Moral Framing and Online Fundraising Outcomes: Evidence from GoFundMe Campaigns

Ji Eun Kim[1], **Libby Hemphill**[1]

[1]School of Information, University of Michigan
jieunkim@umich.edu, libbyh@umich.edu

Abstract

This study examines the relationship between moral framing and fundraising outcomes, encompassing both monetary and social support, through an analysis of 14,088 GoFundMe campaigns. Focusing on three moral foundations (care, fairness, and ingroup loyalty), we quantified the presence of these frames within campaign appeals. Our results show that the association between moral framing and campaign success varies significantly across fundraising categories. Specifically, negatively framed appeals emphasizing the harm experienced by the help-seeker were associated with a higher volume of donations and comments in the Emergency and Memorial categories, yet related to fewer donations in animal-related campaigns. Highlighting harm was not associated with a higher average donation amount per donor. Ingroup loyalty framing was positively associated with both donation volume and supporter engagement. This research extends the literature on moral communication in crowdfunding and offers practical implications for the design of online fundraising platforms to better facilitate prosocial interactions between fundraisers and supporters.

Introduction

Online fundraising platforms have become a vital resource for help-seekers, enabling them to communicate with supporters and raise funding for diverse needs. However, while these platforms host a vast number of campaigns, attention and financial support remain highly skewed (Chakraborty et al. 2019). The majority of campaigns fail to meet their goals; in our GoFundMe dataset, only 22% of campaigns reached their goals. This underscores the significant challenge fundraisers face in capturing donor interest. Hence, identifying the factors that distinguish successful appeals from unsuccessful ones offers critical insights for both fundraisers and platform designers on how to frame and communicate campaigns to improve success rates.

This study empirically examines how the moral framing of fundraising appeals relates to both monetary and social support outcomes within the context of GoFundMe, a leading crowdfunding platform. GoFundMe serves as a space for individuals to seek assistance for various personal and social challenges (see Figure 1). The platform processed over 42 million donations to individuals and 23 million contributions to non-profit organizations in 2024 (GoFundMe 2025), with an average donation rate of two per second and a mean contribution of $77. On this platform, fundraisers set a fundraising goal and post a story along with photos to articulate their needs. Donors may contribute financially, either publicly or anonymously, and provide social support by posting comments. Using a dataset of 14,088 GoFundMe campaigns, we analyze the associations between moral framing and three key performance metrics: the total number of donations, the average donation amount per donor, and the volume of comments.

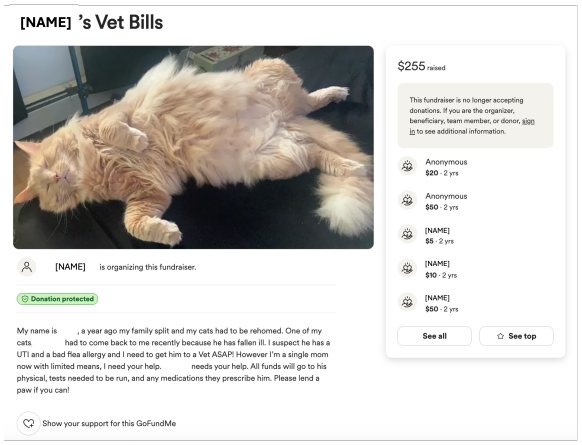

Figure 1: Example GoFundMe campaign. Fundraisers provide background information and explain the purpose of their campaign by uploading a narrative story accompanied by photos. Donors may contribute financially and leave supportive messages. Both fundraiser and donor identities have been removed from the screenshot.

Our linear regression analyses indicate that emphasizing the harm experienced by the help-seeker is associated with a greater volume of donations and comments across various categories; however, these negatively framed appeals do not show a significant relationship with higher average donation amounts. In contrast, positively framed appeals highlighting ingroup loyalty are linked to both an increase in donations and comments as well as larger average donation amounts

per donor.

This study provides both theoretical and practical contributions. Theoretically, our findings advance the literature on donation behavior and dynamics by examining the relationships between the moral framing of campaign appeals and both fundraising and social support outcomes across diverse fundraising categories. Practically, these insights can inform the development of support tools that assist fundraisers in crafting effective, context-specific appeals tailored to their unique goals and situations

Related Work

Determinants of Successful Online Fundraising

According to Hoover et al. (2018), the factors associated with the success of charitable donations can be categorized into three primary types: individual differences, situational characteristics, and solicitation framing. While existing research has established that the first two categories are significantly linked to fundraising outcomes, this research focuses on the relationship between solicitation framing and campaign success, particularly within the context of text-based appeals.

First, individual differences relate to donor characteristics and the shared group identities between givers and recipients that influence donation decisions. For example, on GoFundMe, both men and women are more likely to donate when the majority of existing supporters are of the opposite sex; furthermore, female supporters often express greater empathy in their support messages (Sisco and Weber 2019). Recipient characteristics also play a role, as donors tend to contribute more to those with the same last names (Sisco and Weber 2019) and show a preference for supporting ingroup members over outgroup individuals (Nilsson, Erlandsson, and Västfjäll 2016). Additionally, female-led campaigns generally reach their fundraising goals more quickly (Esperança et al. 2025).

The second category comprises situational characteristics, which include factors such as social pressure and reputation-seeking. Research has shown that individuals engage in helping behaviors to achieve a "warm glow" from altruism or to avoid the guilt associated with inaction (Cunningham, Steinberg, and Grev 1980). These motivations are linked differently to situational factors, such as the level of personal responsibility. For instance, individuals are more sensitive to feelings of guilt in high-responsibility situations, as they may feel a greater sense of culpability for failing to fulfill a duty uniquely assigned to them. Conversely, in low-responsibility contexts, they tend to experience a warm glow, potentially because the act of helping is perceived as more voluntary and exceeding their baseline moral obligations (Erlandsson, Jungstrand, and Västfjäll 2016). Other studies suggest that people seek to enhance their reputation through prosocial actions. For example, individual donation amounts tend to increase when the average donation amount is visible (Sisco and Weber 2019), and people may emulate the behavior of high-status donors to associate themselves with a higher social rank (Kumru and Vesterlund 2010).

Finally, solicitation framing involves the strategic construction of appeals to encourage donations and social support. Analyzing these strategies requires examining both visual and textual information. Some researchers have focused on image-based features, analyzing photos within campaign appeals to measure their association with fundraising performance. For example, Rhue and Robert (2018) found that campaigns featuring happy facial expressions are associated with higher donation counts, whereas neutral expressions are linked to fewer donations. Another study has examined the textual information that contributes to campaign success; for instance, longer campaign appeals are associated with faster goal attainment (Esperança et al. 2025). In addition, the use of "netspeak" in appeals is positively associated with fundraising outcomes, while excessive use of second-person pronouns is linked to less favorable results (Jang and Chu 2022).

