

On the Impact of Choice Architectures on Inequality in Online Donation Platforms

Abhijnan Chakraborty*
MPI for Software Systems

Nuno Mota*†
EPFL

Asia J. Biega
MPI for Informatics

Krishna P. Gummadi
MPI for Software Systems

Hoda Heidari
ETH Zürich

ABSTRACT

Online donation platforms, such as DonorsChoose, GlobalGiving, or CrowdFunder, enable donors to financially support entities in need. In a typical scenario, after a fundraiser submits a request specifying her need, donors contribute financially to help raise the target amount within a pre-specified timeframe. While the goal of such platforms is to counterbalance societal inequalities, biased donation trends might exacerbate the unfair distribution of resources to those in need. Prior research has looked at the impact of biased data, models, or human behavior on inequality in different socio-technical systems, while largely ignoring the *choice architecture*, in which the funding decisions are made.

In this paper, we focus on (i) quantifying inequality in the project funding in online donation platforms, and (ii) understanding the impact of platform design on donors' behavior in magnifying those inequalities. To this end, we borrow decomposable measures of income inequality from economics, and apply it to identify candidate factors contributing to inequality on the DonorsChoose website. Analyzing longitudinal changes in the website design, we show how the platform design impacts the relative contribution of the different factors. Our work motivates the need for a careful investigation of the impact of choice architectures on user decisions, in donation platforms in particular, and in online platforms more generally.

CCS CONCEPTS

• Human-centered computing → Empirical studies in HCI.

ACM Reference Format:

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 *Proceedings of the 2019 World Wide Web Conference (WWW '19), May 13–17, 2019, San Francisco, CA, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3308558.3313663>

*Both authors contributed equally to this work.

†This work was done while the author was interning at MPI for Software Systems.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.

<https://doi.org/10.1145/3308558.3313663>

1 INTRODUCTION

Online donation platforms, such as DonorsChoose, Chuffed, GlobalGiving, CrowdFunder or Facebook fundraising, provide effective platforms for connecting donors to entities in need of financial support. In a typical scenario, after a recipient submits a request specifying her need, donors can contribute financially to help raise the target amount within a pre-specified timeframe. The core mission of such platforms is to counter existing social inequalities. Biased donor behavior, however, may skew the distribution of donations towards or against certain recipients, exacerbating some of these inequalities as a result. Consider DonorsChoose.org, a platform where US-based school teachers can create funding requests to support the educational needs of their classrooms. Since public schools in the US are primarily supported by local taxes [24], schools located in poor neighborhoods are often severely underfunded. If the donors have a preference to donate to schools in their own region, the funding inequality combined with biased donation behavior could lead to a *rich get richer* effect, where wealthy schools easily receive additional funding in the form of online donations, while schools in poor neighborhoods remain impoverished [25].

At first glance, such biased outcomes may seem to be the sole artifact of donors' preferences rather than that of the platform mediating donations. After all, donors are rational beings, and the donation platforms offer them full freedom of choice. However, 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 without limiting their choices: Donation websites present requests within a *choice architecture* [33], influencing donors' decision-making and *nudging* them towards choosing certain alternatives. Examples of nudges [32] on DonorsChoose include the *default* ranking strategy used in project listings, *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 nudges in socio-technical systems, while our empirical contributions focus particularly on online donation platforms.

Prior research works have looked at how different factors impact long-term retention of donors [1] or the project financial success [14, 16, 34]. The likelihood of funding has also been shown to be influenced by the language used in project descriptions [22], the gender of teachers [26], the social media activity associated with the project [20], or the amount of early donations [29]. However, our work highlights that the way a platform presents its users with choices will have significant impact on their donation decisions.

We, therefore, advocate for a careful investigation of choice architectures and the ethics of their impact on guiding user decisions. What we focus on in this paper more specifically, is how the choice architecture of the website might nudge donors towards choosing projects with certain attributes, and as a result, contribute to inequality in donations along the dimension of these attributes.

