


FOUNDERS PLEDGE REPORT

IMPACT INVESTING



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.¹ 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

¹ A note on the authors of this report: John Halstead is part of the research team at Founders Pledge. Hauke Hillebrandt contributed to this report while a researcher at the Center for Global Development and [Lets-fund.org](https://lets-fund.org).



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 high-impact charities are available on our [research page](#).

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 only likely to be the best approach in specific circumstances. Impact investing might be a good option for people doing VC or angel investing, who accept financial sacrifice, and who have an informational advantage over other investors that allows them to reliably identify promising opportunities in a high-impact cause area. A good example 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, we think that such opportunities are likely to be difficult to find, and when they cannot be found, donating now or investing to give are likely to be the best option.

It is true that 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.




Acknowledgements

For comments on draft versions of this report, we are grateful to Tom Adams, Richard Batty, Stephan Chambers, Darrell Duffie, Mathijs van Dijk, Robin Greenwood, Martijn Kaag, Mark Kritzman, Kellie Liket, Joseph Mela, Antti Petajisto, Stefan Schubert, Lennart Stern, Ben Todd, Philip Trammell, and Ivo Welch.



Report structure

1. Introduction	7
Introduces the key concepts of impact investing	
2. Principles of Impact Investing	11
Outlines five principles to guide successful impact investing	
3. How to evaluate impact investing	28
Discusses how to evaluate the three key constituents of counterfactual impact: enterprise impact, investment impact and non-monetary impact	
4. An overall assessment of impact investing	40
Discusses whether impact investing is the best approach for people aiming to have maximal social impact	
5. Recommended reading	48
Lists some high-quality articles relevant to impact investing	



1. Introduction

Impact investing is an increasingly popular approach to doing good. In 2016, \$8.72 trillion, or about 21% of the \$40.3 trillion invested in professionally managed assets in the U.S., was labelled as “socially responsible investing” (though we think much of this is not very strict).² This is a 14-fold increase since 1995.³ One apparent attraction of impact investing is that it allows investors to have social impact while earning a return on their investment, which can be reinvested for further social good. A donation to charity, in contrast, yields no return for the investor. Impact investing also favours entrepreneurial solutions to problems, rather than charitable solutions, which can sometimes be inefficient.

In this report, we outline six Principles of Impact Investing, which should guide impact investors aiming to have substantial social impact. We then set out a framework for evaluating impact investing, and conclude by comparing impact investing to alternative ways of doing good: *investing to give* – investing for profit and then donating the proceeds to charities later – and *donating now*.

Defining key terms

We first need to clarify the key concepts surrounding impact investing.⁴ 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

² US SIF Forum for Sustainable and Responsible Investment, “Report on US Sustainable, Responsible and Impact Investing Trends: 2016 Trends Report Highlights,” 2016, <https://www.ussif.org/files/Trends/US%20SIF%202016%20Trends%20Overview.pdf>.

³ US SIF Forum for Sustainable and Responsible Investment.

⁴ This conceptual framework is mostly borrowed from Paul Brest and Kelly Born, “Unpacking the Impact in Impact Investing,” Stanford Social Innovation Review, August 2013, https://ssir.org/articles/entry/unpacking_the_impact_in_impact_investing.



illustrated by the following example. Suppose I see a woman having a heart attack and perform CPR.⁵ 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.

⁵ This example is from William MacAskill, *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference* (London: Guardian Books, 2015), 83–84.



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 3.3:⁶

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

⁶ Some of these are mentioned in Brest and Born, "Unpacking the Impact in Impact Investing."



donating straight away, or *investing to give*: socially neutral investing and donating the profits later.



2. Principles of Impact Investing

With the conceptual background laid out, we will now enumerate six Principles of Impact Investing:

1. Support companies that benefit (poor) consumers or produce positive externalities
2. Choose a high-impact cause area
3. Support companies in uncrowded markets
4. Work in inefficient markets and expect financial sacrifice
5. Work on problems that are neglected by other impact investors
6. Work in areas where you have, or can gain, an information or network advantage over other investors

These principles are each (close to) necessary conditions for successful impact investing. If impact investing does not live up to *any one* of these principles, then it is unlikely to have substantial impact. We will discuss each principle in turn.

