Founders Pledge

Giving Multipliers

CAUSE AREA SUMMARY

KOJI FLYNN-DO

AUGUST 2022

Contents

I. Introduction

The framework described in this report was developed in large part by Sjir Hoeijmakers, who was previously a Senior Researcher at Founders Pledge.

We've written before about how the very best interventions and organizations are often hundreds or thousands of times more impactful than the median (Our Approach to Charity). On that basis, we focus our research on identifying and funding the highest-impact efforts to make progress against the world's most pressing problems.

We now think that it is possible to do even better than this by supporting programs that drive more money to the highest-impact funding opportunities available. We call these efforts and organizations ***giving multipliers*****: funding opportunities which, as their main output, cause more money to be given to high-impact interventions and organizations.**

In one sense, this is intuitive: if it is impactful to allocate funding to a given organization, then it should be even more impactful to allocate that funding to a giving multiplier which then drives, say, twice the amount to the original organization of interest. However, the fundraising, rather than implementing, nature of these organizations raises unique questions in their evaluation.

In this investigation, we are only evaluating these organizations on the basis of their giving multiplier. That is to say that we are not evaluating the non-monetary externalities of these organizations, which in some cases are extremely significant. If an organization mainly prioritizes its research output but does some fundraising, it could easily be a stellar all-things-considered funding opportunity despite not being recommended in this investigation.

This also means that non-monetary multiplier-like outputs are outside the scope of this investigation. For example, organizations like 80,000 Hours aim to help create more human capital aimed at solving pressing problems and can be thought of as multipliers of a sort. Charity Entrepreneurship aims to help create more organizational capacity by researching and incubating new charities to address gaps in current efforts to solve various problems. Similarly, research that creates knowledge to multiply the impact of other efforts can also sometimes be thought of as a multiplier. These and other multiplier-like efforts are not evaluated in this investigation.

Ia. A Note on False Precision

Note that the results of the calculations described here should be understood as quick, back-of-the-envelope estimates with considerable uncertainty rather than sophisticated analyses that can be taken at face-value. These preliminary estimates are extremely rough, and could easily be off by a factor of 2 to 3. Here, we aim only to estimate the net multipliers in

extremely broad strokes, which we believe will differentiate those giving multipliers that are clearly exceptional from the rest.

Our funding recommendations do not distinguish between, say, a multiplier of 10x vs 12x, simply because the figures are so unstable. We do think that we can say that there is a real difference between, for example, 15x and 3x, but this should not be read as claiming confidence that the true difference between those is 5x.

To emphasize this, we report only the broad categories under which the giving multipliers we evaluated fall rather than our precise numerical estimates, which we believe are misleading in their implied precision.

II. How We Evaluate Giving Multipliers

IIa. The Community Model

Funding Displacement and Overall Cost-Effectiveness

When a funder chooses to allocate funding to a given organization, this can sometimes have the effect of increasing that organization's funding by the donated amount. On the other hand, it could simply displace other sources of funding that would have been donated to the same organization. This is called **funging**, and it is especially important in the context of giving multipliers.

For example, if Cheerful Charity can productively absorb $100, and Brilliant Benefactor will completely fill this room for funding, then Gracious Giver donating $50 to Cheerful Charity has the effect of giving Brilliant Benefactor $50 to use in other ways. Cheerful Charity will receive the same amount of funding either way. Consequently (and perhaps counterintuitively), the impact of Gracious Giver's giving is determined not by how Cheerful Charity will use it, but by how Brilliant Benefactor will allocate their remaining $50.

In the context of funding giving multipliers, there is an additional step. If we allocate $10 to a giving multiplier and that causes $20 to be donated to an implementing organization, that *could* simply result in the implementing organization having an additional $20. On the other hand, it could displace other funding in the way described above. Moreover, that $20 would have been used in some other way — perhaps donated to some other organization, perhaps saved, or perhaps spent on personal consumption. All these effects need to be considered.

Accounting for Funding Displacement with a Unitary Community Model

In this investigation, we consider a community model in which there is a collective of funders who are all interested in maximizing impact and agree on the cost-effectiveness of all their

available funding opportunities. Therefore, they act as a **unitary agent** with a shared budget to allocate across available funding opportunities. In this model, the community will allocate its money such that all the funding opportunities with the highest cost-effectiveness are funded. The "cost-effectiveness bar to clear" (henceforth **the community bar**) for a funding opportunity at a given time is then the cost-effectiveness of the most cost-effective opportunity still available to fund.

