DISCUSSION PAPER SERIES

IZA DP No. 7995

Referral Incentives in Crowdfunding

Victor Naroditskiy
Sebastian Stein
Mirco Tonin
Long Tran-Thanh
Michael Vlassopoulos
Nicholas R. Jennings

February 2014

Forschungsinstitut
zur Zukunft der Arbeit
Institute for the Study
of Labor

Referral Incentives in Crowdfunding

Victor Naroditskiy
University of Southampton

Long Tran-Thanh
University of Southampton

Sebastian Stein
University of Southampton

Michael Vlassopoulos
University of Southampton

Mirco Tonin
University of Southampton and IZA

Nicholas R. Jennings
University of Southampton

Discussion Paper No. 7995
February 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

IZA Discussion Paper No. 7995
February 2014

ABSTRACT

Referral Incentives in Crowdfunding

Word-of-mouth, referral, or viral marketing is a highly sought-after way of advertising. We undertake a field experiment that compares incentive mechanisms for encouraging social media shares to support a given cause. Our experiment takes place on a website set up to promote a fundraising drive by a large cancer research charity. Site visitors who choose to sign up to support the cause are then asked to spread the word about the cause on Facebook, Twitter or other channels. Visitors are randomly assigned to one of four treatments that differ in the way social sharing activities are incentivised. Under the control treatment, no extra incentive is provided. Under two of the other mechanisms, the sharers are offered a fixed number of points that help take the campaign further. We compare low and high levels of such incentives for direct referrals. In the final treatment, we adopt a multi-level incentive mechanism that rewards direct as well as indirect referrals (where referred contacts refer others). We find that providing high level of incentives results in a statistically significant increase in sharing behaviour and resulting signups. Our data does not indicate a statistically significant increase for the low and recursive incentive mechanisms.

JEL Classification: C93, D64, L31, M31

Keywords: crowdfunding, referral marketing

Corresponding author:

Mirco Tonin
School of Social Sciences
University of Southampton
Southampton SO17 1BJ
United Kingdom
E-mail: m.tonin@soton.ac.uk

1. INTRODUCTION

Within just a few years, we have seen a tremendous rise of social networks, which made sharing information easy and natural. A message that is shared on social networks can reach a large audience without the help of traditional media outlets. Furthermore, the message may be targeted towards the right audience as people sharing the message are likely to have friends with similar interests. A recent social experiment, the DARPA Network Challenge [Defense Advanced Research Projects Agency 2010], provided an example of how powerful this method of sharing information can be: a seemingly impossible task of locating 10 red weather balloons placed at secret locations throughout the US was solved within 9 hours.

Reaching audiences through social media is particularly important for organisations that do not have a large marketing budget. For example, many charitable drives or crowdsourcing projects [Douceur and Moscibroda 2007] rely on existing participants to recruit new donors or members. The Guardian[1] reports that 22% of all donations on a popular charitable giving platform, JustGiving, originate from posts current donors make on Facebook.

In this work, we ask the question of how to incentivise social media sharing behaviour. To this end, we undertake a natural field experiment where we test the performance of three incentive mechanisms. The compared mechanisms encourage people to post a message on Facebook, Twitter or via other channels. The mechanisms differ in the way they reward signups resulting from a user's social media post:

— 1 extra point for each contact who signs up (**low** treatment),
— 3 extra points for each contact who signs up (**high** treatment),
— 1 extra points for each contact who signs up, 1/2 for each contact of a contact, 1/4 for each contact of a contact of a contact and so on (**recursive** treatment),
— and a base case treatment providing no additional incentive (**control** treatment).

The experiment is run in the domain of charitable giving, where the goal of sharing a message is to generate support for a charitable cause described on a website that we designed. The users that arrive at the website are asked to sign up to "donate a click". Each click generates 1 point that together with the points generated by our referral mechanisms are used to measure the success of the fundraising campaign. Once a threshold of 1,000 points is reached, £500 is donated to support the drive. Once the next target of 10,000 points is reached, another £500 is donated.

We include both the 1-point and the 3-point mechanisms to test whether the exact amount of points matters for the referral activity. The *recursive* or multi-level incentive mechanism offers stronger incentives than the 1-point mechanism (as it offers the same number of points for direct referrals, but also rewards indirect referrals). However, how it compares to the 3-point mechanism depends on a user's social network. Indeed, either the recursive or the 3-point mechanism may generate more points for the referrer. This mechanism implements the winning mechanism of the DARPA Network Challenge [Pickard et al. 2011].

We evaluate mechanisms along three dimensions. First, we look at the effectiveness of each treatment at encouraging participants to make referrals, as measured by the number of treatment participants who clicked one of the share buttons our website provides. The second dimension is the number of referred visitors to the page, and the third one is the number of those visitors who signed up.

Our main findings are that incentives make a difference in this context, and the exact amount or magnitude of the incentive is important. The 3-point incentive mech-

[1] http://www.theguardian.com/voluntary-sector-network/2012/jul/03/justgiving-donations-facebook-share

anism outperformed the control treatment on all of the metrics we considered (see details in Section 4). The recursive incentive mechanism was not any more successful than a simple 1-point mechanism. One might expect that the presence of any kind of incentives to share should be enough for most people, while further differentiation into 1, 3, or recursive points should have little bearing. However, our results suggest that offering 3 points results in more engagement than offering 1 point or compensating according to the recursive mechanism.

The paper is organised as follows. Section 2 covers related work. Details of our experimental design and execution of the experiment appear in Section 3. The results are presented in Section 4. The last section discusses our findings.

2. RELATED WORK

We are aware of only one other empirical work comparing incentives for referrals. In a recent study Castillo et al. [2014] also conduct a field experiment in the charitable giving context. The authors describe a field experiment that compared mechanisms offering to donate a fixed amount if a user posts about the cause on Facebook. The authors find that offering $1 for posting on Facebook encourages the action. However, a cost-benefit analysis reveals that under this referral program, each extra dollar donated by referred users costs more than $1 in donations by the company to incentivise referrals. Unlike our study, the referral compensation money is donated based on the posting action, and not based on the impact it generates through donations from referred users. Also, our point-based compensation does not directly map to money. The points collected only result in a donation if a threshold is reached.