The Framing of Fundraising Appeals

Framing serves as an effective tool for persuasion, with various strategies being adopted to encourage charitable giving (Wang et al. 2019). Strategic framing is particularly vital in the context of fundraising, as seeking financial assistance can carry a social stigma and may lead to criticism (Parsell and Clarke 2022). This stigma often creates barriers for vulnerable groups when articulating their needs. Consequently, existing research has explored how fundraisers construct their narratives to justify requests while maintaining dignity and reducing negative impressions. For instance, Radu and McManus (2018) found that survivors of intimate partner violence meticulously frame their stories to project a positive self-image, thereby counteracting the stigma associated with financial dependency. Their strategies include situating personal experiences within the broader issue of domestic violence to invoke a sense of social justice or including imagery of children to validate the necessity of external support.

Existing research on fundraising appeals identifies emotional valence as a primary component of framing. Prior studies suggest that both positive and negative frames can be advantageous; specifically, positively framed appeals are associated with a warm glow effect, whereas negatively framed appeals are linked to feelings of guilt (Erlandsson, Nilsson, and Västfjäll 2018). However, empirical findings remain nuanced and occasionally contradictory, as the relationships between emotions and fundraising outcomes appear to vary based on the specific context and goals of the campaign. Some research underscores the efficacy of positive appeals; for instance, positive sentiments in campaigns promoting social bonding are associated with higher donation counts (Paxton, Velasco, and Ressler 2020). Similarly, appeals emphasizing positive moral values, such as care and loyalty, are linked to increased giving (Hoover et al. 2018). Conversely, an alternative line of research indicates that negative framing may be more effective in encouraging certain helping behaviors. For example, prior research highlights the role of anger in philanthropic behavior, as donors may be more inclined to contribute to restorative charities that seek to remedy the harm inflicted upon a victim (Blader and Tyler

2002; van Doorn, Zeelenberg, and Breugelmans 2017). Furthermore, some evidence suggests that joyful language is associated with smaller individual donation amounts (Rhue and Robert 2018). Jang and Chu (2022) found that while negativity is linked to a higher volume of donors and increased social media engagement, it is associated with a decrease in the average amount contributed per donor.

The varied findings between appeal sentiment and fundraising success, where positive and negative frames appear effective in different contexts, suggest that additional dimensions of an appeal or its specific situation relate to campaign outcomes. Identifying the structural elements of an appeal, beyond simple sentiment, may assist fundraisers in designing more effective campaigns. Consequently, we incorporate moral frames into our analysis to capture these additional dimensions of framing strategies. This approach provides a more comprehensive account of how specific moral narratives are linked to donor engagement.

Moral Foundations Theory (MFT)

We adopt three moral frames derived from Moral Foundations Theory (MFT): care, fairness, and loyalty. The original MFT framework identifies five moral foundations (care-harm, fairness-cheating, loyalty-betrayal, authority-subversion, and sanctity-degradation) to explain variations in moral judgment and development (Graham et al. 2013; Haidt and Graham 2007). These foundations have been widely utilized to analyze how moral judgments influence various social outcomes and have been applied to investigate their relationships with charitable giving (Nilsson, Erlandsson, and Västfjäll 2016; Hoover et al. 2018).

The moral foundation of *care* involves the protection of individuals from harm, while *fairness* pertains to the human need for equitable treatment. The *loyalty* foundation relates to group cohesion and the sense of belonging. The *authority* foundation concerns the respect for hierarchy and tradition, and *sanctity* involves avoiding contamination and pursing spiritual purity. Among these five foundations, we selected care, fairness, and loyalty for this analysis for two reasons. First, previous research suggests that these three frames may influence fundraising outcomes (Blader and Tyler 2002; Hoover et al. 2018; Linos, Jakli, and Carlson 2021; Nilsson, Erlandsson, and Västfjäll 2016; van Doorn, Zeelenberg, and Breugelmans 2017). Second, because authority is conceptually tied to hierarchical relationships and sanctity is often specific to religious contexts, these two frames may not be well-suited for analyzing online fundraising.

Research Questions

We aim to address the following gaps in the current literature. First, existing studies have yielded mixed findings, as the associations between the linguistic features of campaign appeals and fundraising success appear to vary by campaign purpose and need. Therefore, we include multiple fundraising categories in our analysis to investigate how moral framing relates to campaign outcomes across these distinct contexts.

previous research has often left a significant portion of the variance in appeal success unexplained by relying on a single binary measure of sentiment to capture emotional content. To address this limitation, we introduce moral frames to provide a more nuanced dimensionality to campaign appeals. We utilize three binary valence measures for our moral axes: care-harm, fairness-cheating, and loyalty-betrayal. We investigate how the moral framing of appeals is associated with three primary outcomes: (1) the total number of donations, (2) the average donation amount per donor, and (3) the volume of supporter comments.

Our first research question examines the positioning of campaign appeals across these three moral dimensions within four fundraising categories and compares these appeals based on donation outcomes. In light of the existing literature, we anticipate that moral frames in the negative domain (e.g., harm, unfairness, outgroup) will be associated with a higher volume of donations but smaller average amounts per donor. Our second research question examines whether the volume and content of comments vary according to the moral framing of the campaign appeal.

- **RQ1**: How does the moral framing of campaign appeals relate to donation outcomes, specifically the total number of donations and the average amount per donor?
- **RQ2**: How does the moral framing of campaign appeals relate to the quantity of social support, measured by the volume of supporter comments?

Data and Methods

Data

For our primary analysis, we utilized a publicly available dataset[1] comprising 14,859 unique GoFundMe campaigns (Xu, Li, and Zhou 2023). To ensure the generalizability of our findings across different contexts, we conducted supplemental analyses on a dataset from Kiva, a crowdfunded loan platform. These robustness checks yielded consistent results; we provide comprehensive details regarding the Kiva dataset and the corresponding analysis in the Appendix.

To safeguard user privacy, all results are presented in aggregate; furthermore, any representative quotes used for illustrative purposes have been paraphrased. While the original dataset includes campaigns from 13 different countries, the vast majority were initiated by U.S.-based fundraisers. To ensure linguistic and contextual consistency, we excluded campaigns from non-U.S. locations and those with non-English appeals, identifying the latter using the *langdetect* library. This filtering process resulted in a final sample of 14,088 campaigns.