If some other outside funding stream allocates money to an organization that clears the community bar, then the community can allocate the displaced funding to the most cost-effective opportunity still available to fund. That is, the community allocates funding with cost-effectiveness equal to the community bar (the cost-effectiveness of the most cost-effective opportunity still available to fund). **Consequently, any outside funding streams to funding opportunities above the community bar have cost-effectiveness equal to the community bar.** This is discussed in more depth and with more precision in Appendix A: Modeling the Effective Giving Community as a Unitary Agent.

Of course, this model is a large simplification, though we think that it approximates the current effective giving community's approach. That community includes organizations like Open Philanthropy, GiveWell, Effective Altruism Funds, Founders Pledge, and others, as well as numerous individual donors and philanthropists.

Applying the Unitary Community Model to Giving Multipliers

There are at least two useful implications from this model, in which any outside funding streams to funding opportunities above the community bar have cost-effectiveness *equal* to the community bar.

First, if a giving multiplier causes at least more than one “quality-adjusted dollar” (defined below in *The Quality Factor*) to be allocated to implementing organizations per dollar spent operating the giving multiplier charity, then the giving multiplier itself clears the community bar. Second, any additional funding generated by giving multipliers has cost-effectiveness equal to, at most, the community bar. However, a giving multiplier may also cause additional funding to be allocated to funding opportunities that do not clear the community bar.

Taken together, this implies that we can think about giving multipliers in terms of how many quality-adjusted dollars they raise per dollar spent when we evaluate giving multipliers as prospective funding opportunities. There is no need to individually calculate the cost-effectiveness of each funding opportunity to which a giving multiplier causes additional money to be donated.

Limitations of the Unitary Community Model

Though the unitary community model is a helpful simplification, there are several unrealistic assumptions required in its application.

First, we assume that all members in the collective of impact-maximizing funders agree on the cost-effectiveness of all available funding opportunities, and that they all have perfect information and access to the same funding opportunities. In reality, neither of these assumptions are true. Despite agreement on certain fundamental principles, impact-maximizing funders vary in their assessments of cost-effectiveness due to disagreements over practical and philosophical questions. Moreover, some funding opportunities may only be accessible to some funders, and others may be sensitive and hence confidential. As such, there is no true community bar. Instead, the bar for being funded is something more like being appealing enough to be funded by at least one funder.

Second, funging is much less mechanical than the precise one-for-one displacement effect described above, wherein one dollar allocated to a funding opportunity that clears the bar mechanically leads to one dollar being allocated to the next most cost-effective opportunity. Though this is a real effect that generally holds, funding streams can variously “crowd in” or “crowd out” other funding, or simply have no displacement effect. For example, some funders are hesitant to fund the entire budget of an organization for fear of distortionary incentives. Other times, a small organization might grow larger and absorb more money if given sufficient seed funding. So in those instances, one might crowd in funding. In other cases, imperfect information or unresponsive funders can mean that a funding opportunity receives more funding than it can usefully absorb, with no displacement to more cost-effective opportunities.

Finally, in our analysis we do not account for how the money generated by giving multipliers from outside the impact-maximizing community would have otherwise been used. We assume that the money raised by multiplier opportunities — which must come from somewhere — would otherwise be used with negligible impact; either by being donated to significantly less cost-effective opportunities or by not being donated at all.

IIb. The Three-Factor Model

The best giving multipliers are those that achieve high "returns" — in the sense of money moved to impactful funding opportunities — relative to their costs. Abstractly, the multiplier is nothing more than the total counterfactual revenue divided by the total costs. We break this into three factors to create a simple multiplicative model.

We call the ratio of money moved to money spent the **gross multiplier**. One must also account for the opportunity cost of labor. Though the wages of employees at giving multipliers are counted in operating expenses, those employees could also earn money themselves in order to give it away (or do some other impactful work). This opportunity cost should be included in the overall costs of the multiplier opportunity. We call this the **labor factor**. Finally, we are interested not only in the gross volume of money moved, but also in the quality of the recipient funding opportunities: a dollar allocated to an excellent organization can be worth hundreds of dollars allocated to a mediocre one. This is captured by the **quality factor**.

Putting this together, we have our three-factor model:

$$\text{Net Multiplier} = \text{Gross Multiplier} \times \text{Labor Factor} \times \text{Quality Factor}.$$

When evaluating giving multipliers as funding opportunities, we make a ballpark estimate of an organization's expected net multiplier over the next three years based on their track record in the recent past (including e.g these three factors so far) and projections about their future plans.

