On the Desiderata for Online Altruism: Nudging for Equitable Donations

NUNO MOTA, Max Planck Institute for Software Systems, Germany
ABHIJNAN CHAKRABORTY, Max Planck Institute for Software Systems, Germany
ASIA J. BIEGA, Microsoft Research, Canada
KRISHNA P. GUMMADI, Max Planck Institute for Software Systems, Germany
HODA HEIDARI, Carnegie Mellon University, USA

Online donation platforms help equalize access to opportunity and funding in cases where inequalities exist. In the context of public school education in the United States, for instance, financial inequalities have been shown to be reflected in the educational system, since schools are primarily funded through local property taxes. In response, private charitable donation platforms such as `DonorsChoose.org` have emerged seeking to alleviate systemic inequalities. Yet, the question remains of how effective these platforms are in redressing existing funding inequalities across school districts. Our analysis of donation data from `DonorsChoose` shows that such platforms may in fact be ineffective in mitigating existing inequalities or may even exacerbate them.

In this paper, we explore how online educational charities could direct more funding towards more impoverished schools without compromising their donors' freedom of choice with respect to donation targets. Seeking to answer this question, we draw on the line of work on choice architectures in behavioral economics and pose a novel research question on the impact of interface design on equity in socio-technical systems. Through controlled experiments, we demonstrate how simple interface design interventions—such as modifying default rankings or displaying additional information about schools—might lead to changes in donation distributions helping platforms direct more funding towards schools in need. Going beyond online educational charities, we hope that our work will bring attention to the role of interface design nudges in the social requirements of online altruism.

CCS Concepts: • **Human-centered computing → Empirical studies in interaction design**;

Additional Key Words and Phrases: School Funding; Fair Donation; Choice Architecture; Digital Nudge

ACM Reference Format:
Nuno Mota, Abhijnan Chakraborty, Asia J. Biega, Krishna P. Gummadi, and Hoda Heidari. 2020. On the Desiderata for Online Altruism: Nudging for Equitable Donations. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 126 (October 2020), 21 pages. https://doi.org/10.1145/3415197

1 INTRODUCTION

Public schools in the United States offer tuition-free access to primary and secondary education to students from all financial backgrounds. In most states, the public education system is divided into local school districts that are largely funded from local revenue sources (such as local property taxes) and by the state governments, with supplemental funding coming from the federal government [19].

126

Authors' addresses: Nuno Mota, Max Planck Institute for Software Systems, Germany; Abhijnan Chakraborty, Max Planck Institute for Software Systems, Germany; Asia J. Biega, Microsoft Research, Canada; Krishna P. Gummadi, Max Planck Institute for Software Systems, Germany; Hoda Heidari, Carnegie Mellon University, USA.

2573-0142/2020/10-ART126 $15.00
https://doi.org/10.1145/3415197

Although public funds are the primary source of school district revenue, multiple studies and media reports have criticized existing school district boundaries as promoting racial and financial segregation [28, 64, 69]. These reports have brought in a renewed attention to the severe deficiencies of the existing school funding policies and, in particular, their role in exacerbating existing racial disparities [1, 88]. As demonstrated through various protests and strikes [48], teachers across the country see adequate school funding as a necessary condition for providing quality education to their students. Since education is one of the key factors determining an individual's socio-economic outcomes, such as employment prospects and wage-earning capacity, existing inequities in access to high-quality education can be viewed as infringing upon fundamental human rights.[1] Proponents of *Equality of Resources* for education have argued strongly that the playing field must be leveled [11, 41, 78]. Advocates of *Equality of Capability* [65] may go a step further and advocate for more funds *in favor* of disadvantaged students—for instance, to equalize the *capability* of getting into a reputable college.[2]

Heavy reliance on local property taxes to fund public schools can result in schools located in impoverished neighborhoods receiving significantly lower levels of funding from local sources. Such funding disparities can play a major role in the long-term educational and economic outcomes of students. By distributing funds more equally across districts, policymakers can equalize access to quality education – which in turn can end the vicious cycle of poverty for many students and improve social mobility [34, 42]. Although one option may be to redraw the district boundaries to amalgamate wealthy and poor neighborhoods into one school district [22], such proposals require community participation and legal support. Thus, one of the few practical alternatives is for the state or federal governments to fill in the financial gaps created from local revenues, so that all districts are equally well-funded. Our analysis in Section 3.1 shows, however, that the state and federal-level attempts to reduce funding inequalities have fallen short of leveling the playing field.

What else can be done to minimize these inequalities? Private charitable donations are a popular approach to remedying persisting social injustices. According to Giving USA annual reports, US citizens and organizations donated a total of $410 billion to charities in 2017, and donations for educational causes in particular amounted to approximately $58.90 billion [85]. Several online platforms, such as `DonorsChoose.org`, `AdoptAClassroom.org` or `DigitalWish.com`, have been launched to help donors to support educational needs. DonorsChoose alone, as of March 2019, has received over $800 million in donations, through which more than 1.3 million school projects have been funded, positively impacting 33.4 million students in public schools across the USA [24]. Yet, the question remains of how effective these platforms are in redressing existing funding inequalities across school districts. Our analysis in Section 3.2 shows that donations made through these platforms do not effectively counterbalance existing inequities in public school funding, and often perpetuate those same inequities.

While many notions of equity are conceivable in the context of educational donations—such as increasing amounts of donations as school poverty level increases, funding all projects from the poorest schools before funding projects of richer schools, having the same amount of donations per student across poverty levels, as well as notions inspired by the individual fairness principle [26] where similar projects and schools get similar amount of donations—our focus in this paper

[1] Article 26 of the Universal Declaration of Human Rights, United Nations, 1948 [54] states that "Everyone has the right to education. Education shall be free, at least in the elementary and fundamental stages. [...] Technical and professional education shall be made generally available and higher education shall be equally accessible to all on the basis of merit."

[2] Note that compared to their advantaged peers, students who come from a disadvantaged background may not be equally able to convert school resources/funds into valuable opportunities, due to the different social and environmental circumstances under which they live. According to the capability approach, this observation can justify allocating more resources to disadvantaged students.

is on *need*-based principles of allocating resources. More specifically, we explore several ways of *directing more donations to the poorest schools*—which often fail to receive an equitable share of public funding. We impose a crucial constraint on possible interventions, however. While increasing funding equity, an intervention should respect the donors' *freedom of choice* in how they donate their private money. The key question we attempt to address in this paper is thus—*whether online donation platforms like DonorsChoose can contribute to the leveling of school funding disparities without limiting donors' freedom of choice?*

Drawing on a rich body of work in behavioral economics, we argue that the platform can indeed play an important role in shaping donors' behavior. Donation websites present requests within a *choice architecture* [81], influencing donors' decision-making and *nudging* them towards choosing certain alternatives. Examples of nudges [80] on DonorsChoose include the *default* ranking in displaying projects, *decision-making aids* such as search filters, or the *selection of project attributes* disclosed to the donors. The main conceptual contribution of our work is to draw attention to the power of such nudges in remedying school funding disparities. We propose several design nudges that can help direct more donations to schools in more need without compromising donors' freedom of choice. We empirically demonstrate the effectiveness of those nudges through controlled human-subject experiments.

We face two major challenges in measuring the impact of the interface design changes on donors' behavior. First, we have to ensure that only a single design element is changed while all the other elements of the platform environment remain fixed, including the set of available projects. Second, while A/B testing by the platform is feasible, such an approach raises ethical concerns in a setting where the outcomes influence the real-world financial resources of teachers and schools in need [7]. To circumvent these two issues, we propose an experimental methodology where paid crowdsourcing workers are asked to make *virtual* donations on a proxy website mirroring the real platform. Such an approach allows us to quantify the influence of individual design elements on project selection rates by changing the design in the proxy server, all the while avoiding manipulation of real donations. More specifically, our controlled experiments use a proxy server mimicking the DonorsChoose website, and seek to investigate whether the project selection rates can be changed through website design interventions (i.e., *choice architecture*) along the dimension of poverty level. Importantly, we show that it is possible to significantly increase the rate of donations to highest-poverty schools by displaying the poverty information in project listings, or changing the default ranking of projects to poverty-based (details in Section 4).

In summary, our work makes the following three contributions: we (i) highlight the inequality in student funding in the US public school system, (ii) draw attention to the inadequate counter-balancing effect of online charitable giving, and (iii) empirically demonstrate the role of choice architectures and interface design nudges in encouraging more equitable donations. Our results highlight the fact that donation platforms have an active role to play in shaping donors' choices and aligning them with the social values. More generally, this paper aims at drawing attention to the role of digital nudges and choice architectures in socio-technical systems. We hope that the paper paves the way for future studies aiming to employ nudges effectively on other platforms that seek to contribute to our collective values, such as fairness and equality of opportunity.