The fundraising campaigns in our final GoFundMe dataset were initially classified into five primary categories: Memorial (30.6%), Medical (23.7%), Animal (19.5%), Emergency (16.9%), and Financial Emergency (9.3%). We subsequently merged "Emergency" and "Financial Emergency" into a single category (Emergency), resulting in four distinct groups for analysis. Memorial campaigns primarily raise funds for funeral expenses and assist bereaved families; the Medical category includes projects aimed at covering healthcare costs; and the Animal category supports vet-

[1] https://zenodo.org/records/8287320

Moral domain	Bias	Term count	Example seed terms
Care	Vice	74	abuse, attack, damage, destroy, harm
	Virtue	57	benefit, care, defend, empathic, peace
Fairness	Vice	39	discriminate, dishonest, injustice, prejudice, unfair
	Virtue	38	equity, fair, honest, justice, reciprocal
Loyalty	Vice	29	deceive, enemy, foreign, immigrant, imposter
	Virtue	59	ally, community, family, fellow, group

Table 1: Moral frames adopted for the study and example seed terms representing each semantic construct.

Moral domain	Bias	Campaign appeal
Care	Vice	[NAME] and [NAME] were involved in a terrible car accident on [DATE]. They've got serious injuries. [NAME] fractured his neck, while [NAME] broke his leg. [...]
	Virtue	We protect animals' right to live and their access to medical care, food, water, and shelter.
Fairness	Vice	My friend [NAME] was attacked by a man and a woman in the [PLACE] parking lot near [PLACE]. They hurled racist slurs at her and physically assaulted her. [...]
	Virtue	Please help me support [NAME] and her kitties in getting back into their home. Any contribution would mean a lot. [...]
Loyalty	Vice	We provide shelter and medication to stray animals.
	Virtue	[NAME] served as a [OCCUPATION] in [PLACE]. His family is celebrating his 35 years of dedication and contributions to the community, a role he cherished deeply as he loved serving those around him. [...]

Table 2: Examples of GoFundMe campaign appeals with varying moral foundation scores. Note that excerpts have been paraphrased to protect user privacy.

erinary bills as well as the rescue and care of animals. The consolidated Emergency category focuses on supporting individuals who have injuries or property loss due to disasters such as fires or floods. Our dataset comprises 1,142,406 donation transactions and 52,550 comments. Approximately 22% of campaigns in our sample successfully reached their specified fundraising goals.

Measuring Moral Values Using Word Embeddings

We utilized the code from Mokhberian et al. (2020) to computationally construct moral frames. Mokhberian et al. (2020) employed FrameAxis (Kwak et al. 2021) to construct moral axes and compute bias scores for news headlines and articles on each axis. This method involves measuring the cosine similarity between document terms and the specified axis. FrameAxis has previously been applied to quantify moral expression in social media posts (Priniski et al. 2021) and to identify partisanship in political news (Mokhberian et al. 2020). Each axis consists of two opposing poles defined by antonym pairs; we used terms from the Moral Foundations Dictionary (Graham, Haidt, and Nosek 2009) to define each pole. For example, the "harm" pole of the harm-care axis consists of terms such as "abuse" and "damage," while the "care" pole is defined by words like "compassion" and "safe." Term counts and examples for each moral frame are provided in Table 1. To operationalize the harm-care axis (i.e., the care frame), we subtracted the mean vector of harm-related terms from the mean vector of care-related terms. All word vectors were obtained from the fastText model pretrained on the Common Crawl corpus (Mikolov et al. 2018).

Following the construction of the moral axes, we characterized each campaign appeal by measuring its specific bias toward these dimensions. The bias score for a campaign appeal on the harm-care axis, for instance, is calculated as the weighted average of the cosine similarities between the constituent words of the appeal and the axis. Consequently, an appeal containing numerous terms with high cosine similarity to the care pole is categorized as more relevant to care than to harm. In our setting, a positive bias score indicates that the appeal aligns more closely with the care pole, whereas a negative score suggests a higher concentration of harm-related language. Representative examples of campaign appeals with varying moral foundation scores are provided in Table 2.

	Animal	Emergency	Medical	Memorial
Number of donations	31.51 (39.34)	57.27 (88.31)	123.95 (135.11)	99.91 (127.14)
Average donation amount per donor	53.89 (34.62)	86.89 (54.71)	107.32 (55.77)	96.12 (50.50)
Number of comments	1.40 (2.93)	2.27 (4.79)	5.77 (9.09)	4.88 (7.46)
Care score	−0.01 (0.01)	0.00 (0.01)	−0.01 (0.01)	0.00 (0.01)
Fairness score	0.00 (0.01)	0.00 (0.01)	−0.01 (0.01)	−0.01 (0.01)
Loyalty score	−0.02 (0.01)	−0.01 (0.01)	−0.01 (0.01)	0.00 (0.01)
Campaign appeal length	225.73 (182.26)	218.38 (205.69)	285.76 (233.43)	159.16 (116.69)
Number of photos	0.74 (2.27)	0.36 (1.48)	0.39 (1.49)	0.23 (1.41)
Fundraising goal	4.7k (9.1k)	12.0k (29.1k)	327.9k (17.3M)	14.2k (19.0k)
Number of campaigns	2,746	3,693	3,335	4,314

Table 3: Descriptive statistics for the GoFundMe dataset. Values represent means, with standard deviations provided in parentheses.

Variables

The outcome variables in this study comprise the number of donations, the average donation amount per donor, and the number of comments. We excluded the total funds raised from this analysis due to its high correlation with the number of donations (Spearman's $\rho = 0.91$, $p < 0.001$). We conducted multiple linear regression analyses to examine the main and interaction effects of fundraising categories and moral foundation scores on each continuous outcome variable. The moral foundation scores for the three moral frames were calculated as described in the previous subsection.

We added several control variables to our models to account for factors potentially associated with fundraising and social support outcomes. First, we utilized VADER, a sentiment analysis tool (Hutto and Gilbert 2014), to determine the sentiment scores of the campaign appeals. Based on their VADER compound scores, appeals were categorized into three distinct groups: positive, neutral, or negative. Positive appeals accounted for 87.8% of the sample, while neutral and negative appeals represented 0.7% and 11.4%, respectively. We also measured the length of each appeal by counting tokens after splitting an appeal by whitespace. We further controlled for the number of photos included in each appeal and the stated fundraising goal amount. Descriptive statistics for all continuous variables used in this analysis are presented in Table 3.

Results

We conducted multiple regression analyses to address our research questions, ensuring that our dataset met the necessary assumptions for such models. For the categorical variables (fundraising category and campaign appeal sentiment), we used "Animal" and "Negative" as the reference categories, respectively. To account for skewed distributions, a log transformation was applied to the number of donations, the average donation amount per donor, the number of comments, campaign appeal length, the number of photos, and the fundraising goal amount. Table 4 presents the results for three regression models: the number of donations (Model 1), the average donation amount per donor (Model 2), and the number of comments (Model 3).