The Gross Multiplier

The **gross multiplier** is calculated by taking the gross volume of counterfactual donations generated by a giving multiplier and dividing by the operating expenses. (For more on what we mean by counterfactual in this context, see *Appendix B: How We Think About Counterfactuals for Giving Multipliers*.)

$$\text{Gross Multiplier} = \frac{\text{Total Money Moved} \times \text{Prop. Counterfactual}}{\text{Operating Expenses}}.$$

So, if an example giving multiplier counterfactually generated $100 at an operating cost of $10, the gross multiplier would be $100/$10=10.

Because giving multipliers themselves are high-impact funding opportunities (assuming they have a net giving multiplier above 1x), if an organization counterfactually fundraises for its own expenses from sources outside of the impact-maximizing giving community, we include that in the Total Money Moved term. For example, we would *not* include funds raised from Founders Pledge members in Total Money Moved, but we *would* include funds raised from the general public.

Approach for Pledging vs. Fundraising Organizations

Some of the organizations evaluated in this investigation aim for their fundraising efforts to translate into donations in days, weeks, or months. These organizations are engaged in fairly straightforward fundraising, and can be evaluated as such. Effektiv Spenden, The Life You Can Save, and Effective Altruism Australia are examples of this model. Other organizations have longer time-horizons, in that they aim to recruit recurring donations via pledges to give some proportion of one's income or wealth yearly. Giving What We Can, One for the World, and the Jewish Effective Giving Initiative are examples of this model.

These two models call for different approaches. In the case of fundraising organizations, we can more easily calculate their gross multiplier by simply dividing their counterfactual money moved by their expenses.

In the case of pledging organizations, we instead calculate the Net Present Value of a given year's cohort donations over five years (so, the pledge year and the four following years), then divide by the operating expenses of that year, plus the expected costs of ongoing pledger stewardship over that same timeframe. Different organizations will spend different proportions of their expenses on stewardship, growth, etc., so we calculate the costs of ongoing stewardship separately for each organization. The five-year time-cutoff is somewhat arbitrary; evaluating until the end of time is impractical, and excluding future years' costs and donations altogether misses out on the main function of pledging organizations.

The Labor Factor

The **labor factor** attempts to account for the significant opportunity costs associated with staff time, in addition to the financial costs of the giving multiplier. If a staff member is trying to maximize the impact of their work, these opportunity costs are plausibly in the same order of magnitude as the staff member's salary, and potentially (much) higher. Suppose, for example, that such a staff member could instead earn-to-give twice as much as their current wages, then give away 50% of their salary to high-impact funding opportunities. In such a case, they would already approximately be "earning back" their salary.

In this investigation, we use having taken the Giving What We Can pledge as a proxy for impact-maximizing staff time, as that indicates an interest and commitment to effective giving.

The labor factor is given by:

$$\text{Labor Factor} = \frac{1}{1 + \text{Prop. GWWC} \times \alpha},$$

where α is the multiple by which impact-maximizing staff members could out-earn their current wages. (*Appendix C: More Precision on the Labor Factor*, explains why the labor factor takes this form.)

We determine this factor by considering the general earnings potential of people in the impact-maximizing community, as well as specific attributes of those working at organizations devoted to fundraising or outreach on effective giving.

As outliers like Dustin Moskovitz have demonstrated, potential earnings are much higher than twice a staff member's current salary. (Bloomberg lists Moskovitz with a net worth of $11 billion at the time of writing.) Moreover, some in the Effective Altruism community have argued that working directly on pressing problems can be significantly more impactful than earning-to-give. On the other hand, staff have self-selected into their current occupations. If they were interested in maximizing the impact of their work and chose to work at a giving multiplier, then perhaps their current roles are the most impactful careers for these staff, due to personal skillset or other factors.

We therefore have considerable uncertainty surrounding the opportunity cost, but believe that it is prudent to account for an average ~2x salary equivalent of opportunity costs for impact-maximizing staff members, with large uncertainty.

Returning to the same example giving multiplier as above, suppose that seven of the ten employees had taken the Giving What We Can Pledge and we use $\alpha = 2$. Then, the labor factor would be:

$$\frac{1}{1 + \frac{7}{10} \times 2} \approx 0.42.$$

A more sophisticated analysis would attempt to estimate the multiple by which impact-maximizing staff members could out-earn their current wages for each organization being evaluated. In this investigation, we have applied $\alpha = 2$ for all organizations.