2 RELATED WORK

2.1 Analyzing crowdfunding platforms

Prior works have looked into various aspects of how platform characteristics influence the amount of funding different projects receive. Wash and Solomon [86] have shown that returning money to donors if a project does not reach its financial goal by a fixed deadline, leads to a larger total amount of donations since people are not afraid to donate to riskier projects. According to Meer [49], increasing the platform fees reduces the likelihood of projects getting funded, and matching donations for a given project does not crowd out other similar projects [50]. With respect to the content, different features of projects, donors, and fundraisers have been found to influence long-term retention of donors [2] and the project's financial success [30, 32, 82]. The likelihood of funding has also been shown to be influenced by the language used in project descriptions [51], the gender of teachers [60], the social media activity associated with the project [44], and the amount of early donations [73]. With respect to the fundraiser behavior, it has been found that they tend to learn more efficiently from successful rather than unsuccessful projects [56], or that they need to put a lot of work into establishing the legitimacy of the project and themselves [79].

To understand donations in the context of DonorsChoose, it is important to consider the implications of a crowdfunding paradigm. In this paper, we focus on the inequality in donations going to different schools, and tie it to the poverty levels of the underlying school districts. By contrasting the donations with the funding different schools get from the government sources, we can check whether online charities are successful in minimizing the offline funding inequalities, and the potential role of the platform in the donation allocation.

2.2 Nudges and choice architectures

Attempting to offer a realistic model of human decision-making, Simon [70] introduced the concept of *bounded rationality*, which is related to "the amount of information individuals can consciously keep track of". Due to this cognitive limitation, people rely on many *intuitions* and *heuristics* to simplify the decision-making process [5, 36], which makes them susceptible to cognitive biases. For example, people tend to follow what others are doing (a phenomenon known as the *herd effect* [45]), and are reluctant to change existing setups (known as the *status-quo bias* [66]). Considering this mental state as suboptimal, Thaler and Sunstein [80] promoted the use of *choice architectures* [35, 71, 81] (i.e., the way choices are presented), and *nudging* [67, 80] (i.e., subtle interventions that produce predictable deviations in behavior) as tools to augment human decision-making.

A long line of cross-disciplinary research has since shown the potential of nudging and choice architectures in multiple applications. For example, the positioning of different food items in a snack bar can influence the food habits (and calorie intake) of the consumers [38], the design architecture of a retirement savings plan can have a significant impact on savings behavior [4], and time-limited discounts on fertilizer greatly increases sales in under-privileged populations [25]. In the context of digital advertising [91], it has been shown that click-through rates depend on a whole group of ads displayed in a block rather than individual ads. Harbach et al. [33] showed that redesigning permission prompts within the Google Play Store facilitated privacy-conscious choices. Brewer and Jones [10] developed a tool to ease seniors' engagement with social media (by providing family-oriented incentives), andMcGee-Lennon et al. [47] demonstrated the potential of auditory queues in increasing long-term memory retention and response times to digital notifications. Along this line of work, our main contribution revolves around understanding how the design of educational charity websites affects donor behavior, and *investigating the potential of nudging in fulfilling social desiderata of online altruism.*

2.3 How people donate

There are two main lines of work studying charitable giving. In the first line, charitable activity is considered to be a rational, utility maximizing process. For example, under a rational choice framework, it has been shown that tax reductions significantly increase monetary contributions to charity [62], and that charitable giving can be modeled through utilitarian objectives (e.g., desire for social influence) [63]. The second line of work considers charitable giving to be an irrational, non-utilitarian behavior. Researchers have identified "personal taste" as the most important criterion for donations [9], and have empirically demonstrated donors' systematic deviations from a theoretical utility-based decision making process [3]—a phenomenon potentially attributed to cognitive biases. Furthermore, studies in the first line of work have a greater focus on the motivations for whether or not and how much to donate, whereas the second line of work shifts attention towards how donors choose between competing options. As such, our work closely aligns with the later.

In a similar vein, other studies into charitable giving have investigated the effects of matching donations [37], whether people tend to support the neediest donation targets [20], whether the display of organization quality ratings increases the likelihood of donation [12], whether people prefer to donate to a specific target or a general charitable fund [27], whether deadlines increase the likelihood of donations [21], or whether approaching donors directly influences the likelihood and amount of donations [29]. In this paper, we combine these aspects to see how changes in interface design can affect the donors donating to schools in higher need. More importantly, the paper brings to the fore the power of interface design to fulfill social objectives of online platforms (such as fairness), without drastically interfering with the ways in which people donate.

3 TACKLING EDUCATIONAL INEQUALITY: THE ROLE OF CHARITABLE GIVING

Public school districts in the USA are funded by local, state and federal governments, and each public school belongs to a single district. This assignment is based on school district boundaries, but recent studies and media reports have raised concerns about the way they are determined (a process termed as 'Educational Gerrymandering' [69]), where the concerns range from racial segregation in public schools [64] to drawing district borders to benefit small, wealthy communities [28]. Given existing district boundaries, our focus in this section is to understand the extent of inequalities in public school funding and whether educational charities can effectively reduce them.

3.1 Inequality in School Funding

3.1.1 Dataset.

From the National Center for Education Statistics (NCES: `https://nces.ed.gov/ccd`) we gathered the breakdown of local, state and federal funding for every public school district in the USA (comprising a total of 19, 639 districts), along with the number of students enrolled in them.[3] We also collected the number of students eligible for free and reduced-fee lunch, which we use to calculate the poverty level of a school district.[4] After removing districts with missing entries or zero enrollments, we are left with information on 16, 720 school districts, which we utilize throughout this paper. To enable meaningful funding comparisons across school districts, we adjust the revenues with the 'Comparable Wage Index for Teachers (CWIFT)' at the district-level [18, 55], making the dollar amounts comparable. This adjustment makes it so that a school district getting more revenue in a costly neighborhood can be properly contrasted with another district getting same revenue in a relatively cheaper area.

In our data, we observe a high variation in the number of students enrolled in different school districts, with poorer districts enrolling a lot more students than their wealthier counterparts. Wealthy districts have a lower fraction of students from low-income families, and current boundaries seem to segregate students depending on their family's wealth. On the one hand, a student from a low-income family is more likely to attend schools with many other students from poor

[3]The data is for 2014-15 school year, the latest available as on October 2019.

[4]To be eligible for free and reduced-fee lunch, a student's family income needs to be within 130% and 185% of the poverty line threshold in different years, e.g., less than $31, 005 and $44, 123 annual income for a family of four in 2014-15 [84].

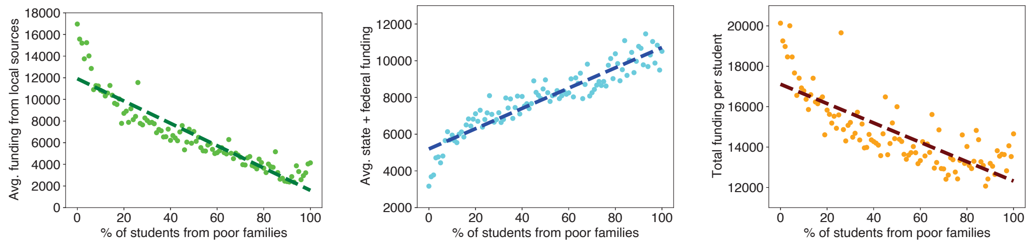

(a) Local funding (b) State and federal funding (c) Total funding

Fig. 1. **Average funding per student for school districts with different poverty levels: (a) funding from local sources is much higher for wealthier districts compared to poorer districts. (b) While funding from federal and state governments seem to reverse this trend, with more funding going to poorer districts, the magnitude of funding is not sufficient to counter-balance the local taxes raised for public school education. (c) As a result, per-student funding (from all sources combined) is much higher for wealthier school districts compared to poorer school districts.**

families. On the other hand, a student from a wealthy family is likely to go to a school with fewer students, most of them from better-off families. To account for the variation in student numbers, we normalize the revenue (funding) received by different districts by the number of students in each district, and consider 'per student revenue' as a comparative measurement.

3.1.2 Distribution of per-student revenue across school districts.

As mentioned earlier, a public school district's funds come from three sources (local, state and federal), with some variation across the states in the relative source contributions. For example, in 23 states, more than half of districts' funds came from state governments, whereas in another 15 states, local revenues constituted more than half of total funds. For the remaining states, neither state nor local revenue sources contributed more than half of total funds. Across the states, on average, about 8% of the total revenues came from federal sources.

