How to Get the Most from Goods Donated to Charities

Extended Abstract

Christopher Culley
University of Southampton, UK
cc2u18@soton.ac.uk

Ji Qi
University of Essex, UK
jq16962@essex.ac.uk

Carmine Ventre
University of Essex, UK
c.ventre@essex.ac.uk

ABSTRACT

The charity sector is assuming a central role in many countries, due to a generalized increase in wealth inequalities and the restructuring of the welfare state. This market, however, exhibits inefficiencies. In this work, we empirically test the adoption of a centralized truthful allocation mechanism without money to charities bidding for donations. Our results show that it is indeed possible to improve the income of the sector by *at least* 50% on average. We further show how the application of *proxy bidding* allows to maintain a significant portion of the welfare improvements without the need of many bids. Our results pave the way for a novel and more profitable model of distribution of donated goods.

ACM Reference Format:
Christopher Culley, Ji Qi, and Carmine Ventre. 2019. How to Get the Most from Goods Donated to Charities. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

In the past few years, the charity sector has taken an even greater societal role in the attempt to fill in the gaps for recession and financial difficulties. According to the UK Charity Commission data [3], the total number of charities has steadily increased in the past few years, their total annual income currently being in the tens of billion pounds. Nevertheless, the financial statements of charities, see, e.g., [2, 6], report the lower donation of goods from the public as one of the biggest *financial risks* for their organizations; this is confirmed by a worrying decrease in stock levels in recent years.[1] Given the current economic climate, the only prospect for a better distribution will arise from new economic models of donations.

Optimization techniques and game-theoretic reasoning can profoundly impact this "market", as this problem can be modelled as a Combinatorial Auction (CA) where the objective is the maximization of the *social welfare* – the sum of charities' "valuations" (income) for the goods. As a result, there is the need to use algorithms that are *incentive-compatible* even in absence of monetary transfers. In this work we study the feasibility to connect the risks that charitable organizations currently face with the theoretical advances on incentive-compatible CAs given in [5]. The objective is to show that this novel system is beneficial or not and supports a call to move away from the status quo.

[1]These observations also apply to other countries, such as, the US – cf., e.g., goo.gl/8eh5uq.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

Algorithm 1: Greedy

Let l denote the number of different bids, $l = nk$.
Let $b_1 \geq b_2 \geq \cdots \geq b_l$ be the non-zero bids and $S_1, \ldots, S_l$ be the corresponding sets.
For each $j = 1, \ldots, l$ let $\beta(j) \in \{1, \ldots, n\}$ be the bidder bidding b_j for the set S_j.
$\mathcal{P} := \emptyset$, $\mathcal{B} := \emptyset$.
For $i = 1, \ldots, l$ do
  If $\beta(i) \notin \mathcal{B} \wedge S_i \cap S = \emptyset$ for all S in $\mathcal{P}$ then
    (a) $\mathcal{P} := \mathcal{P} \cup \{S_i\}$, and (b) $\mathcal{B} := \mathcal{B} \cup \beta(i)$.
Return $\mathcal{P}$.

Our contribution. We envision a web portal wherein people submit their good donations; incentive-compatible CAs (including those of [5]) are run periodically to allocate goods to charities. Before an auction is run, charities bid for the (bundles of) goods they are interested in. We experimentally evaluate the effectiveness of such a portal in the context of UK charities. In absence of hard data, we use an indirect approach to generate synthetic data based on current trading patterns in and income of the charity sector. We entertain the idea that a centralised system could increase the level of donations from the public and also consider higher *stock values* in our experiments. To mitigate the effects of synthetic data, we generated several thousand random instances for combinations of t (t is the minimum number of items we wish to divide universe of donated goods into), value of donations, and random valuations and evaluate how better off charities are on average. The aim is that of testing the system and ascertain whether the theoretical results can translate into practical gains. The best results are achieved by Algorithm 1. Even at current donation levels, the sector is on average better off by more than 50%. The increase of welfare is slightly less than linear with respect to increase of donations from the public, reaching a value of about 220% for tripled donations. We complete this study by showing how to leverage proxy bidding [1] to significantly reduce the data collection from charities. The main message of this work is that a novel, profitable and practical way to allocate donated goods to charities is possible. Truthful auctions need not be solely an object of theoretical study but can significantly impact upon society.

2 IMPROVING THE WELFARE OF CHARITIES

We aimed at establishing which of the algorithms in [5] is better for so-called k-minded charities. These are charities interested in obtaining one bundle of goods out of k possible subsets of donated goods. We here report the results on greedy (Algorithm 1) and Multiplicative Price Update (MPU) [5]. We let $k = 30,000$ and $t = 100$ and present the results in Figure 1 over 20,000 experiments. We

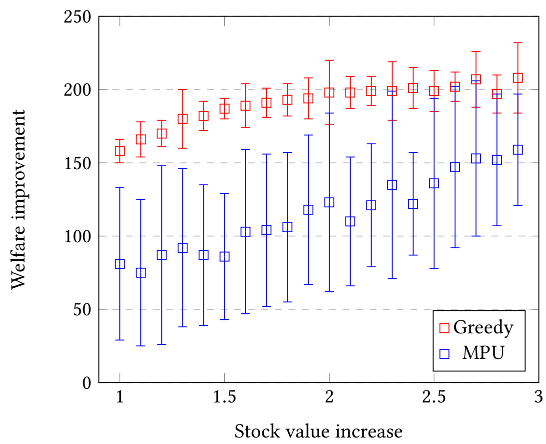

Figure 1: Greedy vs MPU

can conclude that greedy algorithm exhibits a more stable behaviour – steady increase in performances with the the increase of donations and less variance. On average, no matter the stock value level, greedy outperforms MPU significantly. We in fact see that even if a centralised allocation algorithm does not result in an increased donation pot, the competition alone increases the percentage of total value of goods from current levels quite significantly.