Regarding the control variables, both campaign appeal length and the fundraising goal amount were positively associated with all outcome variables. Positive sentiment, however, was positively associated only with the average donation amount per donor ($b = 0.032$, $p < 0.01$) and did not show a significant relationship with the number of donations or comments. Additionally, the number of donations was positively associated with the number of photos included in an appeal ($b = 0.037$, $p < 0.01$).

Moral Framing and Fundraising Outcomes

Model 1 in Table 4 details the associations between the three moral frames and the number of donations, including their interactions with fundraising categories. The main effect of care framing, which is the effect of care framing within the "Animal" reference category, is positive and statistically significant ($b = 2.570$, $p < 0.05$). However, in the Emergency ($b = -5.188$, $p < 0.001$) and Memorial ($b = -3.929$, $p < 0.01$) categories, higher care scores are associated with fewer donations, suggesting that appeals emphasizing the "harm" pole of the axis are linked to a higher volume of donations in these contexts. In contrast, appeals emphasizing ingroup loyalty are associated with increased donation counts. Higher loyalty scores relate to a greater number of donations in the Emergency ($b = 11.864$, $p < 0.001$), Medical ($b = 5.597$, $p < 0.01$) and Memorial ($b = 5.078$, $p < 0.01$) categories. Fairness framing does not show significant effects across categories, suggesting that emphasizing injustice or unfairness does not necessarily lead to a higher number of donations.

Model 2 in Table 4 examines the associations between the three moral frames and the average donation amount per donor, including their interactions across fundraising categories. Among these moral frames, only loyalty framing shows a significant positive association with the average do-

	Model 1		**Model 2**		**Model 3**	
Variable	Number of donations		Average donation amount		Number of comments	
Emergency	0.306***	(0.033)	0.390***	(0.021)	0.099*	(0.039)
Medical	0.881***	(0.035)	0.438***	(0.023)	0.592***	(0.041)
Memorial	0.714***	(0.032)	0.414***	(0.020)	0.547***	(0.037)
Care	2.570*	(1.079)	0.344	(0.690)	0.685	(1.266)
Fairness	−2.266	(1.400)	0.244	(0.895)	−0.433	(1.641)
Loyalty	−0.233	(1.318)	3.154***	(0.842)	−0.969	(1.545)
Emergency × Care	−5.188***	(1.396)	−1.049	(0.892)	−4.845**	(1.636)
Medical × Care	−1.678	(1.528)	0.340	(0.976)	−2.633	(1.791)
Memorial × Care	−3.929**	(1.433)	0.411	(0.916)	−3.917*	(1.680)
Emergency × Fairness	−2.411	(1.771)	0.969	(1.132)	0.798	(2.076)
Medical × Fairness	−0.635	(1.933)	−1.721	(1.235)	1.887	(2.266)
Memorial × Fairness	−3.033	(1.809)	−0.031	(1.156)	1.051	(2.121)
Emergency × Loyalty	11.864***	(1.617)	3.057**	(1.034)	6.034**	(1.896)
Medical × Loyalty	5.597**	(1.754)	−0.815	(1.121)	6.140**	(2.056)
Memorial × Loyalty	5.078**	(1.581)	−0.099	(1.010)	5.819**	(1.853)
Positive sentiment	0.006	(0.019)	0.032**	(0.012)	0.040	(0.023)
Neutral sentiment	−0.069	(0.070)	0.076	(0.045)	−0.091	(0.082)
Campaign appeal length	0.053***	(0.009)	0.031***	(0.006)	0.077***	(0.011)
Number of photos	0.037**	(0.013)	0.006	(0.008)	0.017	(0.015)
Fundraising goal	0.282***	(0.005)	0.121***	(0.003)	0.205***	(0.006)
Intercept	0.671***	(0.066)	2.793***	(0.042)	−1.498***	(0.077)
Number of observations	14,088		14,088		14,088	
Adjusted R^2	0.45		0.31		0.23	

Table 4: Multiple linear regression results for the number of donations (Model 1), the average donation amount per donor (Model 2), and the number of comments (Model 3). Care, Fairness, and Loyalty values represent moral foundation bias scores. Cell values represent coefficients, with standard errors in parentheses. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Moral domain	Bias	Comment
Care	Vice	We are terribly sorry for the heartbreaking loss.
	Virtue	Sending prayers for peace and comfort to your family.
Fairness	Vice	I stand in solidarity with you. The actions of [ORGANIZATION] were unacceptable and a serious misconduct.
	Virtue	I hope you recover quickly and can return to work soon.
Loyalty	Vice	I'm sorry that you and many other innocent minorities have experienced this.
	Virtue	Strength comes from unity, and when there is teamwork and collaboration, great things can be accomplished.

Table 5: Examples of supporter comments with varying moral foundation scores. Note that excerpts have been paraphrased to protect user privacy.

nation amount per donor within the Animal ($b = 3.154$, $p < 0.001$) and Emergency ($b = 3.057$, $p < 0.01$) categories. Care framing does not show a statistically significant association with the average donation amount per donor, suggesting that this dimension may relate more to the decision to donate than to the specific amount contributed.

A notable finding is that while negatively framed appeals highlighting harm were associated with a higher volume of donations, appeals with negative sentiment did not correspond to an increase in the average donation amount per donor. One potential explanation, as suggested by Jang and Chu (2022), is that people may seek to mitigate the negative feeling associated with a refusal to help. In such cases, donors might contribute the minimum amount necessary to alleviate guilt, resulting in a higher frequency of donations but lower individual contributions. This interpretation is supported by our observation that appeals with negative sentiment were associated with a significantly lower average donation amount per donor compared to those with positive sentiment (see Table 4).

Moral Framing and Social Support Outcomes

Model 3 in Table 4 reports the associations between the three moral frames and the volume of supporter comments, including their interactions across fundraising categories. Our results indicate that campaigns in the "Animal" reference category were associated with fewer comments compared to other fundraising categories ($b = -1.498$, $p < 0.001$). Negatively framed appeals emphasizing the "harm" pole of the harm-care axis were associated with a higher volume of support messages in the Emergency ($b = -4.845$, $p < 0.01$) and Memorial ($b = -3.917$, $p < 0.05$) categories. Furthermore, loyalty framing demonstrated a significant positive association with the number of comments in the Emergency ($b = 6.034$, $p < 0.01$), Medical ($b = 6.140$, $p < 0.01$), and Memorial ($b = 5.819$, $p < 0.01$) categories. These findings suggest that appeals centered on ingroup loyalty and collective identity consistently correspond to higher levels of social support.

