allocated accordingly. Keeping a target of \$200 per individual for a fair allocation, the total amount to be allocated was changed from \$3,000 to \$24,600.

Our visual design was refined through unregistered piloting. We switched to a dark background to achieve higher contrast, and traded the coins pictogram for a simpler yellow circle for increased simplicity and readability (Figure 7).

We ran a between-subject experiment where we manipulated the presence and framing of both a text and a visualization presentation format as independent variables. The TEXT-FRAMING variable specifies the framing for the text presentation format: *textGroup* uses a group-focused framing; *textIndividual* uses an individual-focused framing; and *textNone* means that text presentation is not included. The VISUALIZATION-FRAMING variable specifies framing of the visual representation of information, with *visGroup* and *visIndividual* values corresponding to a group and an individual framing respectively, and *visNone* corresponds to no visualization at all.

The full factorial design results in nine conditions. However, we discard the (*textNone*, *visNone*) condition as it results in an empty display (i.e. no information is presented at all). The stimulus included in the display component of the interactive tools (Figure 1-b) for all conditions respectively are depicted in Figure 7-a.

7.2 Participants

Our target sample size was 512 participants (~64 per condition). After a successful pilot with 6 participants which informed visual design refinements, we ran this study on June 26, 2022 until target sample size was reached.

Seventeen (17) participants were excluded. The 495 valid responses were distributed across conditions as per Figure 7-b.

7.3 Measurement & Analysis

We performed the measurements and analyses as described in §4.3 & §4.4. To gauge allocation (un)fairness, we define the *unfairness* metric as the distance to the fair allocation, in dollars. A value of zero means a fair allocation has been made; the further from zero, the unfairer the allocation. For instance, an unfairness value of \$600 means that either \$24,600 or \$23,400 has been allocated to Zefa, both of which are a distance of \$600 away from a fair allocation.

We note that our unfairness metric is identical to the Gini index – a standard measure of income inequality in populations [34] – up to a multiplicative constant. We use this distance metric as we wanted a metric which preserves our exact allocation value for both programs as specific breakpoints might derive additional meaning in the context of allocation strategies. With this measure, we used statistical contrasts to generate the bootstrapped confidence intervals used to answer our research questions.

7.4 Results

Figure 8 shows a summary of our results. The stacked bar chart (a) shows responses per allocation strategy by condition. The plot (b) shows the mean unfairness measurement per condition. From Figure 8-a, we note that fair allocations were not a majority in all conditions, suggesting potential effects in this more granular experimental design, which we investigate through our research questions as follows.

(RQ_{int}) **RQ3.1: Effect of Visualization Framing.** As our primary research question, we were interested in learning the extent to which visualization framing (*visGroup* vs. *visIndividual*) influences the unfairness of allocations across all text-framing conditions. Figure 9-a shows the difference in mean allocation unfairness corresponding to the statistical contrast *visGroup* – *visIndividual*, with *visGroup* (resp. *visIndividual*) corresponding to the union of all text framing conditions where visualization framing is group-focused (resp. individual-focused).

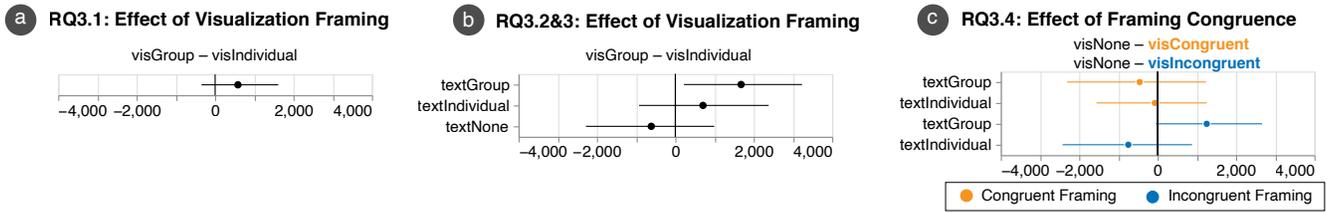


Figure 9: Results of Experiment 3: Mean difference between individual and group framing, overall (a), by text framing (b). Mean difference between present/absent visualization, by text framing and congruence of framings (c). Error bars are 95% CIs.

This represents the influence of visualization framing across all text framing conditions. We found very weak evidence that group framing for visualizations might cause more unfairness in allocations when compared to an individual framing for visualizations.