Regardless of the breakdown of funding sources, about 80% of local revenue comes from local property taxes [46]. Since the collection of property tax is tied with real-estate value, wealthier neighborhoods collect higher amounts compared to poorer neighborhoods. Such reliance on property taxes creates a huge disparity in the amount of funding going to school districts across poverty levels. To clarify this claim, Fig. 1a shows a much higher per student funding for wealthier districts, coming from local sources. When looking at state and federal sources, Fig. 1b shows that this trend is somewhat combated, but the overall contribution of federal sources is low. Plus, there is some inconsistency in the distribution of state funding across states. For example, in Arizona, Georgia, Idaho and Indiana, roughly a similar amount of per student funding goes to wealthiest and poorest school districts, whereas in Louisiana and Montana, a higher per-student funding goes to wealthiest districts. We observe the combined effect in Fig. 1c, where the total per-student funds is much higher for wealthier districts compared to poorer districts.

Such disparity in per-student funding goes against the ideal of *Equality of Resources* for education [11, 41, 78], and can play a major role in students' long-term educational and economic opportunities. By distributing funds more equally across districts, policymakers can equalize access to quality education, which in turn can end a vicious cycle of poverty for many students. However, such policy changes would require prolonged consultations with different stakeholders, and thus have uncertain implementation timelines. As an alternative to this process, private charitable donations have emerged as effective tools to remedy persistent social injustices. Next, we investigate whether online educational charities can help in reducing existing disparities in school districts' funding.

3.2 Role of Educational Charity

In recent years, online donation platforms like MightyCause, Chuffed, GoFundMe, CrowdFunder or GlobalGiving have emerged as popular mediums for connecting donors to those in need of financial resources. In such platforms, requesters submit projects specifying their needs, and donors can contribute financially to help raise the requested amount. Several donation platforms like DonorsChoose, AdoptAClassroom or DigitalWish focus exclusively on educational needs. Ideally, such platforms should contribute to reducing existing inequalities in educational resources. In this subsection, we investigate whether they are effective in mitigating the discrepancies in public school funding.

3.2.1 Dataset.

We particularly focus on `DonorsChoose.org`, a popular donation platform where teachers from public schools across the USA can request donations for their classroom projects. Donors can select projects based on several attributes, such as school location (zip code, city, state), main subject of the project (e.g., Math and Science, Health and Sports, Literacy and Language), or requested resources (e.g., Books, Computers and Tablets, Food, Clothing and Hygiene). However, if a project does not raise its target amount by a certain deadline, it receives no funds, and the donors need to reallocate the donated money to other projects.

To identify the distribution of donations going towards schools in different districts, we analyze the dataset publicly released by DonorsChoose in 2016 [23], which contains information about all projects posted on the platform from 2002 to 2016. During the time period covered by the dataset, DonorsChoose helped raise \$500 million for over 1 million projects posted from $70K+$ schools. Despite the substantial amount of money donated through the platform, not every project gets successfully funded. The data reveals that about 30% of submitted projects fail every year. Moreover, different projects require varying donation amounts and the corresponding schools belong to different poverty levels. Therefore, it is important to understand how schools at different poverty levels perform in terms of attracting donations through DonorsChoose, and whether the donations can reduce funding inequality across public school districts.

3.2.2 Donations going to different districts.

The projects in DonorsChoose are posted from schools in 1489 school districts, which is about 9% of all school districts in the NCES dataset (described in Section 3.1.1). Hence, we group the districts based on the percentage of students from poor households in 10 bins with 10% ranges, i.e., {[0,10), [10-20), ..., [90-100]}. Although the absolute number of requests from poorer districts is much higher than wealthier districts, clearly reflecting the fact that they have a higher need for additional funding, poorer districts also admit lot more students than wealthier districts, leading to a lower amount of requested donations per student (as shown in Fig. 2a). Surprisingly, despite poorer schools having greater need and the level of poverty being available in DonorsChoose (in the project description page), we see in Fig. 2b that a project's success probability to be roughly constant within $69-76\%$ across poverty levels, with a bit lower success rate for poorer districts. As a result, the amount of donations received per student is lower for poorer school districts (Fig. 2c).

One could argue that poorer schools ask for less, and simply receive less as a result. However, the outcome may not be so trivial, as in DonorsChoose, projects asking for higher donations tend to fail with higher probability [14]. In fact, DonorsChoose itself suggests that teachers should break their requirements into smaller projects to increase their chance of success [52]. Thus, asking for more donations may be counter-productive for the poorer schools. We further observe that districts with $10-50\%$ students from low income households seem to perform exceptionally well (only exception being districts with $20-30\%$ of poor students); whereas wealthiest districts (i.e.,

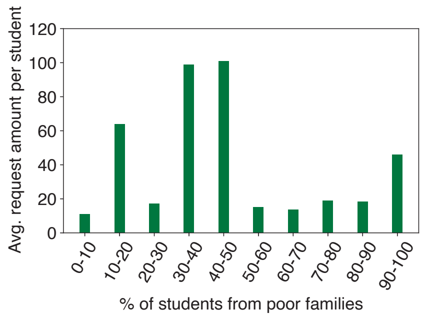

(a) $ requested per student

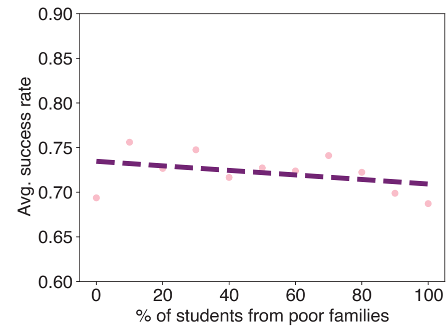

(b) Project success rate

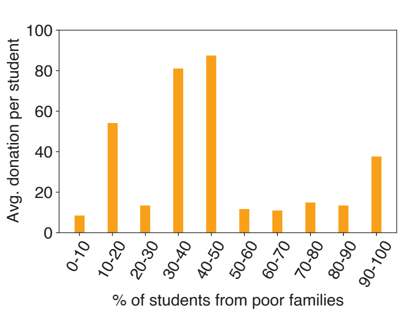

(c) $ received per student

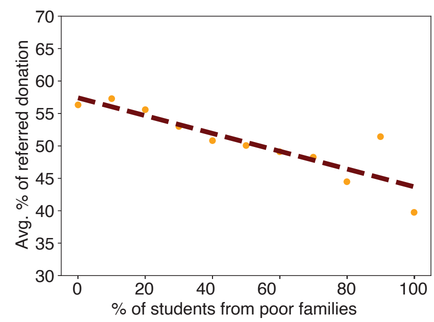

(d) % of referred donations

Fig. 2. **Despite schools from poorer districts having higher need, (a) they request lower donation per student, and (b) project success rate is similar across poverty levels. As a result, (c) the average donation received per student is lower for poorer districts. (d) Teachers can advertise referral links asking people to donate to their projects. Due to the difference in social engagements, the fraction of donations for a project coming from referrals is much higher for wealthier districts.**

with at most 10% students from low income households) are getting less donations per student than poorest districts (i.e., with at least 90% students from low income households). It seems that in DonorsChoose there might exist a tension between *need* (i.e., poverty level) and *know-how* (usually correlated with lower poverty and an important success predictor in crowdfunding platforms [56]). Schools within certain poverty ranges likely strike the right balance between these two components: (i) just enough *need* to motivate effort through the platform; and (ii) just enough *know-how* to produce attractive proposals. We plan to further explore this hypothesis in future work.

3.2.3 Role of teacher activity.

DonorsChoose encourages teachers to publicize their projects by posting them on social media, asking their acquaintances to donate. Such advertisements carry referral links (or teacher specific codes) which enable DonorsChoose to track whether a donation was referred by the corresponding teacher. We analyze these referrals across different projects posted from different districts. Fig. 2d shows that the fraction of donations for a project coming from teacher referrals is much higher for wealthier districts, giving them a much-needed headstart. This scenario is similar to the parent/Parent-Teacher-Organization (PTO) contributions to schools, which are another significant source of disparity in school funding levels, as it is often the wealthiest districts that also get the largest PTO contributions [13, 87].

This further supports our hypothesis of the importance of the know-how—in this case, the ability of social engagement to try and increase funding through referrals. Teachers from wealthier districts may be more active on social media platforms and/or they may be able to reach enthusiastic parents to donate to their respective projects, which teachers from poorer districts may not have access to.

In this section, we observed that platforms like DonorsChoose do not seem to counterbalance existing inequities in public school funding through the way their donors donate – the donations often end up mimicking those same inequities. Moreover, charitable giving constitutes a limited resource when compared to the systemic inequities it attempts to mitigate. Thus, there is a need for the platform to play an active role towards achieving equitable donations. In the next section, we investigate whether by changing the platform design in particular ways, educational charities like DonorsChoose can contribute towards leveling school funding disparities.