Given that negatively framed appeals emphasizing harm are associated with a higher volume of support messages, we further investigated whether supporters provide more positive comments in response to these appeals to compensate for or dilute the negativity. Because the associations between moral framing and outcomes were most pronounced in the "Emergency" category, we segmented campaign appeals from this category into three groups based on the distribution of their moral foundation scores: low (more than one standard deviation below the mean), medium (within one standard deviation of the mean), and high (more than one standard deviation above the mean). We also calculated the moral foundation scores for all comments (see Table 5 for examples) and computed the average comment score for each campaign appeal. Figure 2 illustrates that, across all moral frames, the moral tone of the comments tends to reflect the moral foundation scores of the corresponding campaign appeals. Specifically, appeals with lower moral foundation scores (i.e., emphasizing the negative poles) are associated with comments containing lower scores, while appeals with higher scores relate to comments with higher scores. These results suggest that supporters do not typically utilize comments to counterframe or re-balance negatively framed appeals; instead, the moral framing of the comments tends to align with the framing of the initial appeal.

Discussion

Understanding how moral framing influences people to engage in prosocial behavior is essential for enhancing online fundraising and social support for those in need. This study examined how three moral frames (care, fairness, and ingroup loyalty) are associated with the donation outcomes and social support received by GoFundMe fundraisers. Overall, the findings suggest that appeals highlighting the harm experienced by the help-seeker and emphasizing ingroup identities are positively associated with donor engagement across multiple fundraising categories. This trend was also observed on the micro-lending platform Kiva, where similar framing strategies were associated with an increased volume of loans (see Appendix).

Our results indicate that moral frames are associated with varying outcomes in GoFundMe campaigns, with the specific relationship appearing to depend on the fundraising category. For example, within the "Animal" category, positively

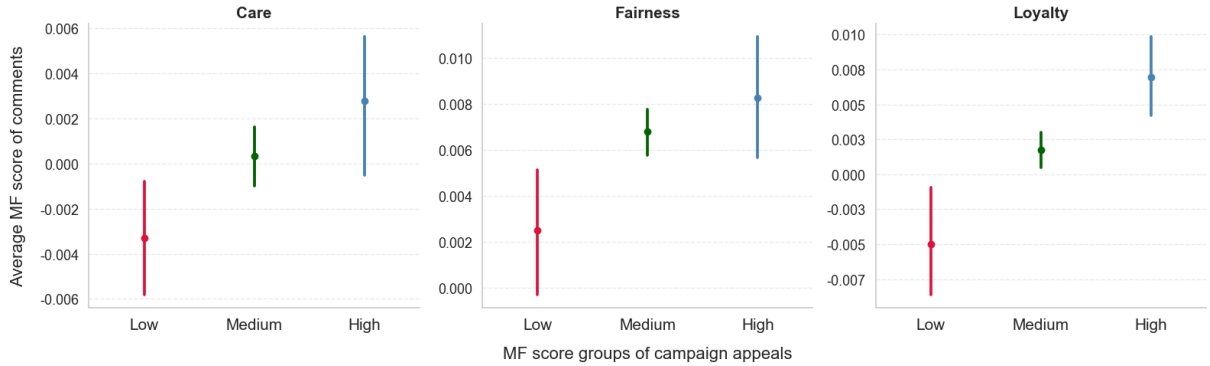

Figure 2: Average moral foundation scores of comments across three moral frames, categorized by the score groups of the corresponding campaign appeals. This analysis specifically uses appeals and their associated comments from the "Emergency" category. Error bars represent 95% confidence intervals.

framed appeals emphasizing care are linked to a higher volume of donations. In contrast, within the "Emergency" and "Memorial" categories, negatively framed appeals highlighting harm are associated with increased financial and social support. Similar patterns are observed on the Kiva platform, where appeals emphasizing harm are associated with fewer loans in the "Agriculture" category but relate to an increased number of loans in the "Food" and "Retail" sectors. These findings suggest that the efficacy of a specific moral frame is not uniform but is instead contingent upon the broader context of the request.

Furthermore, negatively framed appeals emphasizing harm were associated with higher levels of support across multiple categories on both the GoFundMe and Kiva platforms. One potential explanation is that supporters may perceive these situations as significantly more dangerous or urgent, eliciting a stronger behavioral response (van Doorn, Zeelenberg, and Breugelmans 2017). The results presented in Figure 2 support this mechanism by demonstrating that the moral foundation scores of comments generally align with the scores of their corresponding campaign appeals. For example, appeals with low care scores (e.g., Our cousin was the victim of a brutal, unprovoked, and completely random act of violence) tend to receive support messages with similarly low care scores (e.g., I am sickened, sad, and angry that this happened to [NAME]). In other words, appeals emphasizing harm appear to evoke potent emotional reactions from supporters who focus on the immediate adversity faced by the fundraisers. This finding further suggests that, when encountering negatively framed appeals, supporters are more likely to increase their frequency of donation rather than to post comments to provide a positive counterframe.

Unlike care framing, positively framed loyalty appeals were generally associated with higher levels of both monetary and social support. Previous studies have demonstrated that loyalty framing can be particularly effective for fundraising, as individuals are often more inclined to assist those perceived as members of an ingroup rather than an outgroup (Nilsson, Erlandsson, and Västfjäll 2016; Sisco and Weber 2019). However, the consistent positive association between ingroup loyalty and social support warrants further investigation, as it appears to diverge from prior research suggesting that expressions of positivity may reduce the volume of social support received. Specifically, some scholars argue that supporters may feel less compelled to provide assistance when a help-seeker appears to be in a positive or stable emotional state (Rhue and Robert 2018; Walsh and Forest 2021).

Implications

The findings of this research have several theoretical implications. First, this study extends existing literature by detailing how specific moral frames relate to fundraising outcomes. Rather than relying on a single binary measure of positive or negative sentiment, we analyzed the distinct associations of the care, fairness, and loyalty dimensions. This approach provides nuanced insights into the diverse mechanisms through which moral language may motivate donor engagement and social support. Additionally, our results suggest that the relationship between moral framing and campaign success is not uniform but varies significantly across fundraising categories and contexts. Furthermore, we found that supporters tend to mirror the moral tone of an appeal rather than posting comments to counterframe negativity. The concurrent increase in donation frequency suggests a compensation mechanism wherein supporters may prioritize financial contributions over altering the emotional or moral narrative of the appeal.