As our first contribution, we quantify the inequality in project funding and attribute portions of the inequality to different project characteristics (Note that in highly empirical applications, it is difficult to reason about causal effects without making strong assumptions about the data generating process, which limits the practicality of causal estimation approaches). Our approach to this task emerges from its similarity to the problem of quantifying income inequality, specifically to measure the contribution of characteristics such as gender or race to wage inequality. This problem has been studied in economics using decomposable inequality indices [4, 23], whose *additive decomposability property* helps factor out the influence of individual differences. We borrow one of such measures called ‘Generalized Entropy’ to find the contribution of different project characteristics to the overall inequality in funding.

Our second contribution revolves around *understanding how platform design choices affect donor behavior*. This idea is related to the existing literature on nudges [27, 32] and choice architectures [19, 28, 33] in behavioral economics. Research in these areas investigates how people’s choices depend on how information is presented to them. This framing applies to our scenario: we are interested in learning whether the way a donation platform presents project information influences how donors choose projects with specific attributes. We face two challenges when attempting this goal.

First, once we pick 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. Second, while it is feasible to imagine running active experiments (i.e., A/B testing) by the platform, such an approach raises a lot of ethical concerns in a setting where people’s decisions are nudged, and the outcomes influence future real-world financial success of others [3]. To overcome these issues, we analyze longitudinal changes in DonorsChoose platform design, and we estimate the average impact of a design change by comparing the observations right before and after it was first introduced. This way, we are able to quantify the influence of individual website design elements on project selection rates from historical donation data, while avoiding manipulation of future donations.

We run the inequality analysis on the data released by DonorsChoose in 2016, which provides over ten years worth of project and donation information, revealing two interesting findings. First, school location (zip code) is a factor largely contributing to funding inequality. Further analysis confirms that this is because (i) the majority of money is donated within-state (i.e., donors tend to give to their local projects), and (ii) the bulk of donations come from a handful of wealthier states, such as California or New York. Second, we find that the poverty level of the school does not contribute much to funding inequality – projects from high-poverty facilities have equal success rates as those from low-poverty ones, despite research showing that people consciously prefer charitable giving to be need-based [10].




Figure 1: Number of projects posted per year (along with the number of projects successfully funded).

Our longitudinal analysis shows that it is possible to enhance the impact of school location or poverty level in influencing the project selection rates, via changes in the platform design. In the process, we highlight that even though donation platforms are perceived as neutral facilitators of interactions between the donors and the requestors, they do have an active role to play. We envision our work to pave the way for future studies aiming to employ nudges effectively to improve platform-specific goals, such as to maximize total amount of donations, or to ensure the social and ethical requirements of online altruism, such as bringing in fairness in the distribution of donations across projects.

2 DONORSCHOOSE.ORG DATASET

DonorsChoose.org is a popular donation platform where teachers from schools in the United States can request donations for their classroom projects. Such requests range from buying books or tools for specialized experiments, funding of excursions, all the way to replacing worn-down basic classroom equipment. Teachers do not receive money directly; DonorsChoose purchases all the requested resources and ships them directly to the schools [13]. Teachers on average request \$561.57, with the median request being \$413.67¹.

When choosing the projects, donors can select projects based on attributes such as the school location (zip code, city, state), main focus 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 etc.). However, if a project does not raise its target amount by a certain deadline, it receives no funding, and the donors need to reallocate the donated money to other projects. While implementing this ‘all-or-nothing’ rule has been shown to have certain positive effects on donations, like increasing the overall amount of donations [35], there are projects that fail to receive funding as a result.

To identify the magnitude of such unsuccessful projects, in this work, we analyze the dataset publicly released by DonorsChoose in 2016 [12], which contains information about all projects posted on the platform starting from 2002 until the dataset’s release date. During this period, DonorsChoose helped raise over \$500M from over 5M donors. There were over 1M projects created by around 500K teachers from over 70K different schools. The dataset contains many project attributes, such as details about the teacher, requested

¹The median amount is not round because of factors such as sales tax.




Figure 2: CDF of project success rates across different zip codes. The distribution is highly unequal. Around 26% zip codes have 100% success rate (shown as the dashed line), whereas the success rate is much lower in other zip codes.

resources, and the textual description of the request. The dataset also contains donation data, including donation date, amount donated, and location of the donor (state and city; while the zip codes of donors are anonymized for privacy reasons). Since the dataset covers only a part of 2002 and 2016, we remove these two years from our analysis and focus on projects posted from 2003 to 2015.