4 NUDGING TOWARD EQUITABLE DONATIONS

A long line of cross-disciplinary research has looked into the ways people make decisions. As currently understood, there exists a limit on the amount of information one can rationally keep track of, a phenomenon known as *bounded rationality* [70]. Because of this, people often rely on a combination of processes, both conscious (i.e., deliberative thought) and unconscious (i.e., intuitions and/or heuristics) to make decisions [36]. When processing information, this combination becomes necessary to simplify decision making [5, 36], and as a byproduct, people are influenced by choice architectures (i.e., the way choices are presented to them) [35, 71, 81] and nudging (i.e., interventions that produce predictable deviations in behavior) [67, 80]. Even simple changes in the presentation of choices can have significant impact on people's decisions, as shown in domains like financial planning [4], food intake [38] or digital advertising [91].

Following this line of work, we investigate the potential of nudging in fulfilling social desiderata in online charities. Prior works have shown that people consciously recognize the 'economic need' as an important discerning factor for charitable giving [16]. As we observed in Section 3.2, this understanding does not directly translate into donation behavior. We hence hypothesize that donors may also be relying on unconscious processes while donating, and are likely amenable to nudging. In its current design, DonorsChoose presents users with a wide variety of filters to facilitate project browsing (detailed in Section 3.2.1), including a search bar on the website's main search page[5] and every project's funding status (i.e., how much money is still needed to get fully funded). Additionally, there is a dropdown menu allowing a donor to select one of the project sorting criterion: *most urgent*, *lowest cost to complete*, *highest economic need*, *fewest days left*, *most donors*, or *newest*.

In this section, we propose two poverty-related nudges on top of the current design, and run a series of controlled experiments to measure their impact on donation rates across various poverty levels. To compare the effect of the nudges, (i) we run all experiments on a real-time clone of the original platform, (ii) we do not add or remove projects from the platform, and (iii) we do not limit or constrain project search or filtering mechanisms in any way – thus respecting the donors' *freedom of choice*. Likewise, our interventions do not increase the complexity of the underlying choice architecture, nor introduce information previously unavailable on the platform.

4.1 Nudges applied

Depending on the conscious or unconscious impact of the intervention, prior works make a distinction between *overt* and *covert nudges* [31, 75]. If people realize that a nudge is present,

[5] https://www.donorschoose.org/donors/search.html

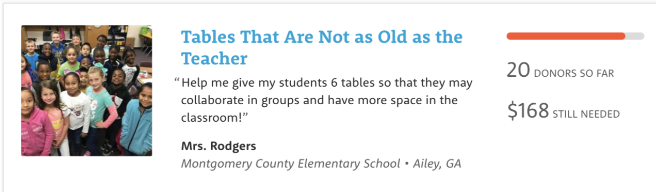

(a) Control: School poverty level is not shown.

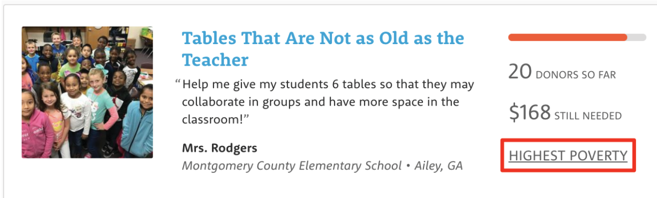

(b) Treatment: School poverty level is explicitly shown.

Fig. 3. **Experimental study: showing poverty level of the schools.**

and rationally account for it, the nudge is *overt*. Otherwise, a nudge would be considered *covert*. Depending on whether a given task more heavily relies on conscious or unconscious reactions, these types of nudges may demonstrate different levels of effectiveness. Overt nudges will have a greater impact on conscious tasks, whereas covert nudges will cater towards more unconscious ones. For our study, we focus on both overt and covert nudges specifically related to poverty:

- **Overt:** *explicitly show the poverty level of the school associated with each project proposal* (Fig. 3).
- **Covert:** *implicitly change the default ranking of shown projects to a poverty-based one* (Fig. 4).

Through a series of online controlled experiments, we want to measure to what extent the above two nudges can increase or decrease donation rates along different poverty levels. While we expect that the change in the choice architecture will have a significant impact on donors' behavior, there are several factors that need to be taken into account in our experiments. Among others, Sunstein [76, 77] notes that circumstantial preferences or a pre-built understanding of an environment (e.g., through past experiences) can incite different behaviors in otherwise comparable users. For example, some nudges may produce: (i) confusion in a target audience (e.g. when specific content is interpreted differently than intended), (ii) *reactance* (i.e., erratic behavior as an attempt to reclaim a sense of freedom), or even (iii) compensating behavior (i.e., behaving in a way that 'accounts for' and consequently 'contradicts' perceived nudges). Thus, there is a need for active experiments to understand how effective our covert and overt nudges are in controlling for donor behavior.

4.2 Nudging donors' decisions

We face two major challenges when attempting to test our interventions. First, after we choose a design element whose influence we want to measure, we have to make sure all other elements of the platform environment are fixed—including other website elements as well as donor characteristics. At the same time, we need to make sure that project offering remains similar across all of our experiments so as not to introduce additional confounding factors. While an approach based on A/B testing is a potential solution, A/B testing in scenarios where actual donors' decisions are nudged and the outcomes influence the real-world financial resources of schools raise ethical concerns.

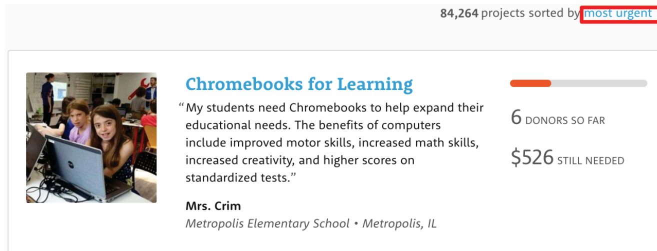

(a) Control: Default ranking is "most urgent".

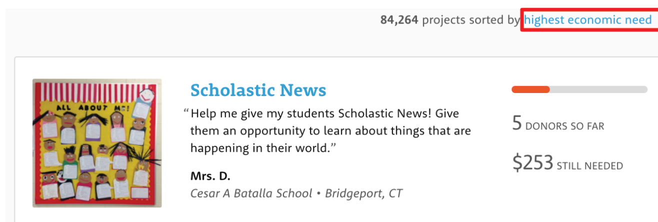

(b) Treatment: Default ranking is "highest economic need".

Fig. 4. **Experimental study: changing the default project ranking order to poverty-based.**

For example, beneficence—*"minimizing any risks to research subjects while maximizing research benefits"*—has been proposed as one of the ethical rules to follow in online experimentation [7].

To overcome these issues, we propose an experimental methodology where paid crowdsourcing workers are asked to make virtual donations which are nonetheless tied to their own money on a smaller scale (we present the compensation details later) on a proxy website mimicking the real DonorsChoose platform. This way, users have access to the entire pool of active projects on DonorsChoose at a particular time, and we are able to quantify the influence of individual website design elements on project selection rates—all the while avoiding manipulation of real donations.

4.2.1 *Setting up the proxy website.*

To conduct the experiments, we need to set up a platform that would act as a real-time proxy for DonorsChoose, while still allowing us to monitor users' activity and modify specific components of the website's design. To achieve this, we created our own web server which relays traffic to and from DonorsChoose, essentially acting as a middleman for our users' interactions with the original website. As an illustration, whenever a user requests a specific page, we make an equivalent request to DonorsChoose and: (i) replace all of the domain references; (ii) add scripts into the HTML for user activity monitoring; (iii) add references to custom JavaScript files responsible for the desired design changes; and (iv) send the final response back to the user. Any resource that should stay unchanged, such as images or project information, is simply forwarded to the user. By forcing all traffic to go through our server, we are able to modify the HTML, control DonorsChoose API calls and feed custom JavaScript files back to the user. We can thus actively redesign UI components and monitor user activity on top of an otherwise authentic experience.

4.2.2 *Study design.*

We recruited 150 participants from Amazon Mechanical Turk (AMT, mturk.com) who were residents of the US and aged 18 years and above. The participants were told that we were developing a system for recommending projects to potential donors, and wish to study their donation behavior and preferences. We gave each participant a virtual budget of $100[6], and asked them to make virtual donations to any project(s) they preferred. We further tied this budget to a $1 bonus they could receive upon completing the experiment on top of their participation compensation. Specifically, the percentage of donated money would inversely determine their bonus in a 1/100 scale (e.g. donating $90 will lead to a final bonus of just ($100 - $90)/100 = $0.1). We also promised and later matched all their selected donations project-by-project on the real platform (in the same 1/100 scale). Thus, the workers' charitable contribution had a real impact. Since the donation is tied to their own money, it not only motivates the workers to carefully select projects they deem worth helping but also helps us filter out non-altruistic participants who will either donate a very small amount or skip the donation process altogether, maximizing their own reward.