The Quality Factor

The **quality factor** is calculated by finding the average quality weight of the money moved. We "bucket" the recipient funding opportunities that receive money through giving multipliers into five weights: 1x, 0.5x, 0.2x, 0.1x, and 0x. These are based on our previous research evaluations and, where these are not available for a given funding opportunity, rough heuristics derived from our current view on the research quality underlying the recommendation. By default, we rate organizations recommended by Animal Charity Evaluators but not by Founders Pledge at 0.2x, previous GiveWell standout charities at 0.1x, and organizations recommended by charity evaluators with lower bars for recommendations at 0.05x.

At Founders Pledge, we separate our research and evaluation into three worldviews and the weights assigned to different funding opportunities are therefore worldview specific. We set the

community bar for each worldview accordingly. See discussion in *The Community Mode*l above for what we mean by the "community bar".

- In the current generations worldview, we take the community cost-effectiveness bar to be roughly 5x that of a donation to GiveDirectly (similar to GiveWell).
- In the longtermist worldview (including climate change), we take the community cost-effectiveness bar to be roughly a current donation to the Nuclear Threat Initiative's biosecurity programs.
- In the animal welfare worldview, we take the community cost-effectiveness bar to be roughly a donation to The Humane League.

$$\text{Quality Factor} = \frac{\sum \$ \text{ to FO}_i \times \text{Quality Weight of FO}_i}{\text{Counterfactual Money Moved}}.$$

So, if the example giving multiplier from before directed $30 to programs rated at 1, $50 to programs rated at 0.5, and $20 to programs rated at 0.1, the quality factor would be:

$$\frac{\$30 \times 1 + \$50 \times 0.5 + \$20 \times 0.1}{\$30 + \$50 + \$20} = \frac{\$57}{\$100} = 0.57.$$

III. Putting It All Together: The (Expected) Net Multiplier

For our example giving multiplier, combining the three factors as above would yield $\text{Net Multiplier} = 10 \times 0.42 \times 0.57 = 2.394$, which is to say that the final **net multiplier**, after adjusting for the quality and labor factors, is 2.394. Based on their plans, projections, and performance, we would then estimate their yearly *expected* net multiplier over the next three years, which is the ultimate term of interest.

Appendices

Appendix A: Modeling the Effective Giving Community as a Unitary Agent

Here, we discuss the community model in more detail.

There is a collective of funders who want to maximize impact by allocating their capital to the most cost-effective funding opportunities available to them. Each member of the collective shares the same information of the cost-effectiveness for each funding opportunity. Consequently, we model them as a unitary agent.

Each funding opportunity in the set of funding opportunities, $i \in F$, has an associated cost-effectiveness function, $c_i(x)$. Each cost-effectiveness function $c_i(x)$ is a function of money spent on funding opportunity i, denoted x_i. Each $c_i(x_i)$ is continuous and exhibits diminishing marginal returns.

The collective of funders has some budget X to spend in total across all $i \in F$. They will allocate the money to maximize cost-effectiveness, such that each funding opportunity i is allocated $x_i \geq 0$. So we have, associated with funding opportunity i: an amount of money spent x_i and a cost-effectiveness function $c_i(x_i)$.

The community will spend all the money above a certain expected cost-effectiveness. At some time t, all the funding opportunities with the highest cost-effectiveness will be funded, up to the point where all of X has been spent.

In general, the "cost-effectiveness bar to clear", B, for funding some opportunity $i \in F$ at some time t would be the cost-effectiveness of *the most* cost-effective opportunity still available $B_t = \max_{i \in F}(c_i(x_i))$. If, at time $t+1$, more capital $\tilde{X}$ was acquired by the community, then the bar at time $t+1$, denoted B_{t+1}, would be the cost-effectiveness c_i associated with $i \in F$ such that $B_{t+1} = \max_{f \in F}(c_i(x_i))$ at time $t+1$.

Example 1: Exogenous Giving Within the Funded Portfolio

Suppose that at time $t-1$ (i.e. prior to the community allocating its capital at time t), someone else spends some amount of money *less* than the amount the community would have allocated, $y_i < x_i$, on any i that the community *would* have funded with x_i dollars.

Then, the community at time t would adjust its allocation accordingly to fund i at $z_i = x_i - y_i$. The surplus of its budget would be spent on the next best thing (that is, the funding opportunity with the highest marginal cost-effectiveness) that is unfunded at time t. Denote that cost-effectiveness $\tilde{B}_t = \max_{i \in F}(c_i(x_i))$ at time t.