Figure 1 shows that the number of projects on the platform has been steadily growing, yet a significant number of projects fail to get funded each year. The data reveals that on average 29.79% of projects fail every year. These failure rates are unlikely to be the result of some projects being low-quality or fraudulent since all the projects posted on the website are vetted and verified by DonorsChoose. The platform also decided to fund all active projects after getting Ripple’s generous \$29M donation in early 2018 [6]. We can thus conclude that all projects on DonorsChoose are worthy of getting funded, and it is important to understand the factors contributing to the inequality of success rates across projects, to ensure that the failure rates do not adversely affect certain types of projects or schools. This problem is the focus of the next section.

3 MEASURING INEQUALITY IN DONATIONS

In this section, we propose a methodology based on the well-studied concept of inequality measurement in economics, to identify whether a given feature (or attribute) contributes to inequality in projects’ chances of success. The proposed method allows us to quantify the inequality contributions *without the need for making any assumptions about the data generating process*.

For a fixed choice architecture, let features X_1, \dots, X_k characterize a project, and Y denote its success (with $Y = 1$ meaning that the project is successfully funded). We have access to observational data of the form $D = \{(x_1^i, \dots, x_k^i), y_i\}_{i=1}^N$ sampled i.i.d. from a joint distribution F . Then, we can utilize D to decide whether feature X_k contributes to the inequality across Y values. Note that we use upper case letters to denote random variables and lower case letters to denote in-sample realizations.

3.1 Decomposable Measures of Inequality

We propose to measure inequality using an *additively decomposable index of inequality* (namely, *generalized entropy*) for the vector




Figure 3: Distribution of project success rates across 4 poverty levels. The distribution is relatively equal.

(y_1, \dots, y_N) . We then estimate the impact of a feature X_k by decomposing the total inequality according to X_k feature values.

We begin with an overview of inequality measurement. Measures of inequality quantify the degree of inequality in a distribution of resources or benefits over a population. More precisely, given a distribution/vector $\mathbf{y} = (y_1, \dots, y_N) \in \mathbb{R}_{\geq 0}^n$, an inequality measure, $I : \bigcup_{n=1}^{\infty} \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$, maps any distribution/vector \mathbf{y} to a non-negative real number $I(\mathbf{y})$. A distribution \mathbf{y} is considered more equal than \mathbf{y}' if and only if $I(\mathbf{y}) < I(\mathbf{y}')$. A long established line of work in economics offers an axiomatic characterization for inequality indices. Virtually all commonly-used inequality indices satisfy axioms such as Anonymity, Population invariance, Transfer principle, Zero-Normalization [30].

Additive decomposability is a structural property of some inequality measures requiring that for *any* partition G of the population into groups, the measure, $I(\mathbf{y})$, can be expressed as the sum of a “within-group component” $I_{\omega}^G(\mathbf{y})$ (a weighted sum of subgroup inequality levels) and a “between-group component” $I_{\beta}^G(\mathbf{y})$ (computed by assigning to each element in a subgroup $g \in G$ the subgroup’s mean benefit μ_g):²

$$I(\mathbf{y}) = I_{\omega}(\mathbf{y}) + I_{\beta}(\mathbf{y}).$$

Following [30], we focus on a family of inequality indices called *generalized entropy indices* that satisfy additive decomposability. For a constant $\alpha \neq 0, 1$, the generalized entropy is defined as follows:

$$\mathcal{E}^{\alpha}(y_1, y_2, \dots, y_n) = \frac{1}{n\alpha(\alpha - 1)} \sum_{i=1}^n \left[\left(\frac{y_i}{\mu} \right)^{\alpha} - 1 \right]. \quad (1)$$

One can interpret generalized entropy as a measure of information theoretic *redundancy* in data. Given a partition G of the population, generalized entropy can be additively decomposed: Suppose subgroup $g \in G$ consists of n_g individuals with the benefit vector $\mathbf{y}^g = (y_1^g, \dots, y_{n_g}^g)$ and mean benefit μ_g . Generalized Entropy can be decomposed as follows:

$$\begin{aligned} \mathcal{E}^{\alpha}(y_1, y_2, \dots, y_n) &= \sum_{g=1}^{|G|} \frac{n_g}{n} \left(\frac{\mu_g}{\mu} \right)^{\alpha} \mathcal{E}^{\alpha}(\mathbf{y}^g) \\ &\quad + \sum_{g=1}^{|G|} \frac{n_g}{n\alpha(\alpha - 1)} \left[\left(\frac{\mu_g}{\mu} \right)^{\alpha} - 1 \right] \end{aligned}$$

²When the partition G in reference to is clear from the context, we drop the superscript G to simplify notation.