RQ3.2: Effect of Presenting Visualization Framing. We were interested in learning the extent to which visualization framing influences allocation fairness across all values in TEXT-FRAMING. Figure 9-b shows the difference in unfairness between *visGroup* and *visIndividual* when presented with different text framings.

To find the effect of visualization framing when text is absent (RQ3.2a), we estimated the contrast $(textNone, visGroup) - (textNone, visIndividual)$. To find the effect of visualization framing when text is shown with an *individualText* framing (RQ3.2c), we estimated the contrast $(textIndividual, visGroup) - (textIndividual, visIndividual)$. Since the CIs for both contrasts clearly cross zero, we find no evidence that group framing causes more or less unfairness in allocations when compared to individual framing for visualizations when no text, or when individual text framing is used.

To find the effect of visualization framing when text is shown with a *groupText* framing (RQ3.2b), we computed the contrast $(textGroup, visGroup) - (textGroup, visIndividual)$. Since the CI is to the right of zero, we find some evidence that group framing for visualizations causes more unfair allocations than individual framing for visualizations, when a group text framing is present.

RQ3.3: Interaction Between Text Framing and Visualization Framing. We wanted to know the extent to which the effect of visualization framing depends on text framing. To do this, we usually assess the overlap between the CIs in Figure 9-b. The overlap between the CIs for *textNone* and *textIndividual*; and for *textGroup* and *textIndividual*, is too large to conclude in an interaction [46]. The overlap between the CIs for *textGroup* and *textNone* is relatively small, so there is some weak evidence that visualization framing has a larger effect (i.e., group visualization framing increases unfairness more as compared to individual visualization framing) when group text framing is used than when there is no text.

RQ3.4: Effect of Presenting Congruent or Incongruent Visualizations. We were interested in learning the extent to which the presence of visualization influences allocation fairness when the framing between text and visualization was congruent and incongruent. Figure 9-c shows the means difference in unfairness between the conditions where visualization framing is congruent and incongruent with text framing. In this chart, positive values can be interpreted as “the presence of visualization results in more fair allocations than when there is no visualization”, in other words, adding a visualization helps; and negative values can be interpreted

as “the absence of visualization results in more fair allocations than when there is a visualization”, i.e., adding a visualization harms.

To find the effect of showing a visualization when the text and visualization have a congruent *group* framing (RQ3.4a) we contrast $(textGroup, visNone) - (textGroup, visGroup)$ (top row in Figure 9-c). To find the effect of showing a visualization when the text and visualization have a congruent *individual* framing (RQ3.4b) we contrast the conditions $(textIndividual, visNone) - (textIndividual, visIndividual)$ (Figure 9-c, second row). As both CIs for the difference of means clearly cross 0, we cannot claim the influence of congruent group or individual framings between visualization and text.

To find the effect of showing a visualization when text information and visualization have an incongruent framing (RQ3.4c) we perform two contrasts. When *textGroup* has an incongruent framing with the visualization, we contrast the conditions $(textGroup, visNone) - (textGroup, visIndividual)$ (Figure 9-c, third row). We find that the difference in means is largely within the positive range. Even though the CI closely crosses zero, we can claim that when text presents allocations per group, adding a visualization presenting these per individual helps achieve fairer allocations. When *textIndividual* has an incongruent framing with the visualization, we contrast $(textIndividual, visNone) - (textIndividual, visGroup)$ (last row in Figure 9-c). As the CI crosses zero, we find no evidence that individual framing for visualizations causes more or less unfair allocations than no visualizations, when text conveys information with a group framing.

RQ3.5 & RQ3.6: Auxiliary research questions. We were interested in the most common stated reasons for the donation allocations (RQ3.5). Since the number of participants using a mixed-allocation strategy differentiated significantly from past experiment, we decided on a further investigation of responses for this category only. We performed an exploratory coding of the responses with 3 codes belonging to one category. The *zefa cause* code indicated whether a participant based their justification around giving more to the Zefa cause than to the Alpha cause; the *alpha individual* code indicated the participant based their justification around giving more to the individuals in Alpha than to the individuals in Zefa; and *balanced* code indicated the participant said they gave similar to both programs. We found that most participants still justified giving more to Zefa due to the higher group size.