On a first glance, one might expect that allocating y_i to i would have the effect of $y_i \times c_i(y_i)$ (dollars times "impact" per dollars, so units of "impact"). But in this model, the overall effect of the allocation of y_i to i would then be $y_i \times \tilde{B}_t$, as opposed to $y_i \times c_i(y_i)$.

Non-Constant Cost-Effectiveness

In the above, we are assuming that c_i and $\tilde{B}_t$ are approximately constant for simplicity. That is, we are assuming that the cost-effectiveness is not diminishing very quickly as more funding is allocated. Sometimes that will not be true, in which case we would have that the the effect is $\int_d^{d+x_i} c_i(x)dx,$ where d is the starting amount of funding allocated.

If, as a result of allocating funding, the cost-effectiveness function c_i of funding opportunity i decreases enough that it ceases to be the maximum, such that another funding opportunity j has a higher cost-effectiveness c_j, then funding would begin to be allocated to j. If c_j then ceases to be the maximum cost-effectiveness, then the funding would start being allocated to the next, and so on. So if y_i gets split into several funding opportunities, then the total effect would just be the sum of the integrals.

For the rest of this Appendix, we'll assume that the cost-effectiveness bar is approximately constant to simplify the exposition.

Defining Giving Multipliers

Define giving multipliers, or multipliers, to be a subset of funding opportunities $m_i \in M \subset F$ that, as their main effects, cause money to be moved to other funding opportunities. That is, for some amount of money u_i allocated to m_i, some other amount of money v_j is generated and allocated to another funding opportunity f_j.

We will call the funding opportunity to which money ultimately flows the *recipient*. In the preceding paragraph, the recipient was f_j.

How much more money?

We will call the ratio of money generated by a multiplier and money allocated to that multiplier $M_i = \frac{v_j}{u_i}$ the *marginal multiplier* of that multiplier opportunity.

The cost-effectiveness of multipliers

To evaluate the cost-effectiveness of a giving multiplier as a funding opportunity, we take the impact generated by allocating v_j to f_j and divide by the cost, u_i. Therefore, the cost-effectiveness of giving to m_i is $c_i(u_i) = \frac{v_j c_j(v_j)}{u_i}$.

We can write the cost-effectiveness of f_j as the bar for funding multiplied by some factor: rB. So, for example, if the multiplier directs money to a funding opportunity that is twice as cost-effective as the community bar, then we would get that $c_j(v_j) = 2B$. We could also write the cost-effectiveness of giving to m_i as $\frac{v_j 2B}{u_i}$.

Implications: Which multipliers are worth funding?

For a multiplier opportunity to be worth funding from the perspective of the collective, it must be the case that the all-things-considered cost-effectiveness clears the community bar. That is, a multiplier opportunity m_i must have $c_i(u_i) = \frac{v_j c_j(v_j)}{u_i} \geq B$.

In this model, the collective will fund:

- any multipliers with recipients that clear the bar and that don't lose money ($r \geq 1, M \geq 1$)
- any multipliers with recipients that don't clear the bar but that have a higher marginal multiplier than the reciprocal of the cost-effectiveness factor ($M \geq \frac{1}{r}$).

Three Cases

The multiplier opportunity can direct money toward funding opportunities with cost-effectiveness that clears the bar, does not clear the bar, or *constitutes* the bar.

Some amount u is allocated to m, causing v to be generated and allocated to f. The cost-effectiveness of m_i, called c_i, depends on the values of the *marginal multiplier* M and the *cost-effectiveness factor* r.

Case One: The recipient *is* the current bar. $r = 1, rB = B$

There are no further effects, so the cost-effectiveness is given: $c(u) = \frac{v}{u} rB = MB$.

If $M \geq 1$, then this multiplier is worth funding.

Case Two: The recipient clears the bar. $r > 1$

This is structurally similar to *Example 1: Exogenous Giving Within the Funded Portfolio*.

Here, the expected cost-effectiveness of the recipient is *above* the bar; $r > 1 \rightarrow rB > B$. That means i would have been allocated some amount x by the community, and so (assuming that $v < x$), the community will respond by allocating $z = x - v$ to i and v to the funding opportunity with the highest marginal cost-effectiveness that is unfunded, which (as before), has cost-effectiveness B.

Consequently, the actual effect of allocating funding v to f is *not* that of allocating v with cost-effectiveness rB, but instead is given, as in Case One, by: $c(u) = \frac{v}{u} B = MB$.

If $M \geq 1$, then this multiplier is worth funding.