Feature	School Location	Amount Requested	Matching Donations	Resource Type	School Type	Primary Focus Subj.	Secondary Focus Subject	Secondary Focus Area	Primary Focus Area	Teacher Prefix	Poverty Level	Teacher Cred.	Grade Level
Normalized Inequality Contribution	0.1372	0.07198	0.02048	0.01399	0.0034	0.00307	0.00291	0.00208	0.0016	0.00108	0.00093	0.00087	0.00058

Table 1: The contribution of different features to inequality in donations. School location has the highest contribution, whereas poverty level has low contribution.

3.2 Measuring Feature Contribution through Inequality Decomposition

We propose estimating the impact of a feature X_k on Y through the following procedure:

- (1) Partition the population (i.e., the set of all projects) into groups G , according to the value of the feature X_k . For instance, suppose X_k represents a teacher’s gender, taking two different values. G partitions the projects into two groups, those with male teachers and those with female teachers.
- (2) Compute and return $I_{\beta}^G(y)/I(y)$ as the impact of X_k on Y .

Note that just because there is a higher between-group component in decomposition according to feature X_k when compared to the decomposition according to X_j does not necessarily mean that feature X_k has more *causal* impact on Y when compared to X_j . Why are we focusing on inequality decomposition then?

3.3 Relationship with Causal Inference

Traditionally, in order to understand the impact of a set of explanatory variables (in our case, X_1, \dots, X_K) on a target variable of interest (in our case, Y), econometricians begin by specifying an “explanatory model” (e.g., a simple regression model). Such a model carefully specifies a counterfactual in order to examine the influence of each supposedly causal factor on the target. In the process of developing such a model, econometricians must make several important choices, including which explanatory variables belong in the model, what functional form should the regression function take (i.e. logarithms, quadratic, etc.), and what the error distribution looks like. These choices are usually guided by domain expertise. For instance, consider a simple regression (which is one of the most commonly used methods in empirical studies): We assume $Y = U + \beta_0 + \sum_{k=1}^K \beta_k X_k$ where U is the error/noise term. Provided that this functional form is correctly specified, and that standard assumptions such as exogenous covariates hold, we can use OLS methods to estimate the coefficients. These estimates then serve as the measure of the causal impact of each attribute on the target variable: the parameter β_k measures the increase in Y attributable to one unit of change in X_k , keeping all the other covariates fixed.

Unfortunately, the above approach is restrictive in many empirical applications. First and foremost, this is because the modeling step can be cumbersome, and conclusions made through explanatory models are often sensitive to model specification [9]. Given these issues, here we base our analysis on an alternative tool, namely, measurement of inequality. [9] have indeed established a strong connection between the two methods under certain conditions. Assuming that the model is correctly specified (e.g., the data is indeed generated through the assumed linear model), they show

that the impacts estimated by OLS and those derived by inequality decomposition are closely related. For the precise expressions, see equation (22) in [9]. In summary, our decomposition approach allows us to measure feature influence without the need for strong assumptions on the data generating process, and is better suited for discrete variables.

4 IMPACT OF CHOICE ARCHITECTURE ON DONATION INEQUALITY

Using the proposed methodology, in this section, we utilize observational data from DonorsChoose to identify the features contributing to inequality in projects’ chances of success.

4.1 What’s Behind Inequality in Donations?

Table 1 illustrates our proposal on the DonorsChoose data corresponding to the year 2015³. Note that the contributions of school location (in terms of zip codes), amount requested, resource type (the type of resource that was requested), and matching donations (whether the amount donated by individual donors to a project are going to be matched by larger corporate donation) are deemed high, whereas the contribution of poverty level is low. Throughout the rest of this section, we narrow down our attention to two of these features: school location and poverty level.