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

For a company to have enterprise impact, it must produce a product that benefits consumers or produces positive externalities. As we argue in section 3, identifying which companies are likely to produce real social benefits is much more difficult than is commonly assumed in the impact investing space.

Benefit (poor) consumers

Companies like Amazon and Uber provide large benefits to their consumers in the form of better service and reduced cost. Many well-off consumers in rich countries are willing to pay for this, which is why these companies have such a high market valuation. However, market demand for a product is not always a reliable guide to its social benefits.

The first reason for this is that consumers may not accurately estimate how good a product would be for their lives. Addictive products are a good example: tobacco product sales totalled \$760 billion in 2016 (excluding China).⁷ Even though market demand for these products is high, they are actively harmful to consumers.

Second, there are immense differences in purchasing power between different consumer groups. The difference is most stark at the international level: as of 2015, a person below the US poverty line earning \$11,000 per year was still richer than 85% of people in the world,


⁷ British American Tobacco, "The Global Market," accessed September 27, 2018, http://www.bat.com/group/sites/UK__9D9KCY.nsf/vwPagesWebLive/DO9DCKFM.



even adjusting for purchasing power of a dollar in different countries.⁸ GDP per capita in the US is around \$57,000, whereas in Uganda it is \$615.

This counts in favour of producing products that benefit poor people in low- and middle-income countries, rather than the relatively poor in rich countries. The reason for this is that money has diminishing marginal utility. Figure 1 shows life satisfaction plotted against household income in different countries:


Figure 1. Household income against life satisfaction



Source: The Economist, '[Money can buy happiness](#)', May 2013. (Data from Wolfers and Stevenson, "[Subjective well-being and income: is there any evidence of satiation?](#)", NBER Working paper (2013))

The x-axis on this chart uses a log scale for household income: each doubling of household income produces around 0.5 point increase in life satisfaction. So, increasing the household income of an Indian family by \$2,000 is about as good as increasing the income of a US household by \$32,000. So, all else equal, creating consumer value in India has more impact than creating consumer value in the US. This suggests that larger social gains can be had from targeting people in low- and middle-income countries than people in rich countries.

⁸ MacAskill, *Doing Good Better*, 21–22.



Produce positive externalities

Many products generate social good that is not captured in the transaction between the company and the consumer. In the parlance of economics, these products have *positive externalities*. For example, meat-alternative companies such as Impossible Foods or Quorn produce benefits for animals because they reduce meat consumption, thereby sparing animals from factory farming. Tesla started by making luxury cars in order to fund research into better batteries, aiding the transition to low-carbon transport, in turn producing benefits for future generations. Neither of these benefits accrue (primarily in the case of Tesla, or at all, in the case of Quorn) to the consumers of these products.

Principle 2: Choose a high-impact cause area

Having acknowledged that companies can be socially impactful in the two ways mentioned above, it is also crucial to consider what kind of multiplier you can get on impact by choosing the cause area to work on. Impact investors face a vast array of possible projects that could benefit poor consumers or produce positive externalities, including climate change, food waste, plastic in the oceans, global poverty, reducing animal suffering, improving agriculture, and so on. How should you decide which of these to work on? Thankfully, tools now exist that make answering this question easier.

A tool we find useful at Founders Pledge to select causes is the Importance, Tractability and Neglectedness (ITN) framework.

- **Importance:** How many individuals are affected by the problem? How badly does the problem affect them?
- **Tractability:** Relative to how large the problem is, how easy is it to solve it? Can we realistically make meaningful progress at this time?
- **Neglectedness:** How much attention does the problem get now, and what resources are currently spent on solving it? In other words, is it a crowded area?

Sometimes, the attention problems receive is not commensurate with their scale. For example, around \$2 billion is spent on surgical procedures to prevent male baldness, while \$547 million is spent trying to prevent malaria, which kills hundreds of thousands of people every year.⁹ Figuring out how much attention a problem receives is crucial because crowded problems are likely to face *diminishing returns* in that the best opportunities to do good (the ‘low-hanging fruit’) have already been taken.

The ITN framework can in principle be used to rank all causes. For example, see [this visualisation](#) of importance and neglectedness within the area of preventing animal suffering.

⁹ Dick Ahlstrom, “Baldness Trumps Malaria? Funding Shows Something out of Kilter,” The Irish Times, accessed July 25, 2018, <https://www.irishtimes.com/business/innovation/baldness-trumps-malaria-funding-shows-something-out-of-kilter-1.3109417>.