Case Three: The recipient does not clear the bar. $r < 1$

In this case, the specific value of r matters. The cost-effectiveness is given: $c(u) = \frac{v}{u} rB = MrB$.

For this to be worth funding, it must be that $rM \geq 1$, or equivalently: $M \geq \frac{1}{r}$. That is, it must be the case that the amount of money generated by m is sufficient to offset the reduced cost-effectiveness of the recipient.

Appendix B: How We Think About Counterfactuals for Giving Multipliers

Because of our interest in the impact-maximizing funding community, we evaluate counterfactuality at the community level in this investigation. **We call donations 'counterfactual' when they would not have been generated in the scenario where the giving multiplier in question did not exist**. We do *not* account for the possibility that some other giving multiplier organization would take the place of the giving multiplier in question because the community-level costs would likely be roughly equivalent. Of course, it could still be the case that the replacement-level giving multiplier would do a slightly better or worse job than the actually-existing giving multiplier, but we omit this consideration for simplicity.

We evaluate counterfactuals from the perspective of the impact-maximizing collective. From this perspective, a giving multiplier is simply an organization that the community funds in order to fundraise from some specific audience. To illustrate, suppose we were in a scenario where (actually existing) giving multiplier Observable Organization did not exist and, consequently, giving multiplier Emergent Establishment were created to fundraise from a similar audience. In that case, the overall multiplier generated would very likely be quite similar, and so we would call the donations generated by Observable Organization counterfactual. If Emergent Establishment came about in the counterfactual, that would still cost the community operating expenses. In this investigation, we consider the overall costs and returns to the collective of funders.

On the other hand, sometimes giving multipliers are attempting to raise funds from overlapping audiences. Suppose that two giving multiplier organizations, First Fundraiser and Second Solicitor, are both fundraising from the same audience and that hearing from either of them is necessary and sufficient for a member of that audience to donate. First Fundraiser was the first to be in touch with all members of this audience. Then, donations from that audience through Second Solicitor would have been generated by First Fundraiser even if Second Solicitor did not exist. Hence, we would not call donations through Second Solicitor counterfactual.

In a more realistic setting, if a person is giving via One for the World and would have given through GiveWell anyway, these donations would *not* be counterfactual. If this person would only have donated in the world where a similar organization sprouted up in the absence of One for the World and did similar outreach, then these donations *would* be counterfactual.

Appendix C: More Precision on the Labor Factor

In our three factor model, we calculate the net multiplier as the product of the gross multiplier, the labor factor, and the quality factor. However, the net multiplier is actually just the counterfactual money moved divided by total expenses. We can represent that as:

$$\text{Net Multiplier} = \frac{\text{Total Money Moved} \times \text{Prop. Counterfactual}}{\text{Total Costs}} \times \text{Quality Factor.}$$

Total costs are made up of the monetary expenses and the opportunity cost (of labor). The opportunity cost of labor is the multiple by which impact-maximizing staff members could out-earn their current wages (α) times their wages. As discussed before, we use having taken the Giving What We Can pledge as a proxy for impact-maximizing staff time, as that indicates an interest and commitment to effective giving. We therefore have, roughly:

$$\text{Opportunity Cost of Labor} = \alpha \times \text{Prop. GWWC} \times \text{Wages}$$

So the net multiplier can be represented:

$$(1)\ \text{Net Multiplier} = \frac{\text{Total Money Moved} \times \text{Prop. Counterfactual}}{\text{Operating Expenses} + \alpha \times \text{Prop. GWWC} \times \text{Wages}} \times \text{Quality Factor.}$$

By contrast, our three factor model is represented:

$$\text{Net Multiplier} = \text{Gross Multiplier} \times \text{Labor Factor} \times \text{Quality Factor}$$

$$(2)\ \text{Net Multiplier} = \frac{\text{Total \$} \times \text{Prop. Cf}}{\text{OpEx}} \times \frac{1}{1 + \text{Prop. GWWC} \times \alpha} \times \text{Quality Factor.}$$

(1) and (2) are equivalent *if* we assume that operating expenses are equal to total wages. While this is not strictly true, most of the organizations evaluated here are focused on outreach and marketing and hence do not have many other expenses beyond labor. Since wages are a subset of all operating expenses, our three factor model (to varying extents, depending on organizational specifics) underestimates the true net multiplier by dividing by a slightly larger number.

For a few organizations, e.g. GiveWell's outreach operations, a large proportion of their operating expenses are devoted to digital advertising. In those cases, we use (1) to calculate the net multiplier.