

Philosophy of Finance Fellowship: Fall 2021
Markets and the Environment

Seminar 4: ESG and Impact Investing
Guest: Christina Alfonso-Ercan

In our fourth seminar, we look at impact investing and its potential to benefit the environment. We read excerpts of two studies on the effectiveness of impact investing, and we are joined by a distinguished guest who built a successful advisory business in the ESG space.

The first reading selection is a study done by two researchers working in non-profit consulting businesses. Hillebrand and Halstead examine the real world impact of impact investing. They identify the principles by which it can be effective, and also insist on including the opportunity costs associated with impact investing. They compare donating outright with impact investing and conclude that donating can have 10x the beneficial effects of impact investing on its own. They also provide good definitions and close with a case study of the Acumen Fund and its claims to have reduced atmospheric carbon. Their key findings are that impact investors must meet two requirements: enterprise impact and “additionality”.

The whole report can be found here <https://lets-fund.org/impact-investing/>

Donating effectively is usually better than Impact Investing

Authors: Hauke Hillebrandt, PhD; John Halstead, DPhil

Executive Summary

Impact investing – investing in, or divesting from, for-profits for the purpose of social impact – is an increasingly popular approach to doing good. It seems to offer the promise of a double bottom line: direct social impact and profits that you can keep or reinvest in other socially beneficial businesses. A donation to charity, in contrast, yields no monetary returns and can only be spent once. In this report, we discuss whether impact investing is indeed a promising approach for people who want to have social impact.

Impact investors face two distinct challenges:

- Investors must find companies with **enterprise impact** – companies that make a positive difference to the world.
- Investors must have **additionality** – they need to make a difference to the performance of those companies, either through providing additional capital (known as *investment impact*) or through providing non-monetary support, such as advice or access to networks.

For both of these challenges, it is crucial to consider the *counterfactual*. That is, we have to ask: what would have happened had we not invested? Will a given solar power company merely displace another near-identical solar power company? Will my capital merely displace another investor? This marks a crucial difference between investing for profit and investing for impact. When investing for profit, we do not need to consider these kinds of questions. If the solar power company I invested in is making a \$100 million profit, it doesn't matter whether an identical solar power company would have sprung up one week later if the company did not exist. And if I made a substantial profit from my investment in the company, the fact that someone else would have acquired those profits had I not done so is irrelevant. When aiming for social impact, however, these questions are fundamental.

When we are deciding whether to impact invest, we must also consider the opportunity cost of impact investing. In the same way, if we want to make a profit, we wouldn't compare the return on our investment to what we would have got *if we had done nothing*. Instead, we would compare our ROI to what we could have done otherwise with the money: if I chose an investment with a 3% return, but another available investment had an 8% return, then I would have made a mistake. The same is true if our aim is to have social impact.

If our aim is to do the most good, there are two alternatives to impact investing:

- **Investing to give** – Investing for profit to donate later to effective charities
- **Donating now** – Donating the money to effective charities now

Having social impact through donations is much more difficult than many people imagine, and it is easy to miss out on huge impact multipliers in philanthropy. However, if done carefully, the social benefits of these alternative approaches can be substantial. Reviews of our recommended giving opportunities are available on our [research page](#).

You can [watch John Halstead give a talk based on this paper on Youtube](#).

Key points

The key findings of this report are:

1. Finding an impactful company is hard

The most promising companies will produce positive externalities or benefit consumers in poor countries, and focus on high-impact cause areas, such as global poverty and health, animal welfare, or climate change. However, evidence suggests that it is difficult to identify in advance which social programmes will work: the path from action to social impact is usually not as you would expect. Socially beneficial businesses have to solve two very difficult optimisation problems simultaneously – turning a profit and having impact. Consequently, finding viable companies with enterprise impact will not be straightforward. Our research suggests that many impact investors seem not to carry out rigorous or analytical impact evaluations.