To find how each condition affects task duration (RQ3.6) we plot the CIs for the mean allocation time for all conditions (see OSF repository). Participants took between 35-85 seconds to reach their decision, and those in the $(textNone, visIndividual)$ and $(textIndividual, visNone)$ conditions reached a decision the fastest.

7.5 Discussion

We took a closer look at the effect of framing. Overall, our findings suggest that individual framing, either for visualization or text, tends to yield slightly fairer allocations when compared to group framing. This result reinforces the assumption that framing information, in a way that helps donors to identify the amount allocated to individuals, motivates them to split the resources equally among all the beneficiaries. We discuss specific effects in more detail.

We found that visualization framing can play a role on resource allocations in specific situations (RQ_{int}). We found that including an individual-framed visualization is able to curb unfairness compared to group-framed visualizations, when the visual representation supplements text conveying allocations to groups. We suspect that the visualizations that displayed allocations per individual helped participants to identify the amount that each beneficiary would receive, making them to allocate the money more fairly than when presented with a group-framed text only. However, it is important to note that showing visualizations, either individual- or group-framed, without an accompanying text may have prompted participants to make more unfair allocations. Therefore, it seems that the presence of text influences participants to opt for a fair allocation strategy.

We also found that providing stand-alone text with an individual framing may produce similar fair allocations than including individual-framed visualizations combined with any text framing (RQ_{int}). However, using group-framed text without including a visualization influences participants to have more unfair allocations. This finding is important since not only including a visualization with individual framing did not curb the effect of text framing (i.e. visualization did not harm), it also helped in cases where text was presenting allocations with a group framing. Therefore, incorporating certain visualization framings may bring a net benefit in terms of promoting fairness, and at worse, not influence the donor. This confirms that choices are prone to framing effects [35], more specifically, that particular visualization framings have an influence on how people make decisions (§2.3). This finding adds contrast to prior work within our domain which has limited findings for visualizations which try to influence prosocial behaviours [54].

8 DISCUSSION AND FUTURE WORK

We discuss our findings, limitations, and research implications.

8.1 Summary of Findings

This work explored the extent to which data visualization and information framing can promote fair resource allocation in humanitarian decision-making. We found evidence that allocations are more fair overall when resource allocation information is framed around individuals rather than groups (RQ_{frm}). Our findings on the effects of adding visualizations are less conclusive (RQ_{vis}), but suggest that visualizations may be more effective at promoting fairness when they are individual-framed and accompanied by a textual representation (RQ_{int}). Those findings contribute to advance the understanding of framing in the context of information visualization and shed light on possible research directions.

Our experiments are a first attempt to understand the interplay between visualizations and information framing in the context of

an interactive decision-support tool for humanitarian resource allocation. We identified five different strategies people use to allocate resources (see §4.3). Overall, we found that most participants chose *fair* or *mixed* allocations. Mixed allocations give slightly less weight to individuals in the larger group. Thus, the presence of a large number of mixed allocations suggests the presence of cognitive biases that may influence participants to favor the smaller population group (possibly exacerbated by a larger population ratio between the groups in Experiment 3). However, the larger number of fair allocations suggests that in our particular experiment setting, *compassion fade* is relatively mild. Other cognitive biases such as the *drop-in-the-bucket* effect and the *diversification bias* do not seem to have played a major role in participants' decisions, since strategies favoring individuals from the smaller group and *naive* allocations giving equally to both programs were relatively uncommon.

We find it interesting that, while individual framing achieved more fair allocation overall, the absence of strong bias which we could expect when emphasizing entativity somehow challenges our premise. Cognitive biases may have had a more nuanced role in influencing allocation, however, this and prior work has not compared different sensitivities which participants may have for each bias. For example, some participants may have only been affected by *drop-in-the-bucket* while others may have been only affected by *compassion fade*. This work does not aim to prescribe specifically which biases visualizations are able to curb. As such, we recommend further investigation of the specific biases discussed in §2 in regards to resource allocation visualizations and framing.

Despite the relatively large proportion of fair allocations, unfair allocations were nonetheless common, and thus there is clearly room for improvement. It appears that the prevalence of fair allocations can be affected to some extent by how allocation information is presented. Our first and second experiment did not find conclusive evidence that adding a visualization to text using a congruent framing makes a clear difference. However, we found converging evidence that individual-framing causes a higher proportion of fair allocations, independently of the presentation format used. Thus it seems that individual-framing should be the choice if the goal is to promote fair allocations, while adding a visualization at least does not seem to harm. Thus, designers can choose to add a visualization to provide other benefits. We discuss those implications further in section 8.3.