The methods and findings of this research may be applied to the study of support-seeking behavior more broadly. For instance, the moral frames analyzed in this study overlap with collective action frames, suggesting that similar mechanisms may shape how individuals mobilize support in social movements and civic campaigns. Benford and Snow (2000) argue that core framing tasks include identifying problems, proposing solutions, and motivating collective action. Our adoption of moral frames resonates with these elements, particularly the diagnostic and motivational tasks. For example, care framing functions diagnostically by attributing a specific problem or harm to an external cause (e.g., My

bunny was attacked by a neighbor's cat). Conversely, loyalty framing serves a motivational role by mobilizing ingroup members through appeals to community cohesion and shared identity (e.g., It's important for us to come together as a community). Building upon this connection, future research could utilize word embedding models to quantify these frames in the context of social movements, thereby extending the utility of moral framing analysis beyond the domain of online fundraising.

The findings of this research have several practical implications, particularly for campaign strategies and the design of fundraising activities and platforms. First, this study provides insights into what constitutes an effective fundraising appeal. Fundraisers can tailor their message framing according to their specific purposes and goals. For example, highlighting harm may attract more support in emergency situations. Meanwhile, emphasizing ingroup loyalty can be effective across all types of campaigns. Furthermore, these findings suggest potential features for fundraising platforms to optimize framing effectiveness. Platforms could offer tools or recommendations to assist fundraisers in framing their campaigns, along with analytics to track the impact of different framing strategies on donations and engagement. In addition, platforms may try enhancing their social support systems through targeted intervention features. For instance, developers could implement prompting systems that offer commenters a variety of pre-drafted, multi-framed message templates (e.g., care-based versus loyalty-based options). By allowing supporters to select or adapt a frame that resonates with their intent, platforms can better facilitate high-quality emotional support.

Limitations and Future Work

We acknowledge several limitations of this work and highlight key opportunities for future research. First, the dataset lacks granular information regarding the fundraisers and supporters on the platform (e.g., gender, socioeconomic status) which are factors previously shown to relate to fundraising outcomes. Access to such supporter-level data would help us identify which moral framing strategies are effective for different supporter segments. Furthermore, securing more comprehensive demographic or behavioral data could enable future studies to employ quasi-experimental designs, such as propensity score matching; this would allow for the identification of comparable campaigns and a rigorous measurement of the causal impact of moral framing on fundraising success.

Second, the findings of this study may not generalize to reward-based crowdfunding platforms, such as Kickstarter or Indiegogo, which primarily support creative or entrepreneurial projects. These platforms often impose restrictions on charitable or personal-cause fundraising. Furthermore, their persuasive framing typically centers on product specifications and reward attributes rather than moral foundations. Accordingly, reward-based campaigns are not directly comparable to the donation-based appeals examined in this study. Future research could investigate whether moral framing remains an important factor in social entrepreneurship projects that bridge the gap between these two crowdfunding models.

Another limitation of this study is the difficulty in isolating the mechanisms underlying our observations. For instance, after observing that negatively framed appeals emphasizing harm were associated with a higher volume of supporter comments, we examined the moral foundation scores of those messages. We initially hypothesized that negatively framed appeals would relate to more positive comments as supporters sought to counterbalance the negativity; however, our analysis did not support this hypothesis. This finding underscores the necessity for a deeper examination of the quality and linguistic content of social support in crowdfunding contexts. Future research would benefit from carefully designed experiments or surveys to explore these dynamics in greater detail.

Lastly, this study utilized existing resources, specifically MFT and its corresponding dictionary, to model the language and framing of campaign appeals through binary oppositions. While these frames are widely applied across diverse disciplines, future research could develop customized dictionaries and novel dimensions specifically tailored to the nuances of online fundraising. For instance, Best and Arseniev-Koehler (2023) constructed four stigma dimensions (immorality, negative personality traits, disgust, and danger) to analyze the negative depiction of diseases in major newspapers over time. In addition to moral frameworks, incorporating socioeconomic dimensions such as class, gender, and race, as adopted by Kozlowski, Taddy, and Evans (2019), would be a valuable direction for subsequent studies. Because these dimensions are often associated with donation dynamics and outcomes, identifying how the positioning of an appeal along these axes relates to fundraising success may offer deeper insights into the drivers of online philanthropy.

Future studies of fundraising campaign appeals may benefit from the incorporation of additional framing models, such as the collective action frames utilized in social movements and civic campaigns, as these frameworks share similar objectives and framing tasks. For example, incorporating prognostic frames (i.e., those that articulate plans or solutions) could enrich our understanding of how a clear course of action relates to a supporter's perception of efficacy. This, in turn, may correspond to more successful fundraising outcomes.

Conclusion

This research investigates the relationship between the moral framing of GoFundMe campaign appeals and the monetary and social support received by fundraisers. By adopting the FrameAxis approach, we quantified the moral foundation scores across three distinct dimensions: care, fairness, and loyalty. Our findings indicate that appeals emphasizing the harm experienced by the help-seeker and ingroup loyalty are associated with a higher volume of donations and supporter comments across several fundraising categories. Furthermore, the results suggest that an emphasis on ingroup loyalty relates to a higher average donation amount per donor. These insights contribute to a more nuanced un-

derstanding of how moral language relates to prosocial engagement in digital environments.

References

Benford, R. D.; and Snow, D. A. 2000. Framing Processes and Social Movements: An Overview and Assessment. *Annual Review of Sociology*, 26(2000): 611–639.

Best, R. K.; and Arseniev-Koehler, A. 2023. The Stigma of Diseases: Unequal Burden, Uneven Decline. *American Sociological Review*, 88(5): 938–969.

Blader, S. L.; and Tyler, T. R. 2002. Justice and Empathy: What Motivates People to Help Others? In *The Justice Motive in Everyday Life*, 226–250. Cambridge University Press.

Chakraborty, A.; Mota, N.; Biega, A. J.; Gummadi, K. P.; and Heidari, H. 2019. On the Impact of Choice Architectures on Inequality in Online Donation Platforms. In *The World Wide Web Conference*, 2623–2629.

Cunningham, M. R.; Steinberg, J.; and Grev, R. 1980. Wanting to and having to help: Separate motivations for positive mood and guilt-induced helping. *Journal of Personality and Social Psychology*, 38(2): 181.

Erlandsson, A.; Jungstrand, A. Å.; and Västfjäll, D. 2016. Anticipated Guilt for Not Helping and Anticipated Warm Glow for Helping Are Differently Impacted by Personal Responsibility to Help. *Frontiers in Psychology*, 7: 1475.

Erlandsson, A.; Nilsson, A.; and Västfjäll, D. 2018. Attitudes and Donation Behavior When Reading Positive and Negative Charity Appeals. *Journal of Nonprofit & Public Sector Marketing*, 30(4): 444–474.

Esperança, J. P. A.; Paz, F. L.; Curto, J. J. D.; and Ferreira, F. A. 2025. Determinants of prosocial crowdlending success. *Annals of Operations Research*, 1–23.

GoFundMe. 2025. Year in Help 2024. https://www.gofundme.com/c/gofundme-2024-year-in-help. Accessed: 2025-05-07.