Figure 2 shows the CDF of project success rates across different zip codes. As evident from the figure, zip code has a high contribution to inequality. For around 26% of the zip codes, all the projects were funded, whereas for other zip codes, the percentage of funded projects varied anywhere between 50% to 90%. Figure 3 shows the distribution of project success across different poverty levels (schools in DonorsChoose are classified into four poverty categories: Highest, High, Moderate, and Low, depending on how many students come from low-income families). As evident from the figure, the project success rates are similar across different poverty levels, explaining the low value of inequality contribution for this feature.

4.2 Behavioral Mechanisms Underlying the Inequality Contributions

We suspect that the high inequality contribution of school location is in part due to the donors’ preference to donate locally, that is, to donate to schools in their own zip code, city or state. The dataset does not contain donors’ zipcode (due to privacy considerations), but our analysis at the city and state level indeed reveals a strong

³As mentioned earlier, the DonorsChoose dataset consists of project funding information of multiple years. To exclude external factors such as the amount of traffic to the website, large fundraising activities in specific years, or varied amounts of matching donations from our analysis, we focus on the projects posted on a particular year.




Figure 4: Percentage of donations going to different school states. There is huge disparity in the donation distribution.




Figure 5: Fraction of schools belonging to highest poverty level in different states. Donation distribution in Figure 4 does not correlate with the school poverty level distribution.

tendency of donors to fund local projects: almost 80% of the total amount of funding is donated within the same state, and around 20% of the total amount is donated within the same city as the donor’s. Only 20% of the total donations are made to schools outside donors’ states. This tendency, combined with the fact that the amount of total donation varies widely across states (see Figure 4), explains the high inequality contribution of the school location. As we will discuss in Section 4.3, the locality bias might be further encouraged by the presence of certain nudges on the website (e.g., the option to search projects by location). This is despite the fact that the donation distribution in presence of the location bias does not correlate with the distribution of highest poverty schools in different states (see Figure 5), which are in higher need of economic support.

From a moral standpoint and to alleviate current social injustice, one may argue that the schools with higher economic need are more deserving of receiving donations, and the website should nudge the donations in their favor. To put it more formally, it may be desirable (ethically and socially) to reduce the contribution of location, but increase the contribution of school’s poverty level to how donations are distributed. Next, we show evidence that choice architecture and nudges indeed impact contribution of these factors.

4.3 Evidence on the Impact of Nudges

To gain insights on the impact of choice architecture and nudges on inequality in different projects’ chances of success, we identify the time points, throughout the years, when a new website design

was introduced in DonorsChoose, using snapshots from ‘The Internet Archive’ (web.archive.org). Then, utilizing the longitudinal donation data, we analyze the effect of these interventions.

To this end, prior works have proposed *Regression Discontinuity Designs* (RDDs) [2, 11, 17, 21]—a quasi-experimental design that reasons about the causal effects of interventions by identifying a threshold above or below which an intervention is assigned. We follow this methodology by comparing the projects *before* and *after* a new design feature is introduced. (The intuition is that other conditions (e.g., the project and donor pool) are likely to be very similar right before and after each intervention.) More specifically, let (x_1, \dots, x_k) specify the features that impact a project’s success (e.g., *school location* or *poverty level*) and t_0 the timestamp when a new design feature is introduced. Let y specify whether a project is successful or not. We fit the following model to the data consisting of projects introduced right before and right after the design change of interest:

$$y = a_0 + a_1x_1 + \dots + a_kx_k + I(t > t_0)(b_0 + b_1x_1 + \dots + b_kx_k). \quad (2)$$

The estimated coefficients b_1, \dots, b_k specify if the intervention has impacted the importance of feature x_1, \dots, x_k on the success of a project. Unfortunately for our data set, linear regression is not a good fit—making the conclusions from the RDD design unreliable.