As a final step, participants completed a follow-up survey in which we collected their demographic information, and asked about their donation preferences in more detail. 64% of the participants were men and the average age was 32. On average, participants spent around 10 minutes completing this task, and received $3 (before bonus calculation) for participating in the experiment.

Since we wanted to test two interventions, we randomly assigned the participants into three groups (with 50 people in each): one control group and two treatment groups. Each treatment group interacted with a version of the website where a single design element was changed, whereas the control group received the default DonorsChoose experience. We ran all these experiments at the same time to minimize any variation in the underlying availability or funding status of the projects on DonorsChoose. Each treatment group was exposed to one of the following website changes:

1. Explicitly showing the poverty level of the schools
DonorsChoose measures the poverty level of a school based on what fraction of its students come from low-income families, and categorizes it into four categories: highest, high, moderate, and low. As Fig. 3 shows, we added the information about the poverty level of the schools to project listings. Participants in the treatment group could see the poverty level of the schools before selecting the projects to donate to (Fig. 3b), whereas the control group participants did not see the poverty level of a school unless they went to a specific project page (Fig. 3a).

2. Changing the default project ranking to poverty-based
There are several options to rank projects on the project search page. The participants in the control group saw the projects ranked based on the default *most urgent* criterion (Fig. 4a), whereas the treatment group was exposed to projects ranked based on the poverty level (*highest economic need* in DonorsChoose parlance) by default (Fig. 4b).

4.2.3 Experimental validity.
We acknowledge the limitation of our experimental design in reproducing truly altruistic donor motivations. Even though we attempted to recreate an authentic platform experience as closely as possible, our study participants (i) did not spend their own money, except for the small compensation bonus, and (ii) were aware that they were undergoing an experiment. However, an experimental setup that meddled with real donations would raise an array of ethical concerns. Despite the limitations of our experiments, we can hypothesize about the donor behavior in an alternative, truly altruistic setup through an analysis of prior literature.

First, as documented by Tversky and Kahneman [83], people are cognitively constrained by *loss aversion*, which is the tendency to prevent loss. During our experiments, participants were given the option to donate as much of a predefined budget as they wanted, resulting in a proportionately lower monetary bonus at the end. Thus, participants did not lose money directly by donating,

[6] To make the budget close to the average donation amount in the DonorsChoose dataset ($80).

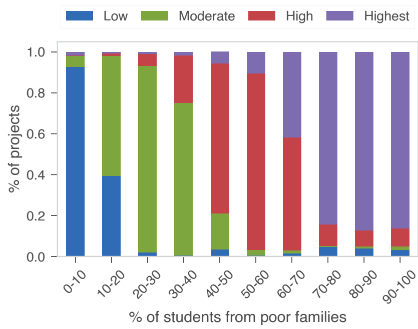

Fig. 5. **Comparing DonorsChoose's poverty categories with % of students from poor households, we see that 'Low', 'Moderate', 'High' and 'Highest' poverty are (respectively) represented by 0-10%, 10-40%, 40-70% and 70-100% poor students.**

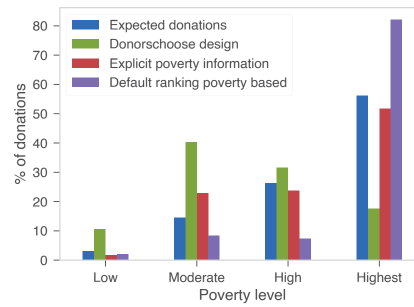

Fig. 6. **Effect of different nudges on the distribution of donation: higher amount of donations go to poorer districts when poverty level is explicitly shown, or the default project ranking is changed to poverty based.**

but rather gained less. If they were using their own money for donating, they may have actually perceived donating as a monetary loss. Although it is likely that donating with one's own money would have resulted in lower donated amounts, we have no reason to believe that it would have changed the way people allocate resources along the poverty level.

Regarding participants' awareness, a prior study by Loewenstein et al. [43] has shown that merely informing users about the presence of nudges does not interfere with their behavior. It is possible that preset defaults would raise psychological reactance in users if they were either perceived as a hindrance to their freedom of choice, or a means to manipulate their behavior away from a platform's purpose. As Cheung and Chan [16] have demonstrated, people consciously recognize economic need as an important factor for charitable giving, so we do not expect our poverty-based interventions to have caused psychological reactance. In particular, we expect the behavior of uninformed donors (who use the platform for the first time) to closely resemble that of our participants.

Since our work focused on the impact of interface design, the validity of our analysis relies on the comparison between multiple experimental conditions. While we are not able to immediately extend all our conclusions to DonorsChoose full ecosystem, our work shows the potential of choice architecture and nudging through the interface design in an online charity setting.

4.3 Experimental findings

During our live experiments, we could not collect the fine grained poverty levels of the school districts. DonorsChoose provides a coarse-grained poverty level that, despite mappable to their datasets' poverty categories (i.e., Highest, High, Moderate and Low), it is not to the NCES dataset. Although limited by these coarse-grained categories, in Fig. 5 we can see the distribution of DonorsChoose's poverty categories over probable NCES' district poverty levels.

Proceeding with our analysis, we first define an *expected donation* baseline for our experiments. As we decided not to run the experiments with sampled projects (and go with all the ones available in the platform), we do not know the exact project distribution overall – aside from the fact it is similar across experiments. However, by looking at historical data, we observe that project availability remains identical over time at each poverty level. We will hence use it as an expectation for the percentage of expected donations, when comparing against the donations in our experiments (Fig. 6).

4.3.1 General observations.

In online controlled experiments, it is standard to compute an *Overall Evaluation Criterion* (also known as *Dependent Variable*, or *Key Performance Indicator*), which gives a quantitative measure of the experiment's objective [39, 40]. We will compare this criterion in the control and treatment conditions to see whether the treatment had any significant impact. For our two experiments, we use *the fraction of donations going to highest poverty schools* as the evaluation criterion.

One interesting matter to note is the difference between the expected project availability and actual project donations for the control group (i.e., donors who encounter the default DonorsChoose design). In Fig. 6, we can see that the number of available projects increases steadily from low to highest poverty, whereas donors seem to have donated mostly to moderate poverty projects. Although not necessarily an undesirable artifact, it shows that donations are not being randomly assigned and there exists some sort of selection process (either imposed by the users themselves or the platform). This difference is most noticeable at the extreme, at highest poverty. From 56% of all the available projects on the platform being from that poverty level, we observe just 18% of donations actually going to that category. In what follows, we will highlight how this changes across our treatment groups.

4.3.2 Impact of explicitly showing poverty level of schools.

When we explicitly showed the poverty level of schools, projects with high and highest poverty got more donations compared to projects with moderate poverty (Fig. 6), more closely resembling our random baseline. Computing the Mann-Whitney U-test [53],[7] we find the differences in the donation distributions of the control and treatment groups to be statistically significant ($p < 0.05$). Thus, we can conclude that showing school poverty information explicitly can push the donors to donate more to poorer schools.

Monitoring other activities of the participants, we did not observe any significant differences in the usage of search filters. However, both the number of unique search pages and unique filters used were reduced by 50% and 40%, respectively. On top of that, the average number of donated project per user went up to 2.7 in the treatment group (from 2.3 in the control group) and the average amount spent in donations went up to $87.7 (from $77.7 in the control group). This signifies that participants were not only donating to more projects, but also donating higher amounts and more easily finding suitable project proposals.

4.3.3 Impact of changing the default project ranking to based on poverty level.

Changing the default ranking of projects to poverty-based led to a considerable change in the donation distribution. Fig. 6 shows that around 82% of donations went to schools with highest poverty levels when the default ranking is poverty-based (compared to 18% when the default ranking is *most urgent*). Mann-Whitney U-test confirms that the difference is statistically significant ($p < 0.05$).

Interestingly, the fraction of users changing the default ranking order remained almost the same (4% in the treatment group against 3% in the control group) but poverty-related filters like 'Special Needs' and 'Warmth, Care & Hunger' saw a decrease in usage up to 20%. Regardless of lower overall search pattern complexity, the average amount of donated money still went up (to $85.9, from $77.7 in the control group). Even if the amount of unique pages visited remained comparable

[7] We did not assume any underlying statistical distributions in the donation outcomes, and hence used the non-parametric U-test to compare the outcomes in the control and treatment groups.