The impact-research organisation 80,000 Hours has developed [a version of the ITN framework](#), which it uses to rank problems according to how promising they are to work on.¹⁰ We are unsure about this and other versions of the ITN framework, but it is a useful heuristic starting point for cause prioritisation.¹¹

In the charitable space, the evidence suggests that merely by choosing the right cause area, you can increase your impact by at the very least one order of magnitude. For example, as of September 2018, our research partner GiveWell estimated that the Against Malaria Foundation saves a life for around \$4,000.¹² In contrast, the UK NHS considers it cost-effective to spend up to around \$26,000 to save a *year* of life.¹³ Similarly large impact multipliers may be available to impact investors who choose their problem area carefully.

According to those that have thought deeply about global prioritisation, three problems that appear highly promising are global poverty and health, factory farming, and global catastrophic risk.¹⁴ Other potentially promising areas may include: improvements in the institutions of science, improvements in mental health, land-use reform, tobacco control, criminal justice reform, and so on.

Principle 3: Support companies in uncrowded markets

One way to assess the counterfactual impact of a company would be to calculate the number of units sold and then multiply by the social benefits of each unit sold. However, this approach is not accurate. The social benefit of a company is given by the difference between the world in which the company exists and the world that would have been brought about if the company did not exist. If a company did not exist, then another similar company might have stepped in and provided a very similar product. So, the effect of a successful company is usually a speed up: it brings the benefits of a product forward in time.

For this reason, the counterfactual social impact of a company depends on the crowdedness of the market it works in. In general, the more crowded a market is relative to its size, the more likely it is that, if your company did not exist, a rival would step in and provide a similar product. For example, in the 2000s, the cleantech market was overcrowded, making it difficult for any one company to succeed despite the vast size of the energy market.¹⁵ This also greatly reduced the potential counterfactual impact of any single cleantech company – if one went

¹⁰ For an explanation of the framework see <https://80000hours.org/articles/problem-framework/>

¹¹ We outline an alternative version in section 2 of our climate change report at www.founderspledge.com/research

¹² See their cost-effectiveness analysis at <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

¹³ <https://www.nice.org.uk/news/blog/carrying-nice-over-the-threshold>

¹⁴ For an introduction to different cause areas see, Effective Altruism, “Introduction to Effective Altruism,” accessed August 24, 2018, <https://www.effectivealtruism.org/articles/introduction-to-effective-altruism/>.

¹⁵ Peter Thiel, “Peter Thiel (Ft. Blake Masters & Matt Cauble) – Lecture 5: Business Strategy and Monopoly Theory,” Genius, accessed August 29, 2018, <https://genius.com/Peter-thiel-lecture-5-business-strategy-and-monopoly-theory-annotated>.



bust, another would quickly fill the niche. In an uncrowded market, in contrast, there is less scope for a rival business to eat into your counterfactual impact.

The counterfactual consideration constitutes a crucial difference between assessing the social impact of a business and assessing its financial value. When assessing the financial value of a business like Airbnb, we can completely ignore the alternative home-sharing business that would later have arisen had Airbnb not existed. If we want profit, then it does not matter whether someone else would have enjoyed our profits if we had not invested.

Principle 4: Work in inefficient markets and expect financial sacrifice

Principles 1, 2, and 3 concerned how companies can have enterprise impact, but enterprise impact is only half the story. *It is possible to invest in a socially impactful company and have zero social impact.* In order to have additionality – to make a difference to the performance of the company – investors must either provide cheaper capital or offer valuable non-monetary support. Principles 4, 5 and 6 concern how investors can have additionality.

The degree to which impact investors can have additionality depends crucially on the markets in which they operate. In this section, we discuss why having investment impact is likely difficult in highly efficient markets, such as large and liquid public stock markets. This is true for both concessionary and non-concessionary investing. Impact investing is likely to have most of its impact indirectly via stigmatisation of targeted firms and industries. There may be more scope for investment impact with VC or angel investing in inefficient private markets, but all investors should expect to sacrifice financial returns for the sake of impact.