In the last experiment, we investigated the effect of information framing in more detail, separating the effects by presentation format. In line with findings from our previous experiments, our results suggest that individual framing plays a clear role in promoting fair allocations. Also consistent with our previous experiments, we found that adding an individual-framed visualization to an individual-framed text (i.e., *congruent individual framing*) does not seem to promote fairer allocations compared to individual-framed text alone. Similarly, we did not find evidence that adding a visualization with *congruent group framing* yields less fair allocations. Thus, it does not seem that adding a visualization to a text reinforces the effect of information framing, if the framing is the same which aligns with past work finding little benefit to visualization in eliciting prosocial behavior [49, 53]. However, we did find evidence that adding an individual-framed visualization to a group-framed text yields to more fair allocations overall.

8.2 Study Limitations

Due to the complex and multi-faceted nature of humanitarian resource allocation problems, there were a number of trade-offs that we had to make to arrive to a tractable study.

First, our simplified task design is far from perfectly capturing real-world situations. In particular, our allocation scenario involves only two charity programs, where all targeted individuals have the same needs and whose life will be impacted equally if they get the same amount of financial aid. Those assumptions are unlikely to hold in virtually any real-world humanitarian resource allocation problem, if only because there are vast inequalities among individuals and because different humanitarian programs are generally far from being equally impactful [27]. In addition, real-world humanitarian resource allocation problems involve complex interactions between multiple organizations and humanitarian programs, whereas our scenario assumed that the individuals from the two groups would not receive other forms of help. Similarly, personal charitable donations interact, and how fair a specific donation is depends on what other people have donated and to whom. It is important to note that our tool is only a stylized tool whose goal is to understand how people make decisions in controlled settings, and that a real decision-making tool should take into account the many additional dimensions of humanitarian resource allocation.

Our work suggests that interactive visualization tools may hinder some cognitive biases in the context of humanitarian decision-making, but visualizations can also mislead and provoke additional biases [64]. The stages of the visualization pipeline can induce perceptual or cognitive misconceptions [52]: curating flawed data, making wrong data transformations, using deceptive visual encodings, misreading, etc. Our studies investigated how visualization design and information framing can contribute to curb cognitive biases related to decision-making, but other biases that affect judgement on perceptual or social levels were not directly addressed [12]. More work is needed to investigate the impact of visualization tools in charitable decision-making and biases ingrained in this process.

It is important to note that we used a simplified notion of fairness, as fairness is a complex concept. We chose to scope and narrow the notion of fairness as providing equal financial aid to all individuals to make our study tractable; and we developed metrics around that definition. Our proposed metric designed for Experiment 3 is equivalent to the standard Gini index; and was therefore a reasonable choice. However, more elaborate metrics that account for the many factors that real resource allocation problems encompass and for the diversity of normative ethical frameworks will be needed to study allocation fairness more thoroughly. Future work should use more holistic measures that capture the multi-dimensionality of allocation fairness, and generalize to an arbitrary number of groups.

8.3 Perspectives for Visualization

This paper brings together theories of framing, cognitive biases, and visualization to explore the effect of design choices for interactive humanitarian resource allocation tools. It is still premature to prescribe visualization design principles or guidelines for reducing unfairness in resource allocations. However, our findings advance our understanding of the problem, potential solutions, opportunities, and challenges in this area.

To start with, our findings add nuance to previous studies which found little influence of visualizations in inducing behaviors that benefit others [54]. We found that visualizations do not harm in our context, and can help in some circumstances. This is an important result as beyond decision making tasks, visualizations have a wealth of potentially beneficial properties – they draw attention [38], they are easier and faster to consume and interpret, and they can even promote self-reflection [18], making them an important asset when communicating about resource allocations.

Although our work only presents preliminary findings about how to design visualizations for resource allocation tools, we have evidence that some design choices may facilitate good resource allocation decisions. Broadly, visualizations should be designed to emphasize help or impact on each individual instead of (or in addition to) on a global scale. It is however important to note that the visualization designs we tested are only one of the many possible designs that could be imagined. In particular, future visualization designs could be imagined that convey richer information. For instance, the pictorial marks representing group populations in anthropographics [53] could be made visually richer so as to capture the diversity of population samples instead of impersonal standard iconography. Other resources such as time or the actual impact on individuals could be represented instead of coins.