	Nudge Type	# total donations	$ total donated	% donations to highest poverty
Default design	–	115	3,885	18
Showing poverty level	Overt	135 (+ 17.4%)	4,385 (+ 12.9%)	52 (+ 34%)
Changing default ranking	Covert	95 (- 17.4%)	4,295 (+ 10.5%)	82 (+ 64%)

Table 1. **Key observations from our two experimental setups, utilizing overt and covert nudges. All results are shown alongside a percentual increase/decrease, using the control group as a reference.**

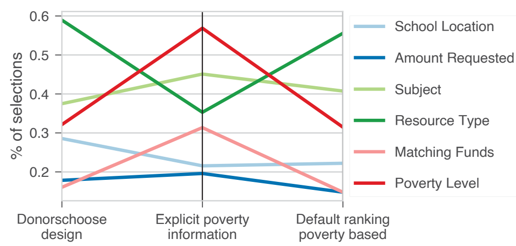

Fig. 7. **Percentage of participants that reported a given feature being important in their decision making, for all the experimental settings described in Section 4.**

to the control group, donations were now vastly within a single, poverty-based sorting criterion. This implies that our participants were once again able to more easily identify suitable projects.

As summarized in Tab. 1, our experiments suggest that nudges can indeed have a big impact on donors' choices. From what we observe, they can mediate between significantly different behaviors and potentially help platforms achieve fair(er) objectives. In charity-related platforms where people's livelihoods may depend on the funding they receive, nudging can easily take on a paternalistic approach in driving their users towards 'good' objectives. That said, the definition for 'good' can be much harder to define, and should require more interactions between all the stakeholders.

4.4 Awareness of nudging and freedom of choice

After completing their respective experiments, all of our participants were presented with the following checklist, regarding their preferences while donating.

Question: Which of the following features played a role in YOUR donation choices? (Check all that apply):

- ■ School poverty (number of students coming from poor families);
- ■ School locality (the school being close to your place of residence);
- ■ Subject of the project (Science, History, Sports etc.);
- ■ Resource requested (Books, Computers etc.);
- ■ Project would get matching donations;
- ■ How much donation was requested.

Each of the above options was tied to one of the project features on the platform. By aggregating all the votes for each experiment, we noticed that responses mostly did not change between the control group and the group with the default ranking changed (Fig. 7). The difference was significant, however, when we explicitly showed the poverty level alongside the projects' descriptions (selection of 'School poverty' increased from 32% to 57%, selection of 'Project would get matching donation' saw a smaller increase from 16% to 31%, whereas the selection of 'Resource requested' decreased from 59% to 35%).

These results demonstrate an interesting trend when compared with the donation distribution results. In the condition with the strongest effect on the distribution of donations (i.e., changing the default ranking), users barely reported any change to the importance of ‘Poverty Level’ in their decision, yet their behavior along this feature clearly deviated from both the control group and our expected baseline. On the other hand, in the condition with the weakest effect on donations (i.e., explicitly showing poverty level), ‘Poverty Level’ clearly was raised as the most important decision-making factor, yet the donations were almost randomly allocated along that feature. Within the context of our experiments, these findings are not unexpected. Based on the notions of overt and covert nudges, we would expect our two interventions to tap into conscious and unconscious thought processes differently. While showing poverty level should engage in more conscious reflection, changing the default ranking would instead tap into unconscious decisions. What we observe through our survey may be the impact of interventions on the users’ awareness.

Following Sunstein [75]’s definitions, if users *know* and *control* the choices they make, they are autonomous and hence free. Similarly to other studies, we presume the freedom of choice of our participants to be respected if control is given to them (i.e., a choice between multiple alternatives is entirely left to their discretion). However, since both conscious and unconscious processes are conditioned by a choice architecture, it may be harder to ascertain whether users know (or are aware of) the choices they make. These results alone are weak evidences on the different effects of overt and covert nudges, but they raise important questions about the potential impact of different interventions in socio-technical systems.

5 CONCLUDING DISCUSSION: TO NUDGE OR NOT TO NUDGE

Leveraging the interface design to achieve equity goals is an under-explored research direction. Recently, multiple works on algorithmic fairness have contributed model interventions applicable in different places of a system’s pipeline, including pre-processing (intervening on the training data), in-processing (intervening directly on the learning algorithm), or post-processing (intervening on the trained algorithm or the results of the algorithm) [6, 15, 57, 58, 72, 74, 92]. For example, in the context of fair search and information finding, fairness interventions modify the learning objectives of ranking functions [72, 92], or reshuffle ranking results [6]. Yet, in an information finding scenario, users do not interact with the model or the data directly, but rather through an interface. In this way, the presentation of the results is the final step that determines whether a system achieves its social or fairness goals. Recent studies present evidence that post-processing based fairness interventions (such as reranking search results) might be ineffective in scenarios where people have strong preferences for specific choices [59]. Leveraging a long line of research in behavioral economics [4, 35, 36, 80, 81], specifically on nudging and choice architectures, we utilize the potential role of interfaces in controlling for these preferences. This allows us to design socio-technical systems that present information to people in a way that does not limit their freedom of choice, while acknowledging that any information presentation will nudge people towards certain choices.

Desiderata of online altruism. By definition, no interface or its underlying choice architecture can ever be “neutral”. Even a blank page will present a choice between waiting for something to load or leaving the page. People naturally reduce complex environments to oversimplified categories [8], so the mere act of interacting with the platform and observing its results will induce certain reactions and conclusions. In online platforms, choice architectures are embedded within the interface designs, and currently they mostly optimize for different monetary objectives such as maximizing customer retention rates or video watch times. However, as we show in this paper, they also have the power to actively nudge their users towards socially desirable goals. It is necessary to recognize that specifying such socially desirable nudging goals is hard. Even in scenarios like charitable giving, a seemingly appropriate setting for paternalistic and socially-oriented nudging, a fair donation distribution could arguably be defined in many ways:

- Satisfying the needs of highest poverty groups first;
- Equalizing funding per student across all poverty levels;
- Equalizing funding for similar projects;
- Preferentially funding projects according to other societal agreed-upon criteria.

It is possible to imagine circumstances under which each of these objectives would be desirable. Proponents of Equality of Resources for education have argued strongly that the playing field must be leveled [11, 41, 78]. Advocates of Equality of Capability [65] advocate for an unequal distribution of funds *in favor* of disadvantaged students (e.g., to equalize the *capability* of getting into a reputable college). In this paper, despite our stance in supporting the later, we do not regard it as the only fair option. Regardless, if our goal is to satisfy a socially desirable platform goal while maximizing individual user *autonomy*, nudging may be a necessity [75].

Ethics of nudging. The question of ethicality has been present in the choice architecture literature due to the seemingly manipulative nature of nudging. Popular understanding of this concept seems to align with Raz's [61] definition, where manipulation is described as anything that "perverts the way [in which] a person reaches decisions, forms preferences or adopts goals". In essence, that is what a nudge is. However, this notion is often criticized for being too broad. Wilkinson [89] expanded on this concept, and redefined manipulation as "intentionally and successfully influencing someone using methods that pervert choice". The author also added that manipulation is something that bypasses an individual's rational capacities. Schüll [68] gave a clear example of this in the context of gambling, and its potential consequences (such as addiction). Coons and Weber [17], despite mostly aligning their views with Wilkinson's, added that manipulation can be achieved through rational interference as well (e.g., misinformation). On the surface, their definition seems more comprehensive than Raz's and feels more intuitive. But if we accept *rationality* and *irrationality* as complementary concepts, the only difference between the two is *intention*. According to Wilkinson [89] and Coons and Weber [17], perverting others' choices is manipulative only if it is intentional.

This definition is by no means generally accepted. For instance, Sunstein [75] believes manipulation to be more closely related to *autonomy* and *self-governance* than it is to *intention*. Winner [90] also discussed this phenomenon through *the politics of artefacts*, where technological development will have direct implications in the way others build perceptions, independently of its intended purpose or use. In this line of thought, what makes something manipulative is its inherent propensity to constrain its targets' *freedom of choice* (either by selectively relaying information or by exploiting underlying biases). More than just among scholars, and in the specific setting of nudging, this notion seems to be shared by a wider audience [31] who, when surveyed about the ethics of overt and covert nudges, referred to overt nudges as 'less manipulative' and more 'autonomy-preserving' – despite both being intended by the choice architect.

According to all these definitions, nudges can be considered manipulative. However, following Sunstein's definition, any existing platform could be considered manipulative in the same way. We argue that for nudging and choice architectures to be ethical, their intentional implementation should: (i) account for users' cognitive biases and compensate for them towards *socially desirable* goals; and (ii) strive to preserve users' *freedom of choice*.

Beyond nudging. As noted earlier in this paper, donors might not donate in a globally optimal way and it may be necessary for platforms to find ways to compensate for inequitable donation outcomes. Even though nudging is a possible way of controlling for donor behavior, the selection of effective interventions remains elusive. Thus, an important question arises as to whether charities should allow for targeted donations, or allocate donations themselves. For instance, by using the percentage of students from poor households as a proxy for school poverty, we can establish a *hierarchy of needs* to satisfy equitably, starting from schools with the highest percentage of students from poor households. To achieve this goal, there exist several alternative mechanisms that could complement nudging:

- **Non-targeted donations** would invariably lead to the most equitable allocation, as the platform would be able to distribute all incoming donations along its defined objective;
- **Partially targeted donations** could allow donors to specify their feature-level donation preferences, without selecting specific projects;
- **Fully targeted donations** could still be allowed, by converting a percentage of every donation into a non-targeted redistribution pool, allocatable by the platform.