The question of referral schemes received a lot of interest in the theoretical community. A model where incentives must be provided for users to propagate a question until a node that knows the answer is reached was studied by Kleinberg and Raghavan [2005]. The use of recursive mechanisms in this context was considered by Cebrian et al. [2012]. A theoretical justification for the recursive MIT mechanism was provided in [Naroditskiy et al. 2012]. Desirable properties of a referral scheme have been posed in [Douceur and Moscibroda 2007]. The question of sybil-proofness of incentive schemes has been considered in a series of papers [Babaioff et al. 2012; Drucker and Fleischer 2012; Chen et al. 2013; Lv and Moscibroda 2013]. However, these papers concentrate on theoretical issues and do not investigate or compare referral mechanisms empirically with real users.

Evidence that personal referrals are important in the fundraising context is presented in [Meer 2011]. More generally, a number of studies have shown that peer pressure has an impact on charitable giving (e.g. [Frey and Meier 2004; Shang and Croson 2005; Smith et al. 2013]).

In a business context, Albuquerque et al. [2012] observe that referral activities of publishers on an online magazine platform bring significant revenue to the platform. One of the recommendations of the paper is to stimulate referral activity with incentive schemes. Schmitt et al. [2011] find that referred customers are more valuable to a bank than non-referred customers. A reason for this is the targeting ability of crowdsourced referrals. Since an existing customer has information about the bank and about his friends, he only refers friends whose needs match what the bank provides. Domingos and Richardson [2001] provide a framework for estimating the value of each customer, taking into account the social influence that customer has on his social network.

3. FIELD EXPERIMENT

In order to compare a range of mechanisms for incentivising referrals, we partnered with Cancer Research UK (CRUK), one of the UK's largest charitable organisations, to carry out a natural field experiment. To this end, we set up a website to raise awareness of one of CRUK's existing fundraising campaigns. The website provided an indirect incentive for visitors to engage with it by adding a point to an overall total displayed on the website whenever a new user signed up. This total was then used to release donations from an outside donor to the CRUK campaign — whenever certain point thresholds were reached, a new donation was made.

After signing up, users were randomly assigned to one of a set of referral mechanisms. Here, some of the users were able to add further points to the common total by successfully referring their social contacts to the website. This allowed us to test whether such additional incentives can increase the probability of a successful referrals, and, by testing various mechanisms, whether the magnitude of the reward is significant and whether recursive schemes offer a benefit in this setting.

In the following, we describe the setup of our field experiment and the tested referral mechanisms in more detail.

3.1. Website

Our website, http://outruncancer.co.uk/, was set up in September 2013 in consultation with CRUK and provided visitors with information about oesophageal cancer and a particular CRUK fundraising campaign, "The Cancer Marathon". This campaign was initiated by six medics at the Southampton General Hospital, who were participating in the New York Marathon in November 2013 to raise money for research on oesophageal cancer.

The stated purpose of our website was to support the campaign by spreading awareness about it and thereby attract further donations from the public. To achieve this and encourage engagement with our site, visitors could contribute one point to a total amount prominently displayed on the website simply by signing up. This did not incur a financial cost for them and involved either entering an email address and password, or signing in through existing Facebook or Twitter accounts. Importantly, we pledged to make several donations of £500 each to CRUK's fundraising campaign, as soon as certain thresholds in the total number of points generated by all users were reached:

(1) £500 when 1,000 points were reached in total
(2) £500 when 10,000 points were reached in total
(3) £500 when 50,000 points were reached in total
(4) £500 when 100,000 points were reached in total

We chose this scheme, in order to cope with varying total amounts of visitors to our site. The first threshold was relatively low, in order to be able to generate at least one donation, while the next thresholds were spaced increasingly further apart, to ensure that the experiment did not have to terminate due to running out of money for donations.

Figure 1 shows the main part of the initial landing page (the full page is shown in Figure 7 in the appendix). Here, a prominent box displayed the current progress to the next donation, a summary of how the website works and a registration panel. Below this (see Figure 7), there were expandable sections offering more details about the campaign and the website. Importantly, as the experiment was designed as a natural field experiment, the website did not disclose that (anonymised) data from users would support research on referral incentives. This was crucial, because knowledge of the research could have altered the users' referral behaviour.

Apart from the dynamically updated total number of points, all visitors to the website received exactly the same information, and there was no explicit mention of referrals. This was done to avoid any priming of the visitors, and to mitigate self-selection bias regarding referral behaviour.

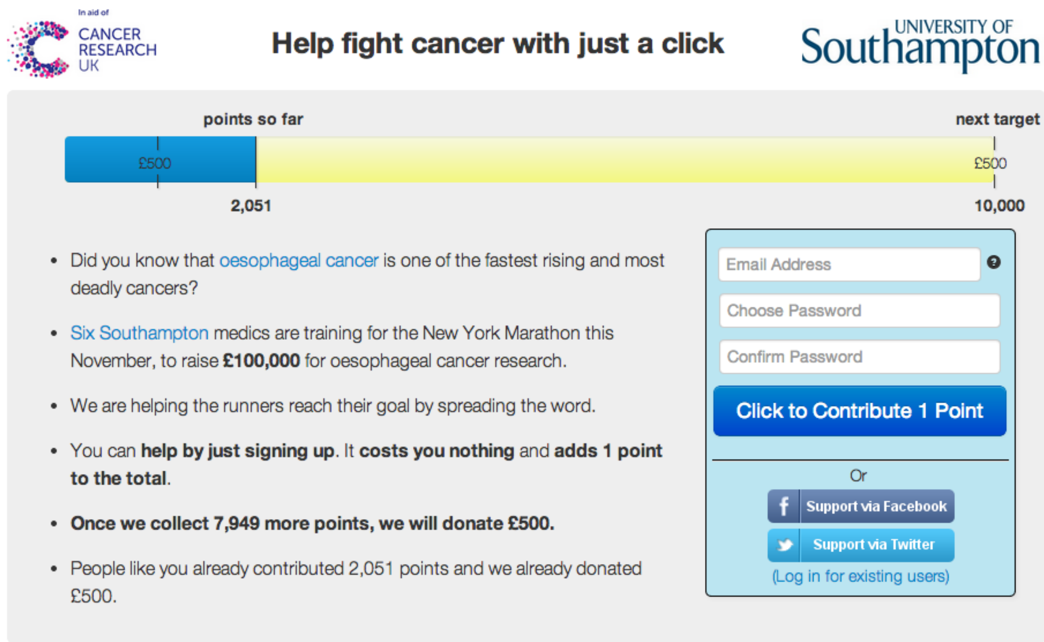

Fig. 1: Main panel of the landing page.