When designing visualizations for future resource allocation tools, one challenging question will be what additional information to include, and how. For instance, would representing geographical locations, genders, races, nature and severity of the problem, be important to convey? Or could it instead be counter-productive to add certain information, due to human biases? The design space for mitigating cognitive bias in visualization [68] helps to critically reflect on such questions, and on possible design strategies to consider, such as temporarily hiding certain information as not all data is always helpful. We see the exploration of these questions as an exciting avenue for future research.

8.4 Perspectives for HCI

Our resource allocation user interface has been designed and deployed as an experimental tool to investigate the role of different design dimensions on resource allocation decision-making. As such, it integrates only very basic functions. In general, our study can be seen as an initial exploration of the vast design space of user interfaces for such tools. There are many opportunities that we see our work could be extended to further support decision-makers.

First, we did not push interactive design beyond implementing a slider to dynamically control the allocation. The HCI literature can provide inspiration to extend our approach to foster particular behaviors, such as by inserting visual nudges [51], adding stickiness [68], or repeatedly nudging the donor [19, 55]. Further, creative, engaging interaction paradigms could activate curiosity and playfulness in tools dealing with otherwise serious problems. There might be potential for such an interface to intrigue and incite people to learn more about causes, their differences, their needs, and how they, as individuals, can help in an effective manner. A systematic review of the tools that are currently employed in large

associations and governments managing multiple charitable programs would also allow to better understand current practices, and challenges that decision makers face with the current tools.

Second, resource allocation tools in the future could benefit from some level of automation. For instance, scraping content online to populate the interface with real-world, timely data, or to plug them into databases such as [usaspending.gov](https://www.usaspending.gov). Future tools could also encompass computational models, for instance, by incorporating predictions which could enable donors to appreciate how their donation can make concrete change.

We also see an opportunity to elevate such resource allocation tools from tool to partners [36], to achieve collaborative intelligence between the decision-maker and the computational tool [70]. A first step would be to look at recent work in improving automation and heuristics-based assistance for the allocation of resources for humanitarian aid [1, 2, 13]. Further, the utility of resource allocation tools is not simply for donors themselves but could also be used for visualizing an automated amount given out by some system, such as one integrating a machine learning model. Such tools provide ways to alleviate potential harms in algorithmically-infused societies [67]. Again, we are excited by the many possibilities to extend our work.

9 CONCLUSION

This work investigated how different aspects of tool design, namely presentation format and information framing, can nudge people into allocating humanitarian resources fairly. Our results indicate that there may be benefits in using text and visualizations with certain framings for such resource allocation tools. We found that, in general, individual-framed resource allocation information (i.e., how much each individual will receive) seems to be more effective in promoting fair allocations, whether information is conveyed through text or through a visualization. We also found that adding an individual-framed visualization to a group-framed text (i.e., how much each group will receive) promotes fair allocations. However, showing individual-framed text appears to be sufficient, and fairness does not seem to be improved by adding individual-framed visualizations. While most of our results are tentative, they are a first step towards understanding the interplay between information visualization and framing for humanitarian resource allocation.

For the sake of experimental control and to make our study tractable, we made a number of simplifying assumptions. We used a fictional resource allocation scenario with two charity programs that are equivalent except for the number of individuals they they target, and we assumed that all individuals have equal needs and will not receive help from other organizations. Future work should explore resource allocation problems that are more realistic, for example by involving charity programs that differ across several dimensions and populations that are more diverse. In addition, we focused on the resource allocation problem for two groups. In real-world examples, charities oftentimes need to allocate resources to more groups. Further studies should also generalize our methods to three or more groups. Finally, none of the designs we tested was able to eliminate occurrences of mixed allocations, and therefore it seems likely that some degree of compassion fade persists across designs. While our study looked at a very specific scenario, we hope it will serve as a starting point for a wider range of research work.

ACKNOWLEDGMENTS

We thank colleagues of the Dynamic Graphics Project (DGP) lab for their input and feedback for our study and manuscript. This work was supported in part by a grant from NSERC (RGPIN-2018-05072) and University of Toronto FA&S Tri-Council Bridge Funding.