2. It is hard to have additionality in large public stock markets

Many impact investors try to affect the stock price of companies in public stock markets, either by boosting the stock price of beneficial companies or by damaging the stock price of harmful companies. These efforts are complicated by socially neutral investors (who only seek profit), who can potentially offset any effects on the stock price. For example, if impact investors divest from an industry, socially neutral investors can move in to buy up the underpriced stock. There is clear evidence of short-term market inefficiency such that impact investors can affect stock prices on the timescale of around 3 months. There is expert disagreement about whether socially responsible investing is likely to have an effect after 6 months and beyond: some economists hold that the effect will be completely offset, some that more than half will be offset, and some that a substantial fraction of the effect might persist beyond 6 months.

Given the size of the market cap of firms targeted by socially responsible investing, it will also be difficult for most investors to have any substantial effect on stock prices in the first place. Moreover, if you invest in a socially beneficial company offering market-rate returns, then you will likely merely displace a socially neutral investor. This means the counterfactual impact of your investment is merely to provide additional capital to the stock market as a whole. For all of these reasons, the direct impact of any single socially responsible investor in large public stock markets is likely to be modest at best. All this being said, genuine strict socially responsible investing is undoubtedly more socially impactful than investing solely for personal profit. Even if the direct effects on stock prices are modest, the indirect effects appear to be more substantial. Thus, the arguments here do **not** give license to ignoring divestment movements solely in order to make money.

3. There is more scope for additionality in VC and angel investing

In inefficient markets with fewer investors and with imperfect information, there is more scope for your investment to make a difference to the company's cost of capital. However, finding and exploiting market inefficiency is difficult. Even in VC and angel investing, the risk that your investment merely displaces someone else's remains a fundamental consideration.

4. There is a trade-off between financial returns and social impact. Investors seeking market-rate returns risk merely displacing socially neutral investors. Consequently, impact investors may need to accept lower returns for the sake of additionality. Impact investors also incur additional costs in identifying, evaluating and supporting the businesses they invest in. If you accept lower monetary returns, then you are giving up money that could be donated to effective charities.

5. Your investment might merely displace another impact investor. Even if you accept subpar financial returns, you need to consider the risk that your investment merely displaces another impact investor who is also willing to accept subpar returns.

6. Impact investing has other benefits. Although they appear to have had modest direct effects on stock prices, divestment campaigns might in the past have helped to stigmatise targeted companies and industries, which in turn has helped to change consumer attitudes and encourage restrictive regulation. Owning the stock of a company also gives you some control over how it operates, allowing you to potentially steer it towards socially valuable ends or to prevent mission drift.

This suggests that, for people aiming to have maximal social impact, impact investing is likely to be the best approach only in specific circumstances. Impact investing might be a good option for people who:

- *Work on an important problem that is neglected by other investors
- *Do VC or angel investing
- *Accept financial sacrifice
- *Have an informational advantage over other investors that allows them to reliably identify promising opportunities

A good example of a case fitting the above criteria would be an investment in a company producing a revolutionary meat-alternative product that is on the brink of financial viability but is, for some reason, ignored by other socially neutral or impact investors. However, when the conditions above cannot be satisfied, investing to give or donating now are likely to be a better bet, if done carefully.

The decision about whether to pursue for-profit or non-profit solutions to problems depends on a few factors. For-profits have some advantages over non-profits in that for-profits tend to be more efficient and customer-focused. However, for products that are not yet market viable, such as public goods, non-profits will be more promising. Non-profits also tend to be more neglected because the incentives to support them (i.e. profits) are lacking. Research on effective charities is improving all the time, allowing donors to have truly outstanding impact for their dollar.

We briefly try to gain an impression of the impact investing space by examining an impact evaluation by an impact investing platform that is a field leader in impact

evaluation. Our investigation showed that donations are likely upwards of 10x more impactful than the impact investing platform, and that there are key gaps in the evaluation carried out by the impact investing platform...

Defining key terms

We first need to clarify the key concepts surrounding impact investing. We define impact investing in the following way:

Impact investing - Investing in, or divesting from, for-profits with the intention of generating social benefit.