3.2. Referral Mechanisms

Once a visitor signed up to the website, they were presented with a page thanking them for their contribution and suggesting a number of ways they can help the fundraising campaign (Figure 2 shows the main part of this page, while the full page is shown in Figure 8 in the appendix). In particular, they were asked to invite their friends via a range of channels: by copying and sending a personalised link, by sharing a link on Facebook or Twitter, or by sending an email with the link.

Here, we provided some users with additional incentives, depending on the treatment they were randomly allocated to. This random allocation took place on signup and used a simple blocked randomisation scheme, whereby the allocation mechanism repeatedly cycled through random permutations of the available treatments (thus ensuring a near-balanced allocation to treatments).

The treatments we considered were as follows:

— **Control:** Users on this treatment were offered no additional referral incentives.
— **Low:** Users on this treatment were told that 1 additional point would be added to the total when they successfully referred someone else to the website (in addition to the signup point generated by that contact).
— **High:** Similar to *Low*, but users were offered 3 additional points.
— **Recursive:** Users were told that 0.5^{n-1} additional points would be added to the total when they successfully referred someone else to the website, including through *indirect* referrals, where n is the distance to the user.[2] Thus, if the user successfully referred another contact, 1 additional point would be added to the total. If that

[2]We say user A indirectly referred B, if B was not (directly) referred by A, but by another user C, which was either referred by A or is also an indirect referral of A. The distance between A and B is the length of the resulting referral chain (a direct referral has distance 1).

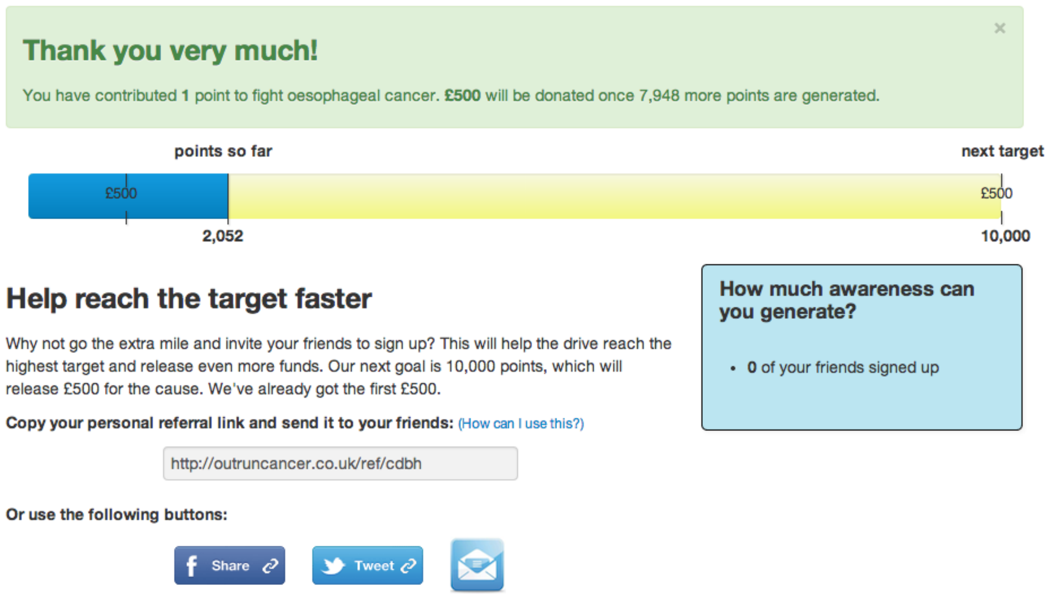

Fig. 2: Referral options after signing up (*control* treatment shown).

referred contact successfully referred someone else, 0.5 additional points would be added, and so on.

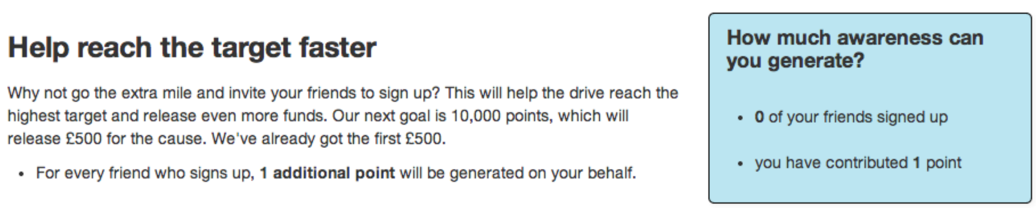

Fig. 3: Explanatory text for *low* treatment. Text for *high* treatment is identical except for the number of additional points (3).

The wordings of these treatments, as shown to the users, are summarised in Figures 3 and 4. Referrals were tracked through a unique ID that was allocated to each user and then included in the referral links generated through the four referral methods on the website (direct link, Facebook, Twitter or email). These links also encoded which of the four methods was used, allowing us to track, for each referred user, who they were referred by and through which referral method. Overall, we recorded the following three key metrics for each user:

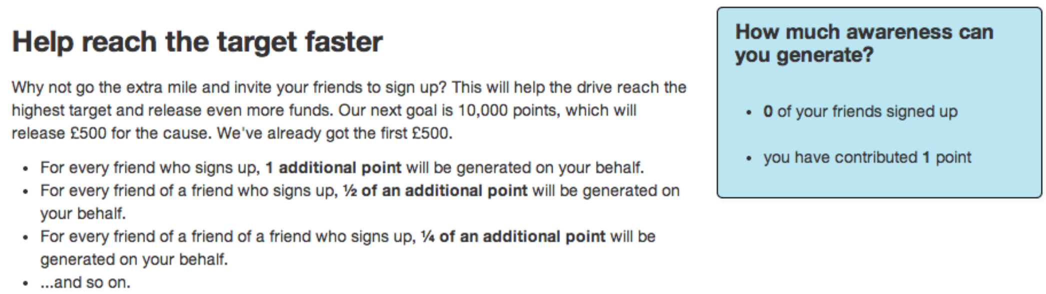

Fig. 4: Explanatory text for *recursive* treatment.