Any of these approaches may have unintended side-effects on the platform's dynamics. On the positive side, non-targeted donations could prove very efficient, and even increase the platform engagement from schools with highest need (as they would immediately benefit from more referrals towards the platform). However, these could simultaneously cause schools with lower need to decrease their engagement, all the while ignoring donor preferences. On the other hand, partially targeted donations would respect donors' preferences, at the cost of a less efficient allocation. These could also condition engagement based on most/less desirable features. Lastly, fully targeted donations may preserve schools' engagement and donor preferences, but they may decrease donors' incentive to donate, as conversion rates get higher.

Future directions. Multiple directions remain open for future work.

- To tackle the issue of external validity, we aim to partner with an existing charity platform – to perform similar experiments at a larger scale. To mitigate the ethical concerns of meddling with real donations, one solution could be to pool donations and distribute them evenly after the experimental phase concludes.
- We also plan to automate the nudge selection process. While the implementation of independent interventions seems trivial, the interactions between different nudges may be quite complex and would require solving optimization problems that select configurations of nudges to satisfy the given objective. Reinforcement Learning seems to be a promising approach for this task, for its effectiveness in pruning complex search spaces.
- Beyond making online platforms more fair, it would be important to better quantify users' awareness of a platform's influence on their behavior, as well as understand how to control this influence when designing a choice architecture.

Overall, we hope that our work will pave the way for future studies aiming to employ digital nudges effectively to contribute to the social and ethical desiderata of socio-technical systems.

ACKNOWLEDGMENTS

We thank the anonymous reviewers whose suggestions have helped us to improve the paper significantly. This research was supported in part by a European Research Council (ERC) Advanced Grant for the project "Foundations for Fair Social Computing", funded under the European Union's Horizon 2020 Framework Programme (grant agreement no. 789373).

REFERENCES

[1] 2019. Separate and Unequal: School Funding in 'Post-racial' America. https://www.topmastersineducation.com/school-funding-post-racial-us

[2] Tim Althoff and Jure Leskovec. 2015. Donor retention in online crowdfunding communities: A case study of donorschoose. org. In *WWW*. 34–44.

[3] Jonathan Baron and Ewa Szymanska. 2011. Heuristics and biases in charity. *The science of giving: Experimental approaches to the study of charity* (2011), 215–235.

[4] Shlomo Benartzi, Ehud Peleg, and Richard H Thaler. 2013. Choice architecture and retirement saving plans. *The behavioral foundations of public policy* (2013).

[5] Buster Benson and J Manoogian. 2018. Cognitive bias codex.

[6] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *ACM SIGIR*. 405–414.

[7] Sarah Bird, Solon Barocas, Kate Crawford, Fernando Diaz, and Hanna Wallach. 2016. Exploring or exploiting? social and ethical implications of autonomous experimentation in ai. (2016).

[8] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences.* MIT press.

[9] Beth Breeze. 2013. How donors choose charities: The role of personal taste and experiences in giving decisions. *Voluntary Sector Review* 4, 2 (2013), 165–183.

[10] Robin N Brewer and Jasmine Jones. 2015. Pinteresce: exploring reminiscence as an incentive to digital reciprocity for older adults. In *ACM CSCW*. 243–246.

[11] Harry Brighouse and Adam Swift. 2006. Equality, priority, and positional goods. *Ethics* 116, 3 (2006), 471–497.

[12] Alexander L Brown, Jonathan Meer, and J Forrest Williams. 2017. Social distance and quality ratings in charity choice. *Journal of behavioral and experimental economics* 66 (2017), 9–15.

[13] Catherine Brown, Scott Sargrad, and Meg Benner. 2017. Hidden Money: The Outsized Role of Parent Contributions in School Finance. https://www.americanprogress.org/issues/education-k-12/reports/2017/04/08/428484/hidden-money

[14] Abhijnan Chakraborty, Nuno Mota, Asia J Biega, Krishna P Gummadi, and Hoda Heidari. 2019. On the Impact of Choice Architectures on Inequality in Online Donation Platforms. In *WWW*.

[15] Abhijnan Chakraborty, Gourab K Patro, Niloy Ganguly, Krishna P Gummadi, and Patrick Loiseau. 2019. Equality of voice: Towards fair representation in crowdsourced top-k recommendations. In *ACM FAT**. 129–138.

[16] C-K Cheung and C-M Chan. 2000. Social-cognitive factors of donating money to charity, with special attention to an international relief organization. *Evaluation and program planning* 23, 2 (2000), 241–253.

[17] Christian Coons and Michael Weber. 2014. *Manipulation: theory and practice.* Oxford University Press.

[18] Stephen Q. Cornman, Laura C. Nixon, Matthew J. Spence, Taylor Lori L., and Douglas E. Geverdt. 2019. *American Community Survey Comparable Wage Index for Teachers (ACS-CWIFT).* Technical Report.

[19] Antonella Corsi-Bunker. 2015. Guide to the education system in the United States. *University of Minnesota* 23 (2015).

[20] Cynthia Cryder, Simona Botti, and Yvetta Simonyan. 2017. The Charity Beauty Premium: Satisfying Donors 'Want' Versus 'Should' Desires. *Journal of Marketing Research* 54, 4 (2017), 605–618.

[21] Mette Trier Damgaard and Christina Gravert. 2017. Now or never! The effect of deadlines on charitable giving: Evidence from two natural field experiments. *Journal of Behavioral and Experimental Economics* 66 (2017), 78–87.

[22] Joy Dodge. 2018. Redrawing School District Lines: Reducing the Link between Educational Inequality and Economic Inequality. *Georgetown Journal on Poverty Law and Policy* 26 (2018), 165.

[23] DonorsChoose. 2016. Download Opendata. https://research.donorschoose.org/t/download-opendata/33

[24] DonorsChoose. 2019. See our impact nationwide since our start in 2000. https://donorschoose.org/about/impact.html

[25] Esther Duflo, Michael Kremer, and Jonathan Robinson. 2011. Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American economic review* 101, 6 (2011), 2350–90.

[26] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *ITCS*. 214–226.

[27] Catherine C Eckel, David H Herberich, and Jonathan Meer. 2017. A field experiment on directed giving at a public university. *Journal of behavioral and experimental economics* 66 (2017), 66–71.

[28] EdBuild. 2019. Non-white school districts get $23 Billion less than white districts, despite serving the same number of students. https://edbuild.org/content/23-billion

[29] James T Edwards and John A List. 2014. Toward an understanding of why suggestions work in charitable fundraising: Theory and evidence from a natural field experiment. *Journal of Public Economics* 114 (2014), 1–13.

[30] Vincent Etter, Matthias Grossglauser, and Patrick Thiran. 2013. Launch hard or go home!: predicting the success of kickstarter campaigns. In *ACM COSN*. 177–182.

[31] Gidon Felsen, Noah Castelo, and Peter B Reiner. 2013. Decisional enhancement and autonomy: public attitudes towards overt and covert nudges. *Judgment & Decision Making* 8, 3 (2013).

[32] Michael D Greenberg, Bryan Pardo, Karthic Hariharan, and Elizabeth Gerber. 2013. Crowdfunding support tools: predicting success & failure. In *CHI Extended Abstracts*. ACM, 1815–1820.

[33] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions. In *ACM CHI*. 2647–2656.

[34] C Kirabo Jackson, Rucker C Johnson, and Claudia Persico. 2015. The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *The Quarterly Journal of Economics* 131, 1 (2015), 157–218.

[35] Eric J Johnson, Suzanne B Shu, Benedict GC Dellaert, Craig Fox, Daniel G Goldstein, Gerald Häubl, Richard P Larrick, John W Payne, Ellen Peters, and David Schkade. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504.

[36] Daniel Kahneman. 2011. *Thinking, fast and slow.* Macmillan.

[37] Dean Karlan and John A List. 2007. Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *American Economic Review* 97, 5 (2007), 1774–1793.

[38] Carmen Keller, Franziska Markert, and Tamara Bucher. 2015. Nudging product choices: The effect of position change on snack bar choice. *Food Quality and Preference* 41 (2015).

[39] Ron Kohavi, Alex Deng, Brian Frasca, Roger Longbotham, Toby Walker, and Ya Xu. 2012. Trustworthy online controlled experiments: Five puzzling outcomes explained. In *ACM KDD*.

[40] Ron Kohavi and Roger Longbotham. 2017. Online controlled experiments and a/b testing. In *Encyclopedia of machine learning and data mining*.