On this definition, impact investing has two defining features: it is firstly about the intention to have impact rather than the actual attainment of impact; and secondly impact investments need not produce financial returns.

Different socially motivated investors will be motivated by different moral values and conceptions of the good society. Some might place great weight on improving the environment, while others will think reducing poverty is more important. In contrast to impact investors, *socially neutral* investors make decisions solely based on their expected financial returns.

The aim of impact investing is to make the world better, but what does this mean? The goal should be to make the world better than it would have been otherwise. In other words, the aim should be to have counterfactual social impact. The importance of the counterfactual is illustrated by the following example. Suppose I see a woman having a heart attack and perform CPR. I save her life, but because I have never performed CPR before, I injure her in the process. It is clear I have done a great thing - I saved the woman's life and she would have died had I not stepped in. Now suppose I had pushed a paramedic out of the way and performed CPR. In this case, my actions did save the woman's life, but I made things worse than they would otherwise have been had I not acted. This example shows that the measure of success is the difference you make relative to if you hadn't done anything, and that considering the counterfactual is crucial to evaluating social impact.

So, investors should not just to ask, "what happened?", but should also ask "what would have happened if I had not invested?"

Counterfactual social impact - The difference between what happens as a result of your investment and what would have happened otherwise.

As we will see, a key concern with impact investing is, to use the analogy above, that investors might be elbowing other investors out of the way, and that companies might merely be elbowing other companies out of the way. In this field, replaceability is a recurring worry.

There are three factors that bear on the impact of impact investing:

1. The impact of the enterprise itself
2. The contribution of the investment to the success of the enterprise
3. The contribution of the investor's non-monetary support to the success of the enterprise

We will discuss each of these factors in turn. For an impact investment to have counterfactual impact, the company invested in must at least have some positive social effects. In other words, it must have *enterprise impact*.

Enterprise impact - The counterfactual impact a business has through its products and operations.

A company can have impact through its products if it improves the lives of consumers or produces positive externalities. For example, a solar power company could improve the welfare of its consumers by reducing electricity costs, and could also produce positive externalities for everyone else by reducing CO₂ emissions. A company can also have impact by benefitting its workers or other actors in its supply chain. For simplicity, in this report, we chiefly focus on the impact companies can have through their products.

Impact investors want not merely to invest in a socially beneficial company, but also to have *additionality*: they aim to make a difference to the performance of the company. They can do this in two ways. Firstly, impact investors' primary aim is usually to increase the capital available to socially beneficial businesses. That is, they aim to have *investment impact*.

Investment impact - The counterfactual impact an investment has on the performance of a company or on the wider market.

Having investment impact is crucially about improving a company's performance compared to the counterfactual. An investment has impact if it provides capital at lower cost than the business would have incurred otherwise. Cheaper capital can enable companies to experiment, scale up, and pursue their social objectives.

Beyond providing capital, impact investors can also have additionality by providing a range of non-monetary benefits.

Non-monetary impact - The counterfactual impact an investor has on the performance of a company or on the wider marketplace through means other than providing capital.

Impact investors can have non-monetary impact in four main ways, which we discuss in more detail in section 3.3:

1. Finding and promoting impact investment opportunities
2. Providing technical assistance and access to networks
3. Securing and protecting the social mission of a company they have invested in
4. Gaining publicity for an advocacy campaign

As well as having social impact, many impact investors also wish to make a financial return from their investments. The financial returns investors aim for range from the concessionary to the non-concessionary.

Non-concessionary investments - Investments that do not sacrifice risk-adjusted financial returns.

Concessionary investments - Investments that sacrifice some risk-adjusted financial returns.

Concessionary investments are on a spectrum from slight financial sacrifice at one end, to a grant to a company at the other. (A grant can be thought of an investment in which the investor loses all their money.) The lower financial returns an investor accepts for the sake of social impact, the higher the opportunity cost of impact investing. The opportunity cost of an investment is what the investor could otherwise have done with the money, which could be donating straight away, or *investing to give*: socially neutral investing and donating the profits later...