To circumvent the issue, we measure the impact of each design change on donors’ behavior via another quasi-experimental design: we estimate the average impact of the change by comparing the




Figure 6: Changes in the inequality contribution of school location and poverty level before and after certain design changes were introduced in DonorsChoose website.

inequality contribution of different features right before and after it was first introduced. Again, the intuition is that other conditions are likely to be very similar right before and after each intervention. Therefore, comparing the inequality contribution after the intervention (treatment group) to that before the treatment (control group) will deliver the local treatment effect.

Next, we look at several nudges directly related to the school location, introduced in DonorsChoose website at different times:

- **2004:** Location filter was added to the homepage and the project search page. We expect this to increase the contribution of school location. Figure 6a confirms this.
- **2012:** Added ‘projects near me’ filter to the homepage. We expect this to increase the contribution of school location. Figure 6a further confirms this.
- **2015:** Location search box was added to both home and search pages. We expect this to increase the contribution of school location, and this is indeed consistent with what we observe in Figure 6a.

We now look at several nudges directly related to the school poverty level:

- **2005:** Filter “High level of poverty” was added. We expect this to increase the contribution of poverty level to inequality (i.e., poor schools now have more chance of getting funds). This is what we observe in Figure 6b.
- **2008:** “Most Urgent” criteria became the default ranking of search results (urgency = highest poverty + lowest cost + fewest days left). We expect this to increase the contribution

of poverty level. This is consistent with our observation in Figure 6b.

- **2012:** “Most Urgent” criteria was redefined (to highest poverty + closest to finish line). We expect this to increase the contribution of poverty level. This is consistent with what we observe in Figure 6b.

These results show the potential of choice architecture to impact donors’ behavior, and we believe that by carefully designing the nudges, the platform can influence the selection rates of projects of different poverty levels and locations in a socially desirable way.

5 CONCLUSION AND FUTURE DIRECTIONS

Recently, there has been a lot of focus on biases in human behavior [7, 22, 26] as well as in algorithmic models [5, 8, 15], and their impacts on different socio-technical systems. Our work sheds light on the power of choice architecture and nudges in shaping the user behavior in online platforms, thus contributing to the observed biases. We conducted an in-depth analysis of the donation platform – DonorsChoose.org, by utilizing longitudinal empirical data. In this process, our work takes an important first step towards adapting nudge frameworks to the social and ethical requirements of online altruism. Many interesting directions remain open for future work:

How can donations be fairly distributed across projects?

To answer this question, we need to first carefully specify a measure of *need* for different projects, which may depend on factors such as the poverty level of the school district, the funding target, project urgency, and others. Given an inter-project-comparable measure of need, one can define fairness as follows: projects with the same level of need should have a similar likelihood of success. Our findings illustrate that the choice architecture can influence project selection rates, and the platform can potentially explore different ways to achieve a fair distribution of donations using nudges.

Ethics of nudges

There has already been much debate about whether and when the use of nudges is ethical [18]. Some refer to nudges as a form of paternalism and restriction of choice and agency, whereas others refer to it as libertarian paternalism – a way of positively influencing people without limiting their freedom of choice [31]. We believe such a discussion should be informed by the context in which nudges are employed. While we did not attempt to tackle these issues in this paper, there are deep ethical concerns that arise when discussing nudges in the context of online charity platforms: For instance, does the goal of promoting more deserving (but less attractive) projects justify the use of the nudges? To what extent? We leave these and many other questions for future work.

Acknowledgments: 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] Tim Althoff and Jure Leskovec. 2015. Donor retention in online crowdfunding communities: A case study of donorschoose.org. In *WWW*. 34–44.
- [2] Pablo Aragón, Vicenç Gómez, and Andreas Kaltenbrunner. 2017. Detecting platform effects in online discussions. *Policy & Internet* 9, 4 (2017), 420–443.