Graham, J.; Haidt, J.; Koleva, S.; Motyl, M.; Iyer, R.; Wojcik, S. P.; and Ditto, P. H. 2013. Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism. *Advances in Experimental Social Psychology*, 47: 55–130.

Graham, J.; Haidt, J.; and Nosek, B. A. 2009. Liberals and Conservatives Rely on Different Sets of Moral Foundations. *Journal of Personality and Social Psychology*, 96(5): 1029.

Haidt, J.; and Graham, J. 2007. When Morality Opposes Justice: Conservatives Have Moral Intuitions that Liberals may not Recognize. *Social Justice Research*, 20(1): 98–116.

Hoover, J.; Johnson, K.; Boghrati, R.; Graham, J.; and Dehghani, M. 2018. Moral Framing and Charitable Donation: Integrating Exploratory Social Media Analyses and Confirmatory Experimentation. *Collabra: Psychology*, 4(1): 9.

Hutto, C.; and Gilbert, E. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1): 216–225.

Jang, H.; and Chu, W. 2022. The effect of message features on donations in donation-based crowdfunding. *Journal of Consumer Behaviour*, 21(6): 1464–1477.

Kiva. 2025a. Developer Home. https://www.kiva.org/build. Accessed: 2025-09-10.

Kiva. 2025b. How Kiva Works. https://www.kiva.org/about/how. Accessed: 2025-09-10.

Kozlowski, A. C.; Taddy, M.; and Evans, J. A. 2019. The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings. *American Sociological Review*, 84(5): 905–949.

Kumru, C. S.; and Vesterlund, L. 2010. The effect of status on charitable giving. *Journal of Public Economic Theory*, 12(4): 709–735.

Kwak, H.; An, J.; Jing, E.; and Ahn, Y.-Y. 2021. FrameAxis: characterizing microframe bias and intensity with word embedding. *PeerJ Computer Science*, 7: e644.

Linos, K.; Jakli, L.; and Carlson, M. 2021. Fundraising for Stigmatized Groups: A Text Message Donation Experiment. *American Political Science Review*, 115(1): 14–30.

Mikolov, T.; Grave, E.; Bojanowski, P.; Puhrsch, C.; and Joulin, A. 2018. Advances in Pre-Training Distributed Word Representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.

Mokhberian, N.; Abeliuk, A.; Cummings, P.; and Lerman, K. 2020. Moral Framing and Ideological Bias of News. In *Social Informatics: 12th International Conference*, 206–219.

Nilsson, A.; Erlandsson, A.; and Västfjäll, D. 2016. The congruency between moral foundations and intentions to donate, self-reported donations, and actual donations to charity. *Journal of Research in Personality*, 65: 22–29.

Parsell, C.; and Clarke, A. 2022. Charity and Shame: Towards Reciprocity. *Social Problems*, 69(2): 436–452.

Paxton, P.; Velasco, K.; and Ressler, R. W. 2020. Does Use of Emotion Increase Donations and Volunteers for Nonprofits? *American Sociological Review*, 85(6): 1051–1083.

Priniski, J. H.; Mokhberian, N.; Harandizadeh, B.; Morstatter, F.; Lerman, K.; Lu, H.; and Brantingham, P. J. 2021. Mapping Moral Valence of Tweets Following the Killing of George Floyd. *arXiv preprint arXiv:2104.09578*.

Radu, M. B.; and McManus, L. 2018. A Qualitative Analysis of Requests for Financial Help via GoFundMe by Victims of Intimate Partner Violence. *Sociological Spectrum*, 38(5): 312–325.

Rhue, L.; and Robert, L. P. 2018. Emotional Delivery in Pro-social Crowdfunding Success. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–6.

Sisco, M. R.; and Weber, E. U. 2019. Examining charitable giving in real-world online donations. *Nature Communications*, 10(1): 3968.

van Doorn, J.; Zeelenberg, M.; and Breugelmans, S. M. 2017. The impact of anger on donations to victims. *International Review of Victimology*, 23(3): 303–312.

Walsh, R. M.; and Forest, A. L. 2021. Can Expressing Positivity Elicit Support for Negative Events? A Process Model

and Review. *Personality and Social Psychology Review*, 25(1): 3–40.

Wang, X.; Shi, W.; Kim, R.; Oh, Y.; Yang, S.; Zhang, J.; and Yu, Z. 2019. Persuasion for Good: Towards a Personalized Persuasive Dialogue System for Social Good. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5635–5649. Florence, Italy: Association for Computational Linguistics.

Xu, X.; Li, J.; and Zhou, F. 2023. MDCC: A Multimodal Dynamic Dataset for Donation-based Crowdfunding Campaigns. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 5417–5421.

Paper Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes
   (e) Did you describe the limitations of your work? Yes, see the Discussion section.
   (f) Did you discuss any potential negative societal impacts of your work? NA, this research aims to support fundraisers and fundraising platforms.
   (g) Did you discuss any potential misuse of your work? NA
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, see the second paragraph of the Data section.
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? Yes
   (b) Have you provided justifications for all theoretical results? Yes
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? Yes
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? Yes
   (e) Did you address potential biases or limitations in your theoretical framework? Yes
   (f) Have you related your theoretical results to the existing literature in social science? Yes
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? Yes, see the Discussion section.
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? NA
   (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   (a) If your work uses existing assets, did you cite the creators? Yes, we have included links to the GoFundMe dataset and the Kiva API. We have also cited the papers relevant to the dataset, the Moral Foundations Dictionary, the word embedding model, and the code used for the FrameAxis analysis.
   (b) Did you mention the license of the assets? Yes, they are publicly available.
   (c) Did you include any new assets in the supplemental material or as a URL? NA
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes, we note that our dataset consists of publicly accessible GoFundMe campaigns, and we collected additional campaigns from the Kiva platform using its public API.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, we did not disclose the identities of any individuals in the paper. We also stated that we presented our results in aggregated form and paraphrased the quotes from campaign appeals (see the Data section).

(f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA

(g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? NA

6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...

(a) Did you include the full text of instructions given to participants and screenshots? NA

(b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA

(d) Did you discuss how data is stored, shared, and deidentified? NA

Appendix (Additional Evidence from Kiva)

Since our GoFundMe dataset consists exclusively of campaign appeals from U.S.-based fundraisers, we collected a comparable dataset from the micro-lending platform Kiva[2] to address this limitation and establish the external validity of our findings. Kiva's institutional mission and its focus on need-based appeals align closely with the characteristics of GoFundMe. However, a key distinction is that Kiva users solicit loans rather than donations. Kiva serves to connect lenders with borrowers in regions where access to traditional banking services is limited (Kiva 2025b). In this model, a borrower applies for a loan with the assistance of Field Partners and provides a narrative regarding the loan's purpose and a repayment schedule. Multiple lenders contribute small amounts until the funding goal is met, after which the borrower utilizes the capital for the stated purposes and repays the loan according to the designated schedule.