...Promising areas for impact investing

To recap, the principles are as follows:

- **Principle 1 - Support companies that benefit (poor) consumers or produce positive externalities**

An example of a company that produces positive externalities would be Impossible Foods, which reduces meat consumption. Companies such as Transferwise benefit the extreme poor by reducing the cost of remittances.

- **Principle 2 - Choose a high-impact cause area**

There is limited work on cause prioritisation, but according to many who have thought deeply about it, three problems that appear highly promising are global poverty and health, factory farming, and global catastrophic risk. Other potentially promising areas include improvements in the institutions of science, mental health, tobacco control, criminal justice reform, and so on.

- **Principle 3 - Support companies in uncrowded markets**

Companies in crowded markets, like the 2000s cleantech market, will have limited counterfactual social impact because they are readily replaceable by rivals.

- **Principle 4 - Work in inefficient markets and expect financial sacrifice**

Impact investors should focus on VC or angel investing in private markets and be willing to accept financial sacrifice.

- **Principle 5 - Work on problems that are neglected by other impact investors**

Impact investment is an increasingly popular space, and impact investors should aim to find areas that are not crowded with other impact investors.

- **Principle 6 - Work in areas where you have, or can gain, an information or network advantage over other investors**

The implications of this principle will depend on the information and network advantage available to each impact investor....

Consider the counterfactual

...[T]he following questions are crucial to bear in mind when assessing investment impact:

1. Would a socially neutral investor have made the investment anyway or nullified the effect of your investment?
 1. Have you invested in an efficient market?
 2. Is your portfolio getting (above) market-rate returns in the long run? Is the space crowded with socially neutral investors?
2. Would another socially motivated investor have made the investment anyway?
 1. Is the space crowded with other socially motivated investors?
 2. Do you have a comparative information advantage?

Investors can avoid investing in markets crowded with socially neutral investors by accepting below market returns. But some investments offering below-market returns might also be crowded because they are attractive to a large number of socially motivated investors. So, there is a case for investing in areas that are unfashionable among socially motivated investors.

We will never be certain about whether someone else would have taken your investment opportunity. At best, we can only make probabilistic judgements. It will always be difficult to quantify this probability. Nevertheless, without trying to quantify it, an evaluation would be incomplete...

4. An overall assessment of impact investing

We conclude by discussing whether impact investing is a promising approach to doing good. In making this judgement, it is crucial to consider the opportunity cost of impact investing. In the same way, if our aim were to make money, we would always consider the opportunity cost of an investment: we wouldn't compare the return on our investment to what we would have got if we had done nothing. Instead, we would compare our ROI to what we could have done otherwise with the money: if I chose an investment with a 3% return, but another available investment had an 8% return, then I would have made a mistake. The same applies if our aim is to have social impact. So, people aiming to do good need to consider what else they could have done with their money. We have seen that impact investing probably involves some financial sacrifice, so it could have fairly substantial opportunity costs.

The two alternatives to impact investing are:

Donating now - Making grants to high-impact charities

Investing to give – Socially neutral investing for profit and donating the proceeds later to high-impact charities.

For an introduction to how to choose between these two options, see [this article](#). The benchmark set by philanthropy, if done carefully, is high. For example, our climate change report suggests that our two recommended climate charities have in the past averted a tonne of CO₂ for a something on the order of \$0.10-\$10. Research by our research partner GiveWell suggests that the Against Malaria Foundation saves a life for something roughly around \$4,000 (as of September 2018).

We will begin by recapping the challenges involved in successful impact investing, will proceed to discuss the respective advantages of non-profit and for-profit approaches, and, to get a sense of the space, will conclude by assessing the approach to impact evaluation used by one of the leading impact investing platforms.

4.1. The challenge of impact investing

Following the six principles we laid out in section 2 is likely to be difficult. First, impact investors have to find a company that produces genuine social impact in a high impact cause area. As we saw in section 3.1, it is usually difficult to identify in advance which projects will be socially beneficial. Finding a socially beneficial company is no guarantee of social impact; impact investors need to have investment impact or non-monetary impact. There is a trade-off between social impact and financial performance because investments in opportunities offering market-rate returns are likely to have low counterfactual impact, and because impact investors face costs that are not borne by regular investors. This increases the opportunity costs of impact investing: profits you could have made by doing regular investing cannot be donated to high impact charities.