REFERENCES

- [1] Emily Aiken, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock. 2022. Machine learning and phone data can improve targeting of humanitarian aid. *Nature* 603, 7903 (2022), 864–870.
- [2] Azrah Anparasan and Miguel Lejeune. 2019. Resource deployment and donation allocation for epidemic outbreaks. *Annals of Operations Research* 283, 1 (2019), 9–32.
- [3] Jonathan Baron and Ewa Szymanska. 2011. *Heuristics and biases in charity*. na.
- [4] Daniel M Bartels. 2006. Proportion dominance: The generality and variability of favoring relative savings over absolute savings. *Organizational Behavior and Human Decision Processes* 100, 1 (2006), 76–95.
- [5] Daniel M Bartels and Russell C Burnett. 2011. A group construal account of drop-in-the-bucket thinking in policy preference and moral judgment. *Journal of Experimental Social Psychology* 47, 1 (2011), 50–57.
- [6] Shlomo Benartzi and Richard H Thaler. 2001. Naïve Strategies in Defined Contribution Saving Plans. *American Economic Review* 91 (2001), 79–98.
- [7] Ralf Bender and Stefan Lange. 2001. Adjusting for multiple testing—when and how? *Journal of clinical epidemiology* 54, 4 (2001), 343–349.
- [8] Sudeep Bhatia, Lukasz Walasek, Paul Slovic, and Howard Kunreuther. 2021. The more who die, the less we care: evidence from natural language analysis of online news articles and social media posts. *Risk Analysis* 41, 1 (2021), 179–203.
- [9] Jeremy Boy, Anshul Vikram Pandey, John Emerson, Margaret Satterthwaite, Oded Nov, and Enrico Bertini. 2017. Showing people behind data: Does anthropomorphizing visualizations elicit more empathy for human rights data?. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 5462–5474.
- [10] John Burn-Murdoch and Pilita Clark. 2017. Climate Change Calculator. <https://ig.ft.com/sites/climate-change-calculator/>.
- [11] Marcus M Butts, Devin C Lunt, Traci L Freling, and Allison S Gabriel. 2019. Helping one or helping many? A theoretical integration and meta-analytic review of the compassion fade literature. *Organizational Behavior and Human Decision Processes* 151 (2019), 16–33.
- [12] André Calero Valdez, Martina Ziefle, and Michael Sedlmair. 2018. Studying biases in visualization research: Framework and methods. In *Cognitive biases in visualizations*. Springer, 13–27.
- [13] Frederick Callaway, Bas van Opheusden, Sayan Gul, Priyam Das, Paul M Krueger, Thomas L Griffiths, and Falk Lieder. 2022. Rational use of cognitive resources in human planning. *Nature Human Behaviour* 6, 8 (2022), 1112–1125.
- [14] Sarah Campbell and Dietmar Offenhuber. 2019. Feeling numbers: The emotional impact of proximity techniques in visualization. *Information Design Journal* 25, 1 (2019), 71–86.
- [15] Lucius Caviola, Stefan Schubert, and Joshua D Greene. 2021. The psychology of (in) effective altruism. *Trends in Cognitive Sciences* 25, 7 (2021), 596–607.
- [16] Chun-Tuan Chang and Yu-Kang Lee. 2009. Framing charity advertising: Influences of message framing, image valence, and temporal framing on a charitable appeal 1. *Journal of applied social psychology* 39, 12 (2009), 2910–2935.
- [17] Fanny Chevalier, Romain Vuillemot, and Guia Gali. 2013. Using concrete scales: A practical framework for effective visual depiction of complex measures. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2426–2435.
- [18] Michelene TH Chi, Miriam Bassok, Matthew W Lewis, Peter Reimann, and Robert Glaser. 1989. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive science* 13, 2 (1989), 145–182.
- [19] Andy Cockburn, Blaine Lewis, Philip Quinn, and Carl Gutwin. 2020. Framing effects influence interface feature decisions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [20] Tina Comes. 2016. Cognitive biases in humanitarian sensemaking and decision-making lessons from field research. In *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 56–62.
- [21] Tina Comes, Niek Wijngaards, and Bartel Van de Walle. 2015. Exploring the future: Runtime scenario selection for complex and time-bound decisions. *Technological Forecasting and Social Change* 97 (2015), 29–46.
- [22] Geoff Cumming. 2013. *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge.
- [23] Geoff Cumming and Sue Finch. 2005. Inference by eye: confidence intervals and how to read pictures of data. *American psychologist* 60, 2 (2005), 170.
- [24] Kahneman Daniel. 2017. Thinking, fast and slow.
- [25] Evanthis Dimara, Gilles Bailly, Anastasia Bezerianos, and Steven Franconeri. 2018. Mitigating the attraction effect with visualizations. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 850–860.