— **Number of distinct referral actions taken:** How many of the four referral methods were initiated by the user on our website. We record this when the user clicks on one of the buttons (or the referral link) shown after signing up.[3]
— **Number of resulting visits:** How many visits were subsequently generated by people following the user's unique referral link (but did not necessarily sign up).
— **Number of resulting signups:** How many new users subsequently signed up through the user's referral link (this is number of the user's direct referrals).

Note that for all four treatments, there is an implicit incentive for users to refer contacts to the website, as all signups generate 1 point towards the total. The *low*, *high* and *recursive* treatments add an additional incentive on top of this. The first two of these, *low* and *high*, are simple schemes that we chose to investigate whether the magnitude of this incentive is relevant — in the first case, it is equivalent to the original incentive offered to the user for signing up; in the second case, it is significantly higher than the original incentive. We included the *recursive* scheme, because it has been used successfully in previous crowdsourcing applications (see Section 2) and because it may encourage users to specifically target contacts with a wide social reach. From a user's perspective, it dominates the *low* scheme, because it also awards 1 additional point for each direct referral, but may generate further points through subsequent indirect referrals. However, the relative attractiveness of *recursive* and *high* depends on the user's social network. For example, a user who has only one friend who is very well-connected, prefers the *recursive* scheme as he will be rewarded half a point on each signup of well-connected friend's friends.

Users were able to log back into the website at any time to view how many of their contacts had signed up and (except for the *control* treatment) how many additional points had been generated through this. We also gave users the option to donate directly to the Cancer Marathon campaign, which we matched one-to-one (up to £10 per user).

3.3. Subject Recruitment

The website went live in September 2013 and over the next two months, 1577 unique users signed up, generating a total of 2,051 points and an additional £418 in user donations. As a result of meeting the first threshold, we donated £500, as well as £408 in matching donations (this is slightly lower than the total amount donated by the users due to the £10 limit per user).

[3]Note that this is not an exact measure of referral messages posted, as users may still cancel the message on the social media website before posting or may never share the link. Instead, it is only an approximate indicator of referral behaviour.

Initial users were recruited from a range of sources. Some publicity was generated by CRUK and the medics running the Cancer Marathon, including through Twitter messages and a local radio interview. However, the bulk of the users were attracted through emails within the University of Southampton, sent to staff and student mailing lists. This email did not mention that the website was part of a research project and did not identify the academics involved in it (see Figure 9 in the appendix).

During the course of the experiment, we received a significant amount of feedback from users, indicating that the points mechanism was not explained clearly and succinctly enough. For this reason, we made a number of changes to the website in mid-October 2013 — primarily to highlight the current points collected to date at the top of the page and to emphasise their purpose. As the experimental conditions changed at that time, we focus only on the 412 participants recruited after the change in the remainder of this paper. These are mainly students from the Faculties of Humanities, Social Sciences, and Health Sciences and their referred contacts, as we emailed their internal mailing lists after the change (and had not previously contacted them).

4. RESULTS

Table I summarises the number of actions taken under each treatment as well as the number of resulting visits and signups. We observe that the 3-point treatment resulted in higher totals across the board. There is no other clear dominance across treatments. Somewhat of an outlier is a low ratio of recursive treatment's visits and signups. There were 103 visits but only 2 signups resulting in a 2% signup ratio. The corresponding ratios for the control, low, and high treatment are 14%, 6%, and 11% respectively. Given the limited number of observations, we could not test for the dependence of the ratio by treatment. However, given its broad range, we conjecture that the variance in the ratio is high, and there is no cause that is treatment specific.

	Control	Low	High	Recursive	Total
Action channels used	13	15	32	28	88
Resulting visits	37	98	218	103	456
Resulting signups	5	6	24	2	37

Table I: Total outcomes per treatment.

We present details of signups and visits in Figures 5 and 6. The histogram in Figure 5 shows how many referrers signed up exactly 1 friend, 2 friends, . . . , 8 friends. For the high treatment, there were 6 referrers who signed up 1 friend, 3 referrers who signed up 2 friends, 2 referrer who signed up 4, and 1 referrer who signed up 8 friends. The histogram of visits in Figure 6 should be read in a similar manner. The observations there are bucketed into equally sized bins. For example, the first bin includes all referrers whose social media shares resulted in the number of referred visits between 1 and 7.

Next, we evaluate the results we obtained from our experiment. In particular, we focus on both the extensive and intensive margins of the behaviour under investigation. The former is a binary (0/1) indicator that reflects whether a certain success measure is achieved. The latter quantifies the degree of success.

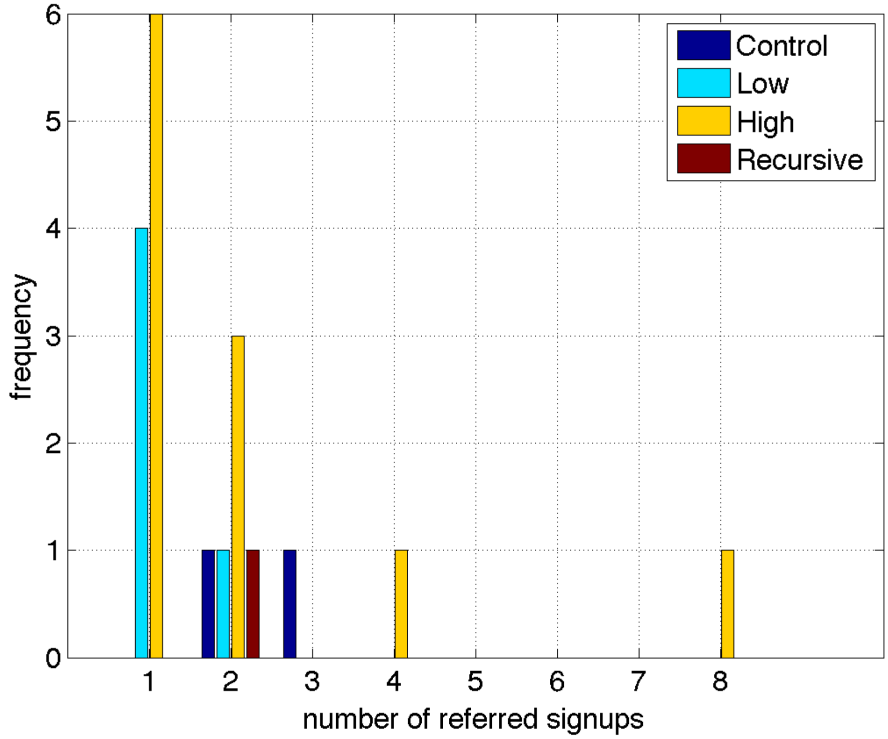

Fig. 5: Histogram of signups per user.