[41] William Koski and Rob Reich. 2008. The state's obligation to provide education: adequate education or equal education. In *Meeting of the Midwest Political Science Association, Chicago*, Vol. 45.

[42] Julien Lafortune, Jesse Rothstein, and Diane Whitmore Schanzenbach. 2018. School finance reform and the distribution of student achievement. *American Economic Journal: Applied Economics* 10, 2 (2018), 1–26.

[43] George Loewenstein, Cindy Bryce, David Hagmann, and Sachin Rajpal. 2015. Warning: You are about to be nudged. *Behavioral Science & Policy* 1, 1 (2015), 35–42.

[44] Chun-Ta Lu, Sihong Xie, Xiangnan Kong, and Philip S Yu. 2014. Inferring the impacts of social media on crowdfunding. In *WSDM*. ACM, 573–582.

[45] Thomas Lux. 1995. Herd behaviour, bubbles and crashes. *The economic journal* (1995).

[46] Joel McFarland, Bill Hussar, Xiaolei Wang, Jijun Zhang, Ke Wang, Amy Rathbun, Amy Barmer, Emily Forrest Cataldi, and Farrah Bullock Mann. 2018. The Condition of Education 2018. *National Center for Education Statistics* (2018).

[47] Marilyn McGee-Lennon, Maria Wolters, Ross McLachlan, Stephen Brewster, and Cordelia Hall. 2011. Name that tune: musicons as reminders in the home. In *ACM CHI*. 2803–2806.

[48] Jennifer Medina. 2019. At Los Angeles Teachers' Strike, a Rallying Cry: More Funding, Fewer Charters. https://www.nytimes.com/2019/01/17/us/lausd-strike-schools.html

[49] Jonathan Meer. 2014. Effects of the price of charitable giving: Evidence from an online crowdfunding platform. *Journal of Economic Behavior & Organization* 103 (2014), 113–124.

[50] Jonathan Meer. 2017. Does fundraising create new giving? *Journal of Public Economics* 145 (2017), 82–93.

[51] Tanushree Mitra and Eric Gilbert. 2014. The language that gets people to give: Phrases that predict success on kickstarter. In *ACM CSCW*. 49–61.

[52] Emilia Murphy. 2018. Get Your Classroom Project Funded: The 4 Essential Steps. https://www.donorschoose.org/blog/get-your-classroom-project-funded

[53] Nadim Nachar et al. 2008. The Mann-Whitney U: A test for assessing whether two independent samples come from the same distribution. *Tutorials in quantitative Methods for Psychology* 4, 1 (2008), 13–20.

[54] United Nations. 1948. Universal Declaration of Human Rights. https://un.org/en/universal-declaration-human-rights/

[55] NCES. 2019. Comparable Wage Index for Teachers (CWIFT). https://nces.ed.gov/programs/edge/Economic/TeacherWage

[56] Chankyung Pak and Rick Wash. 2017. The Rich Get Richer? Limited Learning in Charitable Giving on donorschoose. org.. In *AAAI ICWSM*. 172–181.

[57] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2020. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. In *WWW*. 1194–1204.

[58] Gourab K Patro, Abhijnan Chakraborty, Niloy Ganguly, and Krishna Gummadi. 2020. Incremental Fairness in Two-Sided Market Platforms: On Smoothly Updating Recommendations. In *AAAI*, Vol. 34. 181–188.

[59] Andi Peng, Besmira Nushi, Emre Kıcıman, Kori Inkpen, Siddharth Suri, and Ece Kamar. 2019. What You See Is What You Get? The Impact of Representation Criteria on Human Bias in Hiring. In *AAAI HCOMP*, Vol. 7. 125–134.

[60] J Radford. 2016. The emergence of gender inequality in a crowdfunding market: an experimental test of gender system theory. (2016).

[61] Joseph Raz. 1986. *The morality of freedom.* Clarendon Press.

[62] William S Reece. 1979. Charitable contributions: New evidence on household behavior. *The American Economic Review* 69, 1 (1979), 142–151.
[63] David Reinstein and Gerhard Riener. 2012. Reputation and influence in charitable giving: an experiment. *Theory and decision* 72, 2 (2012), 221–243.
[64] Meredith P Richards. 2014. The gerrymandering of school attendance zones and the segregation of public schools: A geospatial analysis. *American Educational Research Journal* 51, 6 (2014).
[65] Ingrid Robeyns. 2005. The capability approach: a theoretical survey. *Journal of human development* 6, 1 (2005), 93–117.
[66] William Samuelson and Richard Zeckhauser. 1988. Status quo bias in decision making. *Journal of risk and uncertainty* 1, 1 (1988).
[67] Christoph Schneider, Markus Weinmann, and Jan vom Brocke. 2018. Digital nudging: guiding online user choices through interface design. *Commun. ACM* 61, 7 (2018), 67–73.
[68] Natasha Dow Schüll. 2014. *Addiction by design: Machine gambling in Las Vegas*. Princeton University Press.
[69] Genevieve Siegel-Hawley. 2013. Educational gerrymandering? Race and attendance boundaries in a demographically changing suburb. *Harvard Educational Review* 83, 4 (2013).
[70] Herbert A Simon. 1955. A behavioral model of rational choice. *The quarterly journal of economics* 69, 1 (1955), 99–118.
[71] Itamar Simonson. 1989. Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research* 16, 2 (1989), 158–174.
[72] Ashudeep Singh and Thorsten Joachims. 2019. Policy learning for fairness in ranking. In *NeurIPS*. 5427–5437.
[73] Jacob Solomon, Wenjuan Ma, and Rick Wash. 2015. Don't wait!: How timing affects coordination of crowdfunding donations. In *ACM CSCW*. 547–556.
[74] Tom Sühr, Asia J Biega, Meike Zehlike, Krishna P Gummadi, and Abhijnan Chakraborty. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *ACM KDD*. 3082–3092.
[75] Cass R Sunstein. 2015. Nudging and choice architecture: Ethical considerations. *Yale Journal on Regulation, Forthcoming* (2015).
[76] Cass R Sunstein. 2017. Misconceptions about nudges. *Available at SSRN 3033101* (2017).
[77] Cass R Sunstein. 2017. Nudges that fail. *Behavioural Public Policy* 1, 1 (2017), 4–25.
[78] Adam Swift. 2003. *How not to be a hypocrite: School choice for the morally perplexed parent*. Routledge.
[79] Katie G Tanaka and Amy Voida. 2016. Legitimacy work: Invisible work in philanthropic crowdfunding. In *ACM CHI*. 4550–4561.
[80] Richard H Thaler and Cass R Sunstein. 2008. Nudge: Improving decisions about health, wealth, and happiness.
[81] Richard H Thaler, Cass R Sunstein, and John P Balz. 2014. Choice architecture. (2014).
[82] Thanh Tran, Madhavi R Dontham, Jinwook Chung, and Kyumin Lee. 2016. How to succeed in crowdfunding: a long-term study in kickstarter. *arXiv preprint arXiv:1607.06839* (2016).
[83] Amos Tversky and Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5, 4 (1992), 297–323.
[84] Food US Department of Agriculture and Nutrition Service. 2014. Income eligibility guidelines. https://www.govinfo.gov/content/pkg/FR-2014-03-05/pdf/2014-04788.pdf
[85] Giving USA. 2018. Americans Gave $410.02 Billion to Charity in 2017, Crossing the $400 Billion Mark for the First Time. https://givingusa.org/giving-usa-2018-americans-gave-410-02-billion-to-charity-in-2017-crossing-the-400-billion-mark-for-the-first-time
[86] Rick Wash and Jacob Solomon. 2014. Coordinating donors on crowdfunding websites. In *ACM CSCW*. 38–48.
[87] Karen Weese. 2018. Parent-Led Fundraising Makes Some Schools Better but Leaves Others Behind. https://slate.com/human-interest/2018/11/parent-teacher-organizations-education-inequality.html
[88] Gillian White. 2015. The Data Are Damning: How Race Influences School Funding. https://www.theatlantic.com/business/archive/2015/09/public-school-funding-and-the-role-of-race/408085
[89] T Martin Wilkinson. 2013. Nudging and manipulation. *Political Studies* 61, 2 (2013), 341–355.
[90] Langdon Winner. 1980. Do artifacts have politics? *Daedalus* (1980), 121–136.
[91] Dawei Yin, Bin Cao, Jian-Tao Sun, and Brian D Davison. 2014. Estimating ad group performance in sponsored search. In *WSDM*. ACM, 143–152.
[92] Meike Zehlike and Carlos Castillo. 2020. Reducing disparate exposure in ranking: A learning to rank approach. In *WWW*. 2849–2855.

Received January 2020; revised June 2020; accepted July 2020