These effects could be substantial. For example, Norway's \$860 billion sovereign wealth fund divestment from tobacco cost it \$1.96 billion from 2006 to 2015. Two things are notable here. First, as per the argument in Principle 4, this seems to have had little impact on the cash flow of the tobacco industry over this period. Second, this money could have been spent on other things. For example, the Norwegian government could have increased global spending on clean energy R&D by 10% in 2017. Alternatively, if they had spent the money on highly effective global health interventions, then, according to estimates from *The Lancet*, they could expect to have saved around 170,000 lives. They could also have followed the approach of 'mission hedging' and used the money to advocate for regulation of tobacco products. The money that could have been donated would have matched that currently donated by Gates and Bloomberg, the leading philanthropists working in this space. This suggests

that people aiming to have impact need to seriously consider the option of investing to give.

There may be more scope to have counterfactual impact through VC or angel investing, but exploiting inefficiency remains difficult and impact investors have to compete with socially neutral investors and an apparently growing number of other impact investors. The best prospect of counterfactual investment impact likely comes from the chance to subsidise the capital of companies on the brink of viability.

A key barrier to successful impact investing is that socially beneficial businesses have to solve two very difficult optimisation problems simultaneously: running a successful business and having substantial social impact. The evidence suggests that most attempts at doing either of these things alone fail. Doing both at the same time is therefore likely to be especially difficult. In many cases, it probably makes more sense to pursue one's business aims and one's charitable aims separately.

To be a little bit more precise and technical: multi-objective optimisation is harder than single-objective optimisation, and it is usually probably better to optimise for financial returns without social impact constraints with investments that feed your charitable giving, and then to optimise for social impact through non-profits without profit-making constraints.

4.2. For-profits vs. charities

Impact investors aim to have impact through for-profits, whereas philanthropists aim to have impact through charities. Why favour one type of approach over the other? Each approach has its advantages.

The advantages of non-profit solutions

For-profit solutions are likely to fail in certain political and economic conditions, namely for the provision of public goods and beneficial goods with insufficient consumer demand. A final general advantage that non-profit approaches have is that we should expect them to be more neglected.

Public goods

In economics, public goods are defined as those that are both non-excludable and non-rivalrous. 'Non-excludability' means that the cost of keeping non-payers from enjoying the benefits of the good or service is prohibitive. If an entrepreneur stages a fireworks show, for example, people can watch the show from their windows or backyards. Because the entrepreneur cannot charge a fee for consumption, each consumer has an incentive to free ride by allowing others to pay for the show and then watching from their backyard. If the free rider problem cannot be solved, valuable goods and services will remain unproduced. 'Non-rivalrous' goods are those for which consumption by one person does not affect other people's ability to consume the

good. For example, my learning some information does not reduce your ability to learn that piece of information.

Public goods will tend to be underprovided by the market because for-profit firms cannot reap the benefits of providing them.

From the point of view of the impact-focused individual, two of the most important public goods are policy change and research. For any problem that you think is due to political failure, non-profits are much better suited to solving it than for-profits. For example, removing zoning restrictions in major metropolitan areas like San Francisco would be very socially beneficial. Why then is there not a for-profit anti-zoning company that advocates for land use reform? The reason is that even if such a company were to succeed in changing the law, they could not effectively exclude non-payers (i.e. almost everyone in San Francisco) from enjoying the resultant benefits. For many of the major problems facing the world, such as climate change, biosecurity and global poverty, political change is arguably the most effective way forward. If so, this counts in favour of non-profit approaches.

Research is another important public good. While it is true that many companies engage in research, it will still tend to be underprovided by the market because the information produced is a public good. For example, everyone benefits from research into reducing the risk of nuclear war, but it is impossible to exclude non-payers from enjoying the benefits of this research. This is why much research is funded by governments and large foundations.