4.1. Extensive Margin

In this section we examine whether there are differences in the outcomes of interest in the extensive margin. We focus on three binary outcomes: (i) whether a referral action was taken; (ii) whether at least one visit was generated; (iii) whether at least one signup was generated.

	Control	Low	High	Recursive	Total
Action taken	10 (9.8%)	14 (13.5%)	22 (21.4%)	19 (18.4%)	65 (15.8%)
Not taken	92 (90.2%)	90 (86.5%)	81(78.6%)	84 (81.6%)	347 (84.2%)
Total	102	104	103	103	412

Table II: Referral actions per treatment.

With regards to referral actions, we see in Table II that people in our sample used some of the tools we provided to refer other potential participants in 16% of the cases (65 of the total 412). With no incentives, this happened in only 10% of the cases, while the percentage is always higher when referral incentives are present. In particular, the high incentive treatment is successful in eliciting a referral action among 21% of the people, while for the low incentive treatment the corresponding figure is 14%. The recursive treatment is in between, with a figure of 18%. The p-value of a chi-squared test is, at 0.106, close to the conventional 10% confidence level for rejection of the null

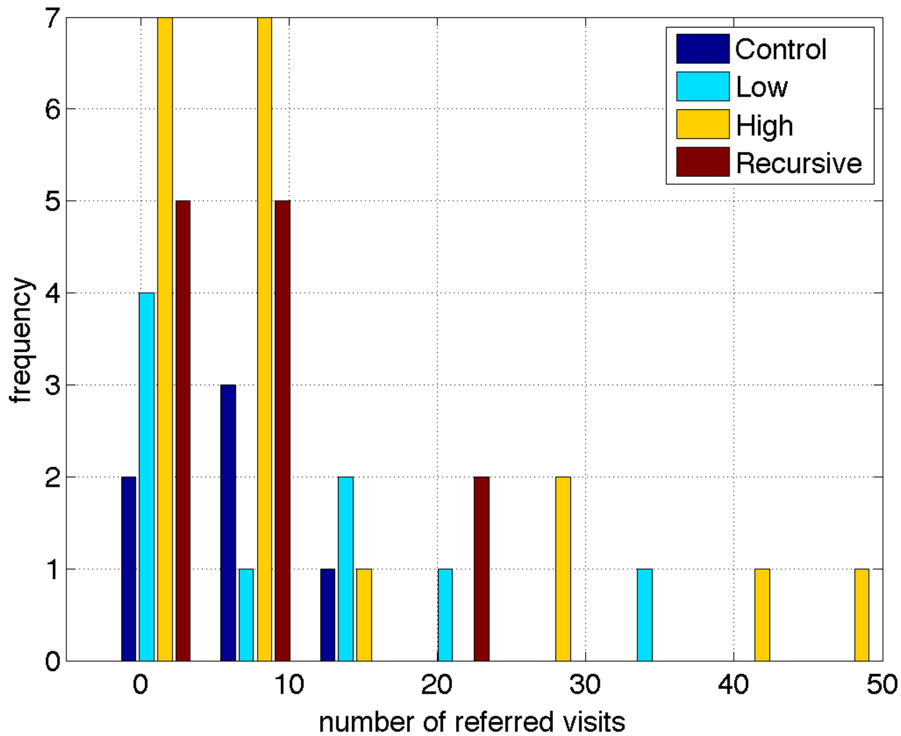

Fig. 6: Histogram of visits generated per user.

hypothesis of independence between the likelihood of undertaking some referral action and the various treatments.

	Control	Low	High	Recursive	Total
Visits	6 (5.9%)	9 (8.7%)	19 (18.4%)	12 (11.7%)	46 (11.2%)
No visit	96 (94.1%)	95 (91.3%)	84(81.6%)	91 (88.3%)	366 (88.8%)
Total	102	104	103	103	412

Table III: Generated visits per treatment.

Out of the 16% of people who used some of the referral tools, 75% were successful in generating some visits to the website. That is, out of 65 people who took a referral action, 46 resulted in at least one new user visiting the website (see Table III). The effect of incentives in generating additional traffic is even more evident, with only 6% of the people in the control group generating some visits, while the figure for the high incentive group, at 18%, is three times higher. Again, the low incentives are less successful (9%) and recursive incentives are positioned in between (12%). A chi-squared test rejects the null hypothesis of independence between the likelihood of generating some visits and the various treatments (p-value=0.028).

The percentage of people who generated additional signups is 5% (see Table IV). Again, our referral treatments seem to be effective in boosting this number. Only 2%

	Control	Low	High	Recursive	Total
Signups	2 (2%)	5 (4.8%)	11 (10.7%)	1 (1%)	19 (4.6%)
No signup	100 (98%)	99 (95.2%)	92 (90.3%)	102 (99%)	393 (95.4%)
Total	102	104	103	103	412

Table IV: Signups per treatment.

of people assigned to the control group generated additional signups, while in the high incentives case, the figure is much higher, at 11%. Low incentives are less effective (5%), while the recursive treatment is the least effective (1%). Again, a chi-squared test rejects the null hypothesis of independence between the likelihood of generating additional signups and the various treatments (p-value=0.004).