3 PROXY BIDDING

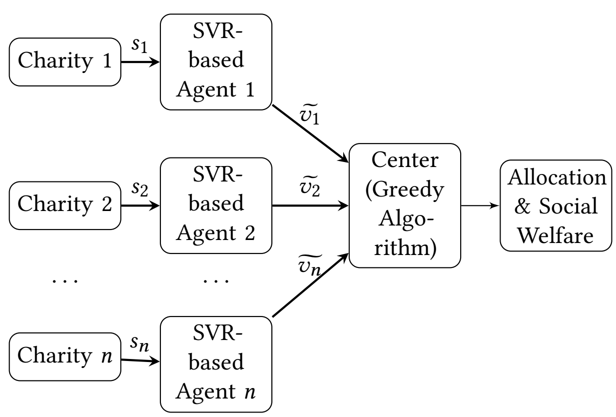

Figure 2: Overall architecture of the proposed ML-based Combinatorial Auction Model.

In this section, we want to consider whether it is possible to retain the significant improvements guaranteed by greedy, while reducing the bidding required from the charities.

In our ML-based Combinatorial Auctions Model, we adopt a Support Vector Regression (SVR) algorithm [4] and choose Gaussian kernel for each agent. This is motivated by the perfomances of this kernel in [1]. As shown in Figure 2, we can simply regard each charity as being equipped with an ML agent and the agent will learn form the bidder to generalize the whole valuation function (denoted $\widetilde{v}_i$) by just eliciting a small number of bids (denoted s_i). The charities

Table 1: Average Social Welfare

k-Minded	Seed (p)				
k = 1000	***0.05***	***0.1***	***0.2***	***0.3***	***1***
SW	1.1202	1.2119	1.2858	1.3274	1.3592
Variance	0.0115	0.0030	0.0042	0.0027	0.0031

[a] p = 1 means no learning is involved.

will then be able to bid quickly. Finally, the greedy algorithm can allocate the goods to each charity according to the predicted bids $\widetilde{v}_i$. Two sets of experiments are conducted with SVR-based proxy agents and without. To maintain experiments manageable, we set $k = 1,000$ and produce $1,000$ random instances with $t = 100$. For each of these instances, we test p at $\{0.05, 0.1, 0.2, 0.3\}$ and run 150 experiments for each value of p. We can then get the average social welfare that the greedy algorithm achieved on the $\widetilde{v}_i$'s. The results are shown in Table 1 and the percentage of social welfare for the different 150 experiments for each p are plotted in Figure 3.

The results in Table 1 draw a very positive picture. Even when $p = 0.05$, the worst case among these settings, the welfare can still be 1.12016 on average, which is improving the current welfare by around 12%. Figure 3 gives us a full view of the experiments.

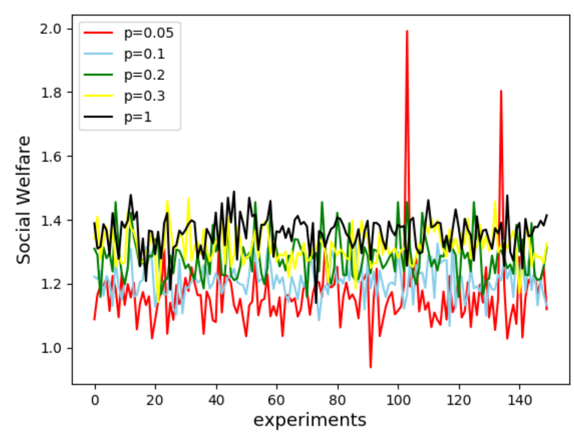

Figure 3: Social welfare with different p

4 CONCLUSIONS

We have presented results for incentive-compatible centralized algorithms to distribute charitable goods under realistic assumptions. We have seen that even without an (with an important, resp.) increase in the value of donations the greedy algorithm results in an increase in the average value of goods a k-minded (additive, resp.) charity gains in a year.

Our results point in the direction of a meaningful use of AI to substitute humans in the bidding phase. Overall our results draw a promising picture and build the foundations for a successful experimentation on real data from charities.

REFERENCES

[1] Gianluca Brero, Benjamin Lubin, and Sven Seuken. Probably approximately efficient combinatorial auctions via machine learning. In *AAAI*, pages 397–405, 2017.

[2] British Heart Foundation. Annual report and accounts 2016, 2016. Available at 'https://www.bhf.org.uk/informationsupport/publications/about-bhf/annual-report-and-accounts-2016'.

[3] Charity Commission. Sector facts and figures, 2017. Available at 'http://www.charitycommission.gov.uk/about-charities/sector-facts-and-figures/'.

[4] Harris Drucker, Christopher JC Burges, Linda Kaufman, Alex J Smola, and Vladimir Vapnik. Support vector regression machines. In *Advances in neural information processing systems*, pages 155–161, 1997.

[5] Dimitris Fotakis, Piotr Krysta, and Carmine Ventre. Combinatorial auctions without money. volume 77, pages 756–785, 2017.

[6] Oxfam. Annual report 2016-17, 2017. Available at 'http://www.oxfam.org/en/about/annual-reports'.