Impact investing in efficient equity markets

According to the efficient capital market hypothesis, asset prices fully reflect all available information, making it difficult for individual investors to achieve better risk-adjusted financial returns than the market.¹⁶ If markets were efficient, then demand curve for stocks would be nearly horizontal (in technical terms, nearly perfectly elastic) because arbitrageurs would return stocks to fundamental values. If so, changes in demand have no effect on the stock price.¹⁷

There is a wealth of evidence demonstrating that capital markets have some inefficiencies over the course of around 3 months.¹⁸ In other words, there is evidence that demand curves are

¹⁶ Burton G. Malkiel and Eugene F. Fama, “Efficient Capital Markets: A Review of Theory and Empirical Work,” *The Journal of Finance* 25, no. 2 (n.d.): 383–417, <https://doi.org/10.1111/j.1540-6261.1970.tb00518.x>.

¹⁷ Jeffrey Wurgler and Ekaterina Zhuravskaya, “Does Arbitrage Flatten Demand Curves for Stocks?,” *The Journal of Business* 75, no. 4 (2002): 583–608, <https://doi.org/10.1086/341636>.

¹⁸ This is often tested by measuring the effect of inclusion in, or exclusion from, a stock index. For recent overviews, see Nimesh Patel and Ivo Welch, “Extended Stock Returns in Response to S&P 500 Index Changes,” *The Review of Asset Pricing Studies* 7, no. 2 (December 1, 2017): sec. 1, <https://doi.org/10.1093/rapstu/rax012>; Antti Petajisto, “Why Do Demand Curves for Stocks Slope Down?,” *Journal of Financial and Quantitative Analysis* 44, no. 5 (October 2009): 1013–44, <https://doi.org/10.1017/S0022109009990317>.



downward sloping in the short-term. There is much less agreement on how efficient markets are on the >3-month timescale.¹⁹ Petajisto (2011) writes of stock market demand curves:

“Today the evidence for downward-sloping demand curves in general is overwhelming. The main unsettled question is about how the slope depends on the horizon: if demand curves are much steeper in the short run than in the long run, then the price impact will reverse quickly...”²⁰

The reason for this disagreement is that it is difficult to statistically assess the longer-term demand curve of stocks. Some recent work suggests that over the course of around 6 months, demand curves are nearly perfectly elastic, such that prices fully revert back over this timescale.²¹ Other work also suggests that probably more than half of the effect reverses over the ~6 month timeframe.²² However, there is also some evidence that price anomalies can persist for much longer than 6 months, suggestive of long-term inefficiency.²³

Difficult to affect stock prices

This has interesting implications for impact investors working in large and liquid public stock markets. Assuming that stock markets are highly efficient on the ~6-month timeframe, if impact investors increase the price of a stock in the short-term, socially neutral investors will likely sell their shares at a premium until the shares partially or completely revert to the equilibrium price in the medium term. Conversely, if enough investors divest from socially harmful companies such that their share prices fall, the odds are that socially neutral investors will buy up the underpriced stock, thereby returning the stock to its equilibrium price in the medium term.²⁴ Unless socially neutral investors revise downward the net present value of the

¹⁹ Note that in the literature, it is often said that price effects are “permanent” if they last on the order of weeks. See for example, Joop Huij and Georgi Kyosev, “Price Response to Factor Index Additions and Deletions,” SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, October 3, 2016), <https://papers.ssrn.com/abstract=2846982>.

²⁰ Antti Petajisto, “The Index Premium and Its Hidden Cost for Index Funds,” *Journal of Empirical Finance* 18, no. 2 (March 1, 2011): 272, <https://doi.org/10.1016/j.jempfin.2010.10.002>.

²¹ Patel and Welch, “Extended Stock Returns in Response to S&P 500 Index Changes.” See also Atif Ansar, Ben Caldecott, and James Tilbury, “Stranded Assets and the Fossil Fuel Divestment Campaign: What Does Divestment Mean for the Valuation of Fossil Fuel Assets?” (Smith School of Enterprise and the Environment, University of Oxford, 2013), 22–23, <https://www.smithschool.ox.ac.uk/publications/reports/SAP-divestment-report-final.pdf>.

²² Robin Greenwood, “Short- and Long-Term Demand Curves for Stocks: Theory and Evidence on the Dynamics of Arbitrage,” *Journal of Financial Economics* 75, no. 3 (March 1, 2005): 607–49, <https