	Control	Low	High	Recursive
Referral actions	-1.293***	0.188	0.499**	0.394*
	(0.17)	(0.23)	(0.22)	(0.22)
Visits	-1.565***	0.202	0.666***	0.372
	(0.20)	(0.27)	(0.25)	(0.26)
Signups	-2.062***	0.398	0.818**	-0.276
	(0.29)	(0.36)	(0.33)	(0.47)

Table V: Regression analysis of the treatments.
Probit regression — robust standard errors are given in parentheses. *** [**] (*) denote significance at 1, [5], (10) % level

These results are confirmed by a regression analysis (Table V), where it is evident that the high incentives treatment is effective in generating significantly more referral actions, more visits, and more signups than the control group. For the other treatments, the coefficients are generally positive but not statistically significant. Note also that with regards to signups, the difference between the high incentives and the recursive treatment is also statistically significant (p-value=0.007), while the difference between high and low is marginally not (p-value=0.116). Finally, for visits, we find a significant difference between the coefficients of the high versus the low treatment (p-value=0.04).

4.2. Intensive Margin

In this section we examine whether there are differences in the outcomes of interest in the intensive margin. We focus again on three outcomes: (i) number of distinct referral actions; (ii) total number of visits generated; (iii) total number of signups generated (as described in Section 3.2).

	Control	Low	High	Recursive	Total
Action channels used	0.1	0.14	0.26	0.23	0.18
Resulting visits	0.36	0.94	2.12	1	1.11
Resulting signups	0.05	0.06	0.23	0.05	0.09

Table VI: Average outcomes per treatment.

Table VI presents averages for the various outcomes across the four treatment groups. It is evident that the high incentives treatment achieves the highest average in all measures, and in particular, with regards to the number of signups. In fact, t-tests of differences in means of the various outcomes between participants in the high and the control group indicate statistically significant differences in all cases (p-values are 0.008 for actions, 0.018 for visits and 0.069 for signups).

Finally, when we perform nonparametric tests (Mann-Whitney test) of differences in the distribution of the various outcomes between participants in the high and the control group we find statistically significant differences in all cases (p-values are 0.018 for actions, 0.005 for visits and 0.012 for signups). When comparing high to low treatment, we find significance differences for visits (p=0.04), while for sign-ups the p-value is 0.11. Also, the difference in the distribution of sign-ups of high and recursive treatments is significant (p-value=0.003).

5. DISCUSSION

Word-of-mouth, referral, or viral marketing is a highly sought after way of advertising. Ideally, it would occur organically because people find the promoted message of high value. It is however notoriously difficult to design viral messages. In this paper we attempt not to optimise the message, but to find out how sharing of a given message can be stimulated.

Our finding that the 3-point mechanism is more effective than the 1-point mechanism points to the need to choose the level of compensation carefully. This is contrary to the intuition that it is not the exact level, but rather the presence of some form of incentives, which has the most effect. Consistent with this intuition, Castillo et al. [2014] find little difference in referral activity between a mechanism that pays $1 and $5 for a sharing action. The corresponding sharing rates are 37% and 39%.

One might expect incentives to not matter at all in the context where points do not directly correspond to donations, but only matter if a sufficient number of points is reached. This is even more plausible if the threshold number of points does not appear to be within easy reach. This was the case in our experiment. All of the observations used in the data analysis came from the users who signed up after the first target of 1,000 points was reached, and when the second target of 10,000 points was far away. Indeed, the total number of points at the end of the experiment was just above 2,000. No user could reasonably expect that their friends would bring 8,000 points no matter which treatment they were part of. The result that users still reacted to the 3-point treatment more than to the 1-point or recursive treatments is even more interesting given the context.

Our mechanisms “pay” for successful referrals and are complementary to the mechanisms that pay for referral actions. Cost-benefit analysis of paying for referral actions results in cautionary lessons [Castillo et al. 2014]. We conjecture that the mechanisms considered in our work can be more cost-effective. In particular, one could guarantee an arbitrary cost-benefit ratio by setting the amount of reward as a fraction of the donation made by the referred donors.

ACKNOWLEDGMENTS

We would like to thank the team at CRUK and, in particular, Hannah Fox and Tim Underwood. This work was funded by the Strategic Interdisciplinary Research Development Fund of the Faculty of Social and Human Sciences at the University of Southampton, and by the EPSRC ORCHID programme grant.

This work is supported by the British Academy Small Research Grant SG131130.

REFERENCES

ALBUQUERQUE, P., PAVLIDIS, P., CHATOW, U., CHEN, K.-Y., AND JAMAL, Z. 2012. Evaluating promotional activities in an online two-sided market of user-generated content. *Marketing Science 31,* 3, 406–432.

BABAIOFF, M., DOBZINSKI, S., OREN, S., AND ZOHAR, A. 2012. On bitcoin and red balloons. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 56–73.

CASTILLO, M., PETRIE, R., AND WARDELL, C. 2014. Fundraising through online social networks: A field experiment on peer-to-peer solicitation. *Journal of Public Economics* 0, –.

CEBRIAN, M., COVIELLO, L., VATTANI, A., AND VOULGARIS, P. 2012. Finding red balloons with “split” contracts: robustness to individuals' selfishness. In *ACM Symposium on Theory of Computing (STOC)*.

CHEN, W., WANG, Y., YU, D., AND ZHANG, L. 2013. Sybil-proof mechanisms in query incentive networks. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 197–214.

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY. 2010. DARPA Network Challenge Project Report.

DOMINGOS, P. AND RICHARDSON, M. 2001. Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '01. ACM, New York, NY, USA, 57–66.

DOUCEUR, J. R. AND MOSCIBRODA, T. 2007. Lottery trees: motivational deployment of networked systems. In *SIGCOMM* (2007-10-23), J. Murai and K. Cho, Eds. ACM, 121–132.

DRUCKER, F. A. AND FLEISCHER, L. K. 2012. Simpler sybil-proof mechanisms for multi-level marketing. In *Proceedings of the 13th ACM conference on Electronic commerce*. ACM, 441–458.

FREY, B. S. AND MEIER, S. 2004. Social comparisons and pro-social behavior: Testing” conditional cooperation” in a field experiment. *The American Economic Review 94,* 5, 1717–1722.

KLEINBERG, J. AND RAGHAVAN, P. 2005. Query incentive networks. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*. IEEE, 132–141.

LV, Y. AND MOSCIBRODA, T. 2013. Fair and resilient incentive tree mechanisms. In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*. ACM, 230–239.