- [26] Pierre Dragicevic. 2016. Fair statistical communication in HCI. In *Modern statistical methods for HCI*. Springer, 291–330.
- [27] Pierre Dragicevic. 2022. Information Visualization for Effective Altruism. *arXiv preprint arXiv:2209.00836* (2022).
- [28] Pierre Dragicevic. 2022. Towards Immersive Humanitarian Visualizations. *arXiv preprint arXiv:2204.01313* (2022).
- [29] Robert J Fisher. 1993. Social desirability bias and the validity of indirect questioning. *Journal of consumer research* 20, 2 (1993), 303–315.
- [30] Steven L Franconeri, Lace M Padilla, Priti Shah, Jeffrey M Zacks, and Jessica Hullman. 2021. The science of visual data communication: What works. *Psychological Science in the public interest* 22, 3 (2021), 110–161.
- [31] James Friedrich and Acacia McGuire. 2010. Individual differences in reasoning style as a moderator of the identifiable victim effect. *Social Influence* 5, 3 (2010), 182–201.
- [32] Iason Gabriel. 2017. Effective altruism and its critics. *Journal of Applied Philosophy* 34, 4 (2017), 457–473.
- [33] Alexander Garinther, Holly Arrow, and Pooya Razavi. 2022. Victim number effects in charitable giving: Joint evaluations promote egalitarian decisions. *Personality and Social Psychology Bulletin* 48, 1 (2022), 34–48.
- [34] Corrado Gini. 1997. Concentration and dependency ratios. *Rivista di politica economica* 87 (1997), 769–792.
- [35] Corentin J Gosling, Serge Caparas, and Sylvain Moutier. 2020. The interplay between the importance of a decision and emotion in decision-making. *Cognition and Emotion* 34, 6 (2020), 1260–1270.
- [36] Jonathan Grudin. 2017. From tool to partner: The evolution of human-computer interaction. *Synthesis Lectures on Human-Centered Interaction* 10, 1 (2017), i–183.
- [37] Neeraja Gupta, Luca Rigotti, and Alistair Wilson. 2021. The Experimenters' Dilemma: Inferential Preferences over Populations. *arXiv preprint arXiv:2107.05064* (2021).
- [38] Fred Hohman, Matthew Conlen, Jeffrey Heer, and Duen Horng Polo Chau. 2020. Communicating with interactive articles. *Distill* 5, 9 (2020), e28.
- [39] Jessica Hullman and Nick Diakopoulos. 2011. Visualization rhetoric: Framing effects in narrative visualization. *IEEE transactions on visualization and computer graphics* 17, 12 (2011), 2231–2240.
- [40] Karen Jenni and George Loewenstein. 1997. Explaining the identifiable victim effect. *Journal of Risk and uncertainty* 14, 3 (1997), 235–257.
- [41] Daniel Kahneman and Amos Tversky. 1981. *The simulation heuristic*. Technical Report. Stanford Univ CA Dept of Psychology.
- [42] Tehila Kogut and Ilana Ritov. 2005. The “identified victim” effect: An identified group, or just a single individual? *Journal of Behavioral Decision Making* 18, 3 (2005), 157–167.
- [43] Wolfgang Köhler. 1967. Gestalt psychology. *Psychologische Forschung* 31, 1 (1967), XVIII–XXX.
- [44] Morgane Koval and Yvonne Jansen. 2022. Do You See What You Mean? Using Predictive Visualizations to Reduce Optimism in Duration Estimates. In *CHI Conference on Human Factors in Computing Systems*. 1–19.
- [45] Donald R Kretz. 2018. Experimentally evaluating bias-reducing visual analytics techniques in intelligence analysis. In *Cognitive Biases in Visualizations*. Springer, 111–135.
- [46] Martin Krzywinski and Naomi Altman. 2013. Error bars: the meaning of error bars is often misinterpreted, as is the statistical significance of their overlap. *Nature methods* 10, 10 (2013), 921–923.
- [47] Louisa Küchler, Guido Hertel, and Meinald T Thielsch. 2020. Are you willing to donate? relationship between perceived website design, trust and donation decisions online. In *Proceedings of the Conference on Mensch und Computer*. 223–227.
- [48] Guowei Li, Monica Taljaard, Edwin R Van den Heuvel, Mitchell AH Levine, Deborah J Cook, George A Wells, Philip J Devereaux, and Lehana Thabane. 2017. An introduction to multiplicity issues in clinical trials: the what, why, when and how. *International journal of epidemiology* (2017).