- [3] 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).
- [4] Alan S Blinder. 1973. Wage discrimination: reduced form and structural estimates. *Journal of Human resources* (1973), 436–455.
- [5] Engin Bozdag. 2013. Bias in algorithmic filtering and personalization. *Ethics and information technology* 15, 3 (2013), 209–227.
- [6] Stephen Burke. 2018. #BestSchoolDay 2018: Every Project Funded! <https://www.donorschoose.org/blog/best-school-day-2018>.
- [7] Abhijnan Chakraborty, Saptarshi Ghosh, Niloy Ganguly, and Krishna P Gummadi. 2016. Dissemination biases of social media channels: On the topical coverage of socially shared news. In *ICWSM*. AAAI.
- [8] Abhijnan Chakraborty, Johnatan Messias, Fabricio Benevenuto, Saptarshi Ghosh, Niloy Ganguly, and Krishna P Gummadi. 2017. Who makes trends? understanding demographic biases in crowdsourced recommendations. In *ICWSM*. AAAI.
- [9] Frank A Cowell and Carlo V Fiorio. 2011. Inequality decompositions – a reconciliation. *The Journal of Economic Inequality* 9, 4 (2011), 509–528.
- [10] 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.
- [11] Sayamindu Dasgupta and Benjamin Mako Hill. 2018. How "wide walls" can increase engagement: evidence from a natural experiment in Scratch. In *CHI*. ACM, 361.
- [12] Donorschoose.org Data. 2016. Research for Education. <https://research.donorschoose.org/t/download-opendata/33>.
- [13] Donorschoose.org. 2019. About US. <https://www.donorschoose.org/about>.
- [14] Vincent Etter, Matthias Grossglauser, and Patrick Thiran. 2013. Launch hard or go home!: predicting the success of kickstarter campaigns. In *COSN*. ACM, 177–182.
- [15] Megan Garcia. 2016. Racist in the machine: The disturbing implications of algorithmic bias. *World Policy Journal* 33, 4 (2016), 111–117.
- [16] Michael D Greenberg, Bryan Pardo, Karthic Hariharan, and Elizabeth Gerber. 2013. Crowdfunding support tools: predicting success & failure. In *CHI Extended Abstracts*. ACM, 1815–1820.
- [17] Scott A Hale, Peter John, Helen Margetts, and Taha Yasseri. 2018. How digital design shapes political participation: A natural experiment with social information. *PloS one* 13, 4 (2018), e0196068.
- [18] Daniel M Hausman and Brynn Welch. 2010. Debate: To nudge or not to nudge. *Journal of Political Philosophy* 18, 1 (2010).
- [19] Eric J Johnson, Suzanne B Shu, Benedict GC Dellaert, Craig Fox, Daniel G Goldstein, Gerald Häubl, Richard P Larrick, John W Payne, Ellen Peters, David Schkade, et al. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504.
- [20] 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.
- [21] Momin M Malik and Jürgen Pfeffer. 2016. Identifying platform effects in social media data. In *ICWSM*. AAAI.
- [22] Tanushree Mitra and Eric Gilbert. 2014. The language that gets people to give: Phrases that predict success on kickstarter. In *CSCW*. ACM, 49–61.
- [23] Ronald Oaxaca. 1973. Male-female wage differentials in urban labor markets. *International economic review* (1973), 693–709.
- [24] U.S. Department of Education. 2017. The Federal Role in Education. <https://www2.ed.gov/about/overview/fed/role.html>.
- [25] Chankyung Pak and Rick Wash. 2017. The Rich Get Richer? Limited Learning in Charitable Giving on donorschoose. org. In *ICWSM*. AAAI, 172–181.
- [26] J Radford. 2016. The emergence of gender inequality in a crowdfunding market: an experimental test of gender system theory. (2016).
- [27] 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.
- [28] Itamar Simonson. 1989. Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research* 16, 2 (1989), 158–174.
- [29] Jacob Solomon, Wenjuan Ma, and Rick Wash. 2015. Don't wait!: How timing affects coordination of crowdfunding donations. In *CSCW*. ACM, 547–556.
- [30] Till Speicher, Hoda Heidari, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. A Unified Approach to Quantifying Algorithmic Unfairness: Measuring Individual and Group Unfairness via Inequality Indices. In *KDD*. ACM.
- [31] Cass R Sunstein and Richard H Thaler. 2003. Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review* (2003), 1159–1202.
- [32] Richard H Thaler and Cass R Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
- [33] Richard H Thaler, Cass R Sunstein, and John P Balz. 2014. Choice architecture. (2014).
- [34] 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).
- [35] Rick Wash and Jacob Solomon. 2014. Coordinating donors on crowdfunding websites. In *CSCW*. ACM, 38–48.