MEER, J. 2011. Brother, can you spare a dime? peer pressure in charitable solicitation. *Journal of Public Economics 95,* 7-8, 926 – 941.

NARODITSKIY, V., RAHWAN, I., CEBRIAN, M., AND JENNINGS, N. R. 2012. Verification in referral-based crowdsourcing. *PLoS ONE 7,* 10.

PICKARD, G., PAN, W., RAHWAN, I., CEBRIAN, M., CRANE, R., MADAN, A., AND PENTLAND, A. 2011. Time-critical social mobilization. *Science 334,* 6055, 509–512.

SCHMITT, P., SKIERA, B., AND VAN DEN BULTE, C. 2011. Referral programs and customer value. *Journal of Marketing 75,* 1, 46–59.

SHANG, J. AND CROSON, R. 2005. Field experiments in charitable contribution: The impact of social influence on the voluntary provision of public goods. *Unpublished manuscript*.

SMITH, S., WINDMEIJER, F., AND WRIGHT, E. 2013. Peer effects in charitable giving: Evidence from the (running) field.

Received February 2014; revised ; accepted

APPENDIX

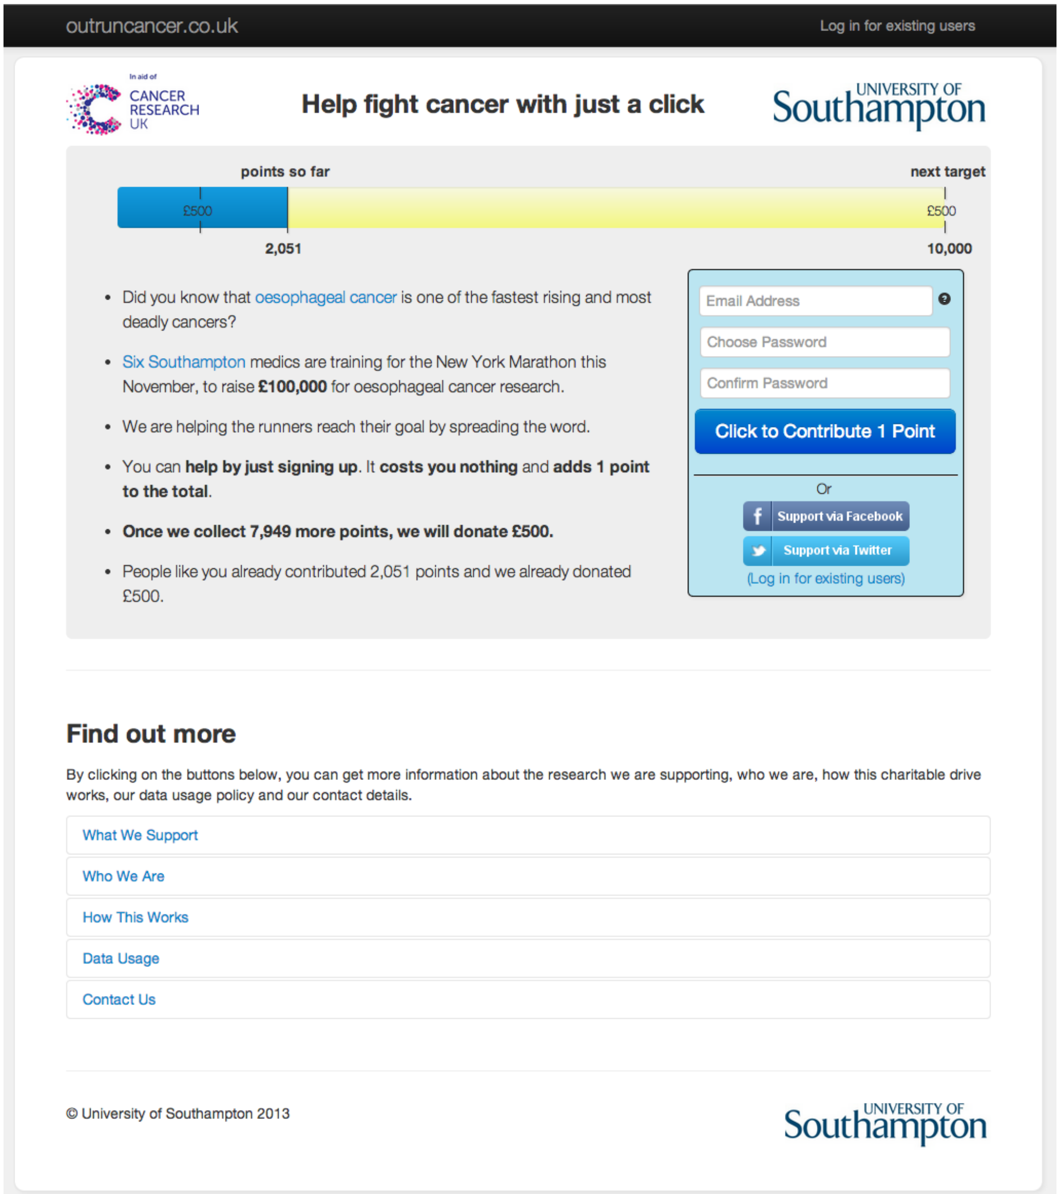

Fig. 7: The landing page shown to new visitors.