Using the publicly available Kiva API (Kiva 2025a), we initially retrieved 19,894 campaigns uploaded to the platform in 2025. We omitted 6,470 ongoing campaigns and 13 unfunded appeals, resulting in a final sample of 13,411 funded campaigns. While a loan request typically expires if the funding goal is not met within the standard 30-day window, the expiration rate on Kiva is notably low.

Our Kiva dataset contains campaigns from a diverse range of categories, with the most frequent being Agriculture, Food, and Retail. Projects within the Agriculture category typically support farmers facing unpredictable weather, variable crop yields, volatile markets, and challenges in meeting basic needs. The Food category comprises projects facilitating the operation of small markets, grocery stores, or restaurants. The Retail category includes borrowers seeking to establish small businesses. The remaining 16 categories (e.g., Sanitation & Hygiene, Services, Clothing, Education, Arts) were grouped into the "Others" category. Descriptive statistics for all continuous variables are presented in Table 6.

[2]https://www.kiva.org/

Our dataset includes campaigns from 55 countries, with the most frequent being the Philippines, Kenya, Madagascar, Uganda, and Ecuador. The remaining 50 countries were grouped into the "Others" category. We included three categorical variables (category, region, and campaign appeal sentiment) using "Agriculture," "Philippines," and "Negative" as the reference categories, respectively. As approximately 97% of the campaigns involved a single borrower, we created a binary variable to distinguish between campaigns supporting a single individual (0) and those supporting multiple borrowers (1). Finally, to address skewness in the data distributions, a log transformation was applied to both the number of loans and the total fundraising goal amount.

Although lenders can post support messages, our analysis was restricted to loan volume due to the limited number of comments available; specifically, our dataset contained only 110 messages. We quantified the moral foundation scores of the Kiva campaign appeals using the same FrameAxis procedure employed in our GoFundMe analysis. Table 7 presents examples of Kiva appeals alongside their corresponding moral foundation scores, demonstrating the variation in moral framing across different loan requests.

The results from the Kiva model largely mirror the patterns observed in the GoFundMe dataset (see Table 8). In the Food and Retail categories, emphasizing harm and ingroup loyalty was associated with an increase in the number of loans received. Conversely, within the Agriculture category, campaigns that utilized positive care framing (i.e., highlighting care over harm) relate to a higher loan volume. These findings further support the conclusion that effective framing strategies are dependent on the fundraising category. This aligns with the GoFundMe analysis, where the Animal category exhibited a distinct pattern compared to other campaign types.

Another notable finding is that loan volume is negatively associated with the number of borrowers per campaign. One potential explanation for this relationship is that supporters may feel a stronger personal connection when helping a single borrower, as opposed to a group. Supporters might also perceive assisting one individual as more tangible or manageable than supporting multiple people, regardless of the stated fundraising goal. This finding raises a question regarding how the number of help-seekers relates to supporter perceptions of efficacy. As the number of borrowers in a campaign increases, the perceived impact of a single donation may diminish, which could ultimately lead to fewer contributions.

	Agriculture	Food	Retail	Others
Number of loans	10.26 (14.57)	12.62 (15.86)	10.97 (16.15)	16.00 (44.93)
Average loan amount per loan	58.92 (84.42)	65.92 (110.71)	73.49 (118.01)	49.07 (77.85)
Care score	0.02 (0.01)	0.01 (0.01)	0.02 (0.01)	0.02 (0.01)
Fairness score	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Loyalty score	0.00 (0.01)	−0.01 (0.01)	−0.01 (0.01)	−0.01 (0.01)
Campaign appeal length	117.54 (38.34)	108.71 (42.31)	97.47 (39.30)	100.65 (47.87)
Number of borrowers	1.15 (1.43)	1.29 (2.17)	1.25 (2.14)	1.15 (1.62)
Fundraising goal	0.4k (0.6k)	0.6k (1.0k)	0.5k (0.9k)	0.6k (2.0k)
Number of campaigns	4,115	2,869	2,443	3,984

Table 6: Descriptive statistics for the Kiva dataset. Values represent means, with standard deviations provided in parentheses.

Moral domain	Bias	Campaign appeal
Care	Vice	[NAME], an experienced farmer and returning client, is planning to increase his herd by buying more cows, despite the persistent issue of foot and mouth disease.
	Virtue	[NAME] is a dedicated and talented seamstress who works hard to support herself and her family. [...] This financial support will enable her to create a safer and more comfortable home for herself and her loved ones.
Fairness	Vice	Most of the school children come from marginalized families who cannot afford their educational expenses. [NAME] has requested a loan to construct a building to accommodate all the students. [...]
	Virtue	[NAME] plans to expand her inventory and reach more customers. [...] She also aims to create more job opportunities and support her community.
Loyalty	Vice	[NAME] is a widow and the mother of one child. She trades in various types of glasses and is applying for a loan to increase her inventory and better satisfy her customers.
	Virtue	[ORGANIZATION] consists of women who are united by solidarity, effort, and commitment. They work as a team, supporting one another to move forward and improve their living conditions. [...]

Table 7: Examples of Kiva campaign appeals with varying moral foundation scores. Note that excerpts have been paraphrased to protect user privacy.

Variable	Coefficient	SE
Food	0.075**	0.024
Retail	0.072*	0.030
Others	0.093***	0.025
Care	2.964***	0.879
Fairness	−0.088	0.875
Loyalty	−1.544	0.916
Food × Care	−6.115***	1.308
Retail × Care	−6.325***	1.392
Others × Care	−0.423	1.151
Food × Fairness	2.134	1.427
Retail × Fairness	3.337	1.819
Others × Fairness	−0.616	1.231
Food × Loyalty	3.700**	1.297
Retail × Loyalty	4.510***	1.368
Others × Loyalty	0.915	1.241
Kenya	0.022	0.019
Madagascar	0.213***	0.020
Uganda	0.352***	0.023
Ecuador	0.147***	0.026
Others	0.178***	0.014
Positive sentiment	−0.041	0.025
Neutral sentiment	−0.014	0.057
Campaign appeal length	0.000	0.000
Multiple borrowers	−0.159***	0.029
Fundraising goal	0.713***	0.007
Intercept	−2.060***	0.049
Number of observations	13,411	
Adjusted R^2	0.59	

Table 8: Multiple linear regression results for the number of loans. Care, Fairness, and Loyalty values represent moral foundation bias scores. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.