Goods that are strongly undervalued by consumers

Sometimes consumer demand for goods does not match up to the social benefits provided by the good. In the extreme case, demand for a highly beneficial good falls to zero when even a small price is charged. For example, the evidence suggests that charging even a small fee for malaria bednets would greatly reduce demand, making it much more effective to distribute the nets for free. This is in part a product of the fact that consumers of malaria nets are very poor and have low willingness to pay, but the complete drop off in demand in response to even small fees may be because consumers underestimate the benefits of bednets. This also seems to be true for other products, including solar lamps and school uniforms. For goods such as these, there will be insufficient demand to sustain for-profit businesses.

Neglectedness of non-profit solutions

The final advantage of non-profit approaches over for-profit approaches is that we should expect non-profit approaches to be more neglected because there are much stronger incentives (i.e. money) to find for-profit solutions. In the US in 2016, around \$390 billion was donated to charity, which was only around 2% of US GDP. Even within this small slice of funding, there is severe misallocation of resources. For example, the number of animals killed in factory farms dwarfs those killed in

laboratories or other sources, but factory farming receives only a small fraction of the funding devoted to animal welfare. We believe that this kind of misallocation is widespread in philanthropy, which opens up the opportunity for careful philanthropists to find highly neglected and important areas. Because of the incentives in well-functioning markets, there is much less scope to find neglected unfashionable areas that are suited to for-profit solutions.

The advantages of for-profit solutions

When neither of the above two conditions apply, for-profits have a couple of advantages over non-profits: in the right market conditions, they have better product feedback and face better incentives. Non-profits are subject to a principal-agent problem. In the for-profit case, the person served (the principal) pays the company (the agent) for the service, whereas in the non-profit case, the person served (the agent) does not pay for the service, and it is instead paid for by philanthropic donors (the principal). For-profits therefore get valuable direct feedback from consumers about whether they are providing a good product: if consumers don't like their product, then their revenue will decline. Non-profits lack this feedback because donors have worse information on the quality of the product than do direct beneficiaries. A non-profit could provide a poor service but still receive ample funding from donors.

It follows that competitive pressure between for-profits encourages them to compete on quality and price in order to win market share. Although there is competitive pressure between non-profits, the organisational incentives are towards competing for donations rather than competing to serve beneficiaries.

It is the principal-agent problem that explains the necessity for rigorous external impact evaluation of charities, as is carried out by Founders Pledge, our research partner GiveWell, and other impact-focused philanthropists such as Hewlett Foundation. This kind of external evaluation is not required in well-functioning markets because the incentive, information and feedback structures are set up differently. This is not to say that non-profits will never provide a good service if they are not monitored. The point is that the incentives and feedback mechanisms are not set up to encourage efficiently providing a good product.

Replaceability in philanthropy

It should also be made clear that philanthropy also faces the problem of replaceable funding. If you fill the funding gap of a charity, the effect of that might be to free up money for another donor who would have filled the funding gap had you not done so. In this case, the marginal effect of your donation would actually be to shift money to what that next donor donates to.

However, there are generally weaker incentives to fill the funding gap of charities since there is no profit reward for doing so. In addition, this dynamic is only at play when your donation fills the funding gap of a charity. At Founders Pledge, a charity's funding

gap is a key consideration relevant to the decision of whether we recommend the charity or not, and we only recommend charities that could productively use additional funds. Therefore, careful philanthropists should be able to significantly reduce the replaceability concern.

4.3. Should you impact invest?

We argued that if you cannot satisfy the six principles of impact investing, you are unlikely to have substantial impact. Thus, when these principles cannot be satisfied, donating now or investing to give are likely to be the better option. However, when the six principles can be satisfied, impact investing might be a good option. In particular, investors with a strong informational advantage in a high impact sector might be able to find promising and neglected opportunities.