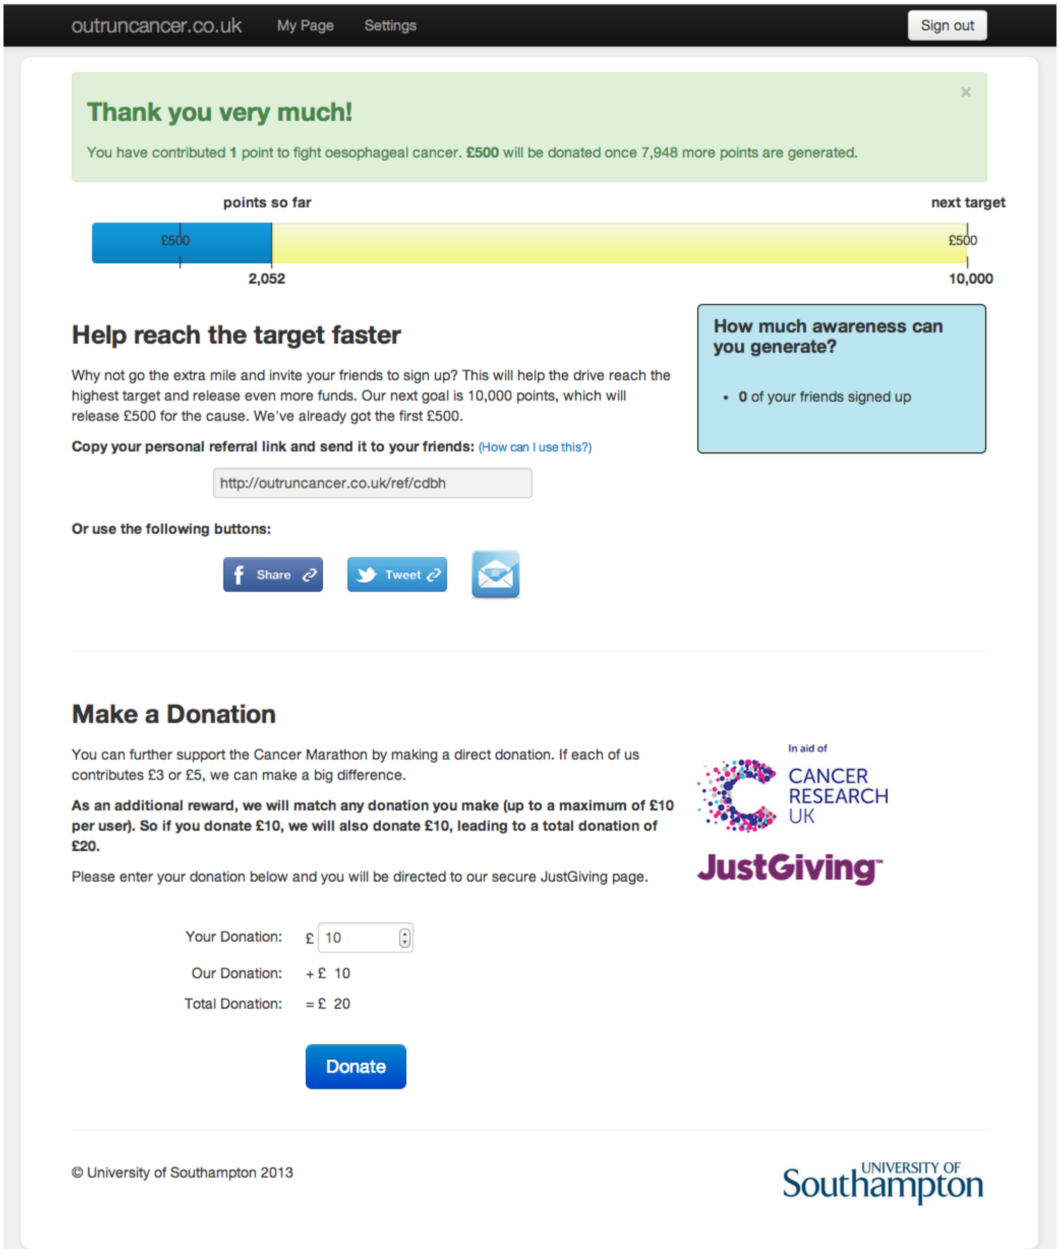

Fig. 8: Call to action after signing up (control treatment shown).

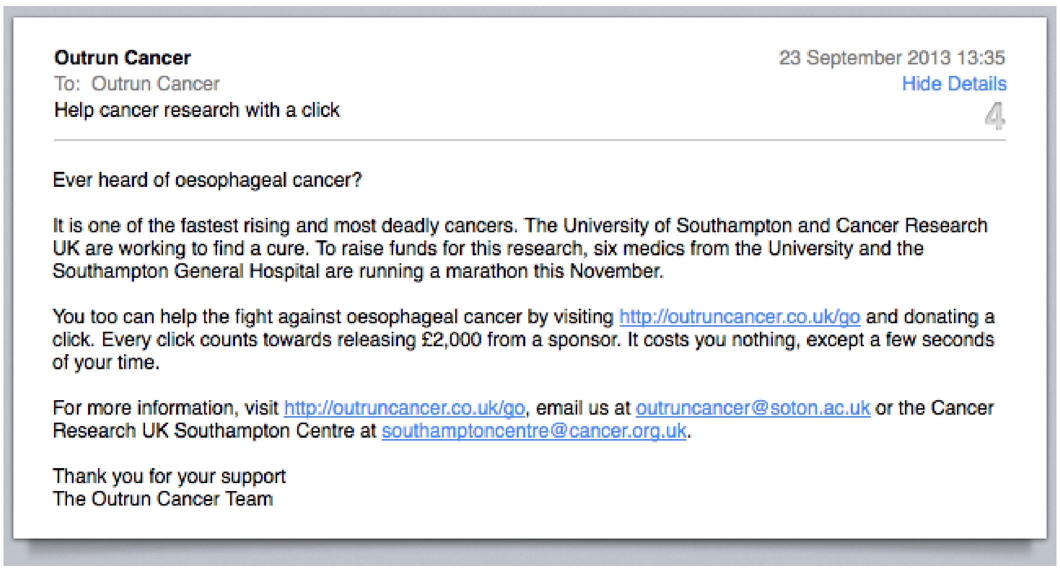

Outrun Cancer
To: Outrun Cancer
Help cancer research with a click

23 September 2013 13:35
Hide Details

Ever heard of oesophageal cancer?

It is one of the fastest rising and most deadly cancers. The University of Southampton and Cancer Research UK are working to find a cure. To raise funds for this research, six medics from the University and the Southampton General Hospital are running a marathon this November.

You too can help the fight against oesophageal cancer by visiting http://outruncancer.co.uk/go and donating a click. Every click counts towards releasing £2,000 from a sponsor. It costs you nothing, except a few seconds of your time.

For more information, visit http://outruncancer.co.uk/go, email us at outruncancer@soton.ac.uk or the Cancer Research UK Southampton Centre at southamptoncentre@cancer.org.uk.

Thank you for your support
The Outrun Cancer Team

Fig. 9: Email to invite participants.