- [49] Johannes Liem, Charles Perin, and Jo Wood. 2020. Structure and empathy in visual data storytelling: Evaluating their influence on attitude. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 277–289.
- [50] Scott R Maier. 2015. Compassion fatigue and the elusive quest for journalistic impact: A content and reader-metrics analysis assessing audience response. *Journalism & Mass Communication Quarterly* 92, 3 (2015), 700–722.
- [51] Justin Matejka, Michael Glueck, Tovi Grossman, and George Fitzmaurice. 2016. The effect of visual appearance on the performance of continuous sliders and visual analogue scales. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 5421–5432.
- [52] Andrew McNutt, Gordon Kindlmann, and Michael Correll. 2020. Surfacing visualization mirages. In *Proceedings of the 2020 CHI Conference on human factors in computing systems*. 1–16.
- [53] Luiz Morais, Yvonne Jansen, Nazareno Andrade, and Pierre Dragicevic. 2020. Showing data about people: A design space of anthropographics. *IEEE Transactions on Visualization and Computer Graphics* (2020).
- [54] Luiz Morais, Yvonne Jansen, Nazareno Andrade, and Pierre Dragicevic. 2021. Can Anthropographics Promote Prosociality? A Review and Large-Sample Study. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [55] Nuno Mota, Abhijnan Chakraborty, Asia J Biega, Krishna P Gummadi, and Hoda Heidari. 2020. On the desiderata for online altruism: Nudging for equitable donations. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–21.
- [56] Marie Neurath. 1974. Isotype. *Instructional science* (1974), 127–150.
- [57] Clair Null. 2011. Warm glow, information, and inefficient charitable giving. *Journal of Public Economics* 95, 5–6 (2011), 455–465.
- [58] Daniel Read and George Loewenstein. 1995. Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied* 1, 1 (1995), 34.
- [59] Jacob Ritchie, Daniel Wigdor, and Fanny Chevalier. 2019. A lie reveals the truth: Quasimodes for task-aligned data presentation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [60] Laura Marie Schons, John Cadogan, and Roupini Tsakona. 2017. Should charity begin at home? An empirical investigation of consumers' responses to companies' varying geographic allocations of donation budgets. *Journal of Business Ethics* 144, 3 (2017), 559–576.
- [61] Paul Slovic, Melissa L Finucane, Ellen Peters, and Donald G MacGregor. 2007. The affect heuristic. *European journal of operational research* 177, 3 (2007), 1333–1352.
- [62] Robert W Smith, David Faro, and Katherine A Burson. 2013. More for the many: The influence of entitativity on charitable giving. *Journal of Consumer Research* 39, 5 (2013), 961–976.
- [63] Steve Stemler. 2000. An overview of content analysis. *Practical assessment, research, and evaluation* 7, 1 (2000), 17.
- [64] Danielle Albers Szafrir. 2018. The good, the bad, and the biased: Five ways visualizations can mislead (and how to fix them). *Interactions* 25, 4 (2018), 26–33.
- [65] Nhat Tran, Malihe Alikhani, and Diane Litman. 2022. How to Ask for Donations? Learning User-Specific Persuasive Dialogue Policies through Online Interactions. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*. 12–22.
- [66] Daniel Västfjäll, Paul Slovic, Marcus Mayorga, and Ellen Peters. 2014. Compassion fade: Affect and charity are greatest for a single child in need. *PLoS one* 9, 6 (2014), e100115.
- [67] Claudia Wagner, Markus Strohmaier, Alexandra Olteanu, Emre Kıcıman, Noshir Contractor, and Tina Eliassi-Rad. 2021. Measuring algorithmically infused societies. *Nature* 595, 7866 (2021), 197–204.
- [68] Emily Wall, John Stasko, and Alex Endert. 2019. Toward a design space for mitigating cognitive bias in vis. In *2019 IEEE Visualization Conference (VIS)*. IEEE, 111–115.
- [69] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 56–65.
- [70] H James Wilson and Paul R Daugherty. 2018. Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review* 96, 4 (2018), 114–123.