This being said, we should not let the best be the enemy of the good. Impact investing might, when the six principle cannot be satisfied, be worse than investing to give, but from a social impact point of view it is still probably better than socially neutral investing alone. Thus, the arguments here should not be thought to justify socially neutral investing over impact investing. Impact investing in efficient markets may often only have a modest effect on the capital available to companies, but a modest effect is better than nothing, and the indirect effects appear more substantial. Our point is that, when the six principles cannot be satisfied, people can probably increase their impact significantly by switching from impact investing to investing to give or donating now.

To establish with greater certainty the merits of investing to give vs. impact investing, it would be useful to compare case studies of what can be achieved by the two approaches in different contexts. To take a first step in that direction, we carried out a brief review of Acumen Fund, a non-profit impact investment fund, which seems to be a field leader in impact evaluation. In their 2017 Energy Impact Report, Acumen state that they invested \$22.1 million in companies that averted 6.4 million tonnes of CO₂, and they are on track to get all of their investment back. Our review showed that their impact evaluation excludes the following key factors:

Acumen do not include staff costs when calculating their cost-benefit ratios. Given that they have around 110 staff working on eight sectors, we can assume that they have roughly 27 staff working on energy, and roughly assume that they are paid \$50,000 per year. Over ten years, this makes staff costs of roughly \$13.5 million, more than half of the total capital invested in the sector (\$22.1 million).

Acumen does not attempt to account for the counterfactual investment impact of its capital – it only discusses this consideration in moderate depth for one company in its portfolio. It claims that its overall \$22.1 million investment has crowded in \$104.5 million from other investors, but it is unclear whether this money would have been forthcoming anyway.

Acumen does not measure the marginal effect of its investments on the performance of the company, but instead compares its own costs to the total benefits produced by all the companies it has invested in.

Acumen has made great strides in impact measurement recently and appears to be a field leader in this respect, but this suggests that there is room for improvement in the field before it begins to accurately estimate counterfactual impact.

Finally, it is worth noting that, if we take the estimates given in Acumen’s Energy Impact Report at face value, the cost-effectiveness of Acumen’s portfolio is worse than both donating now and investing to give. In our climate change report, we roughly calculated that our recommended charities – the Clean Air Task Force and Coalition for Rainforest Nations – would avert a tonne of CO₂ for something on the order of \$0.10 - \$10, though such estimates are of course highly uncertain. If we take \$1 per tonne as a reasonable median estimate, the possible social impact of Acumen’s Energy Fund vs. donating now vs. investing to give is shown in the table below.

	2007 investment	Tonnes of CO ₂ averted (2007-17)	Financial return in 2017	Tonnes of CO ₂ averted
Acumen Energy Fund	\$22.1m	6.4 million tonnes	\$22.1m	6.4m
Donate now	\$22.1m	22m	\$0	0
Invest to give	\$22.1m	0	\$35m	35m

This table assumes that:

- Acumen will reinvest their financial returns, achieving the same social cost-effectiveness over the following ten years (2017-2027).
- ‘Investing to give’ investors can get annual financial returns of 5% over 10 years, which can then be donated to our recommended climate charities in ten years, at a social cost-effectiveness of roughly \$1 per tonne.
- The ‘donate now’ philanthropist donating to our climate charities in 2007 also enjoyed a cost-effectiveness of roughly \$1 per tonne.

This table is of course a great simplification and disfavours Acumen insofar as it does not allow for the possibility that they could continue to reinvest the profits in socially impactful businesses in perpetuity. Still, we would expect their future cost-effectiveness to decline as the low-hanging fruit are taken. Moreover, for the reasons we outlined above, the Acumen figures are likely to be a considerable overestimate and we would guess that their cost-effectiveness would decline by at least one order of magnitude once the factors we discussed are properly accounted for. So, the table suggests that impact investing looks worse than the alternatives in this case. Even on the pessimistic estimate of the cost-effectiveness of our recommended climate charities (\$10 to avert a tonne of CO₂), donating now or investing to give look better than impact investing.

(Note that the table is not meant to show that investing to give is actually better than donating now: there are unaccounted for advantages to donating now, such as diminishing returns and compounding social benefits)...