

## **A conversation with Professor S. Nageeb Ali, May 9, 2016**

### **Participants**

- S. Nageeb Ali – Associate Professor of Economics, Pennsylvania State University
- Holden Karnofsky – Co-Founder and Co-Executive Director, GiveWell

**Note:** These notes were compiled by GiveWell and give an overview of the major points made by Professor Ali.

### **Summary**

GiveWell spoke with Professor S. Nageeb Ali of Pennsylvania State University about game theoretic approaches to donor coordination. Conversation topics included behavioral models of altruism, economic models of donor coordination, and other people to talk to about this subject.

### **Behavioral models of altruism**

Models of donation behavior rely on particular models of altruism. Models of altruism include consequentialism, "warm glow" giving, and social signaling.

#### **Consequentialism**

A consequentialist model of altruism focuses on the outcomes that one can change.

#### **"Warm glow" giving**

A lot of literature over the last 30 years has argued that people do not necessarily take a purely consequentialist approach to giving; rather, many people experience a "warm glow" as a result of personally contributing to a cause.

#### **Social signaling**

In the past decade, a growing body of economic literature has suggested that people may donate to charity partly because they want to be perceived as altruistic. This literature suggests that it may be productive to break down a group of large contributors into small teams of people who regularly donate to similar charities, in order to better coordinate and motivate their continued giving. This strategy would prevent GiveWell from needing to coordinate with a large number of individuals.

### **Economic models of donor coordination**

#### **"Funging"**

If Good Ventures were to "funge" individuals' donations to GiveWell-recommended charities (i.e. fill the funding gaps left after individuals have donated), donors would have a disincentive to give to these charities. This could lead individual donors to give to less effective charities, which would reduce the impact of their donations.

It might not become an issue for a couple years, but eventually people would realize their donations don't matter. This is actually consistent with experimental game theory as well: when there is a change to some institution that affects people's incentives; those issues doesn't crop up immediately, they take some time, but over time, even in contribution experiments or things of that sort, people start converging to what the equilibrium would predict.

A good reference for learning about this topic is Fudenberg and Levine (1997): "Measuring Players' Losses in Experimental Games," Quarterly Journal of Economics.

### **Matching**

Matching grants are frequently used to encourage people to donate more to a charity than they would have otherwise. Economic literature has explored (but has not yet definitively answered) questions including:

- Whether an individual who makes a donation knowing that it will be matched feels a greater sense of "warm glow" than if it were not matched.
- Whether matching donations incentivize people to give half of what they would have otherwise donated.

Matching does not seem like a very good strategy for GiveWell because:

- It would tie Good Ventures' funding decisions to individual donors' decisions, which are not necessarily in line with Good Ventures' preferences.
- Matching grants typically do not give donors a choice of which charity to donate to. Since GiveWell has several recommended charities, it is not clear that matching is a good strategy.

### **Splitting**

"Splitting" is GiveWell's term for a strategy whereby one funder fills what they determine to be their "fair share" of a charity's funding gap. This is the strategy GiveWell used in its 2015 giving recommendations.

Splitting is not a common economic strategy, and Professor Ali was not familiar with the concept before talking to GiveWell. However, he believes that this is a defensible solution to the problem of GiveWell's coordination with donors because:

- Splitting solves the coordination problem without requiring direct communication with each individual donor.
- It allows GiveWell to avoid the problems associated with funging (e.g. giving donors a disincentive to donate to GiveWell's top charities).
- It creates a sequential coordination game, in which GiveWell makes the first move and individual donors follow. This is preferable to simultaneous coordination games, which involve a group of people trying to accomplish a common goal without being able to communicate with each other.

- It coordinates donors by setting a tangible goal for how much funding individuals donors should contribute. Matching grants, by comparison, do not provide such a tangible goal.

Coordination games (which involve a group of people working together toward a common objective) are often complicated by a lack of centralized information. This is not a problem in this case, because GiveWell provides a source of information that donors can use to guide their individual giving decisions toward the objective of doing the most good. Instead, much of the challenge comes from the impracticality of direct negotiation and coordination between all of the parties.

Part of the reason splitting hasn't been researched much is because much of the research is motivated by seeing what's done in practice with respect to fundraising, and a lot of what's done is matching. People have thought less about coordination with a large group.

### *1/n rule*

In cases where individuals are working together to meet an objective but cannot coordinate with each other, the "1/n rule" refers to a strategy wherein each individual addresses 1/n of a problem, where  $n$  is the number of people working toward the objective. This is an appealing solution because while it's not possible to know what other people are doing, the objective would be met if everyone followed this rule. If everyone were to follow a different rule (e.g. contributing more or less than 1/n), the total contributions would be more or less than what is necessary.

Economics research typically only addresses coordination problems in cases where the "game" is not related to the design of a mechanism, as it is for GiveWell.

*Upon reviewing these notes for publication, Professor Ali wrote the following: "Let me re-phrase it in the following way: when people have thought about the design of mechanisms, the fundamental issue that people think about is private information, and not necessarily coordination, because the assumption is that the mechanism designer can suggest how people should behave in an incentive-compatible way."*

*It is the case that economists do think about coordination issues, but in settings where the mechanism / game is exogenous and not part of design. Since our conversation, I have realized that there is a branch of the literature that thinks a bit about coordination from the standpoint of what can be implemented across all equilibria—i.e., a context in which the person designing the mechanism cannot suggest what people do—but much of that analysis is often done assuming all participants share the same information.*

*I think in terms of the issues of coordination that GiveWell is considering, there isn't work that has been done, and I think studying these issues theoretically might be interesting and useful."*

### **Setting targets for individual donations**

Another approach to the coordination problem could be to announce a target amount of money that GiveWell hopes to raise from individual donors without announcing how much it has recommended that Good Ventures donate. (GiveWell finds this approach problematic because it conflicts with GiveWell's commitment to transparency.)

### **Combined splitting and funging approach**

GiveWell could strike a compromise between splitting and funging by:

1. Recommending that Good Ventures contribute 50% of each charity's funding gap.
2. Waiting to see how much individuals donate within a certain amount of time.
3. Recommending that Good Ventures partially fund any funding gaps that are left over for any of the top charities (for example, filling 50% of the remaining funding gap, or alternatively doing something like giving enough to ensure that at least 80% of each overall funding gap is filled).

This model balances the pros and cons of full funging (which ensures that funding gaps are filled but disincentivizes individual donations) and full splitting (which incentivizes individual donations but could result in large funding gaps for top charities). The combined model gives individuals a reasonable incentive to donate and carries a more moderate risk of funding gaps remaining.

In this model, an individual's donation may or may not be consequential, depending on the behavior of other donors.

This model would give GiveWell the flexibility to determine ex post how much of the remaining funding gaps to fund, using up-to-date information about which charities seem to offer the highest marginal returns. The amount to be funded could be determined using an algorithm that GiveWell would make public, in accordance with its commitment to transparency.

### **Other people to talk to**

- Steven Matthews – Professor of Economics, University of Pennsylvania
- James Andreoni – Professor, Department of Economics, University of California, San Diego
- John List – Homer J. Livingston Distinguished Service Professor in Economics and the College and Chairman of the Department of Economics, University of Chicago

*All GiveWell conversations are available at <http://www.givewell.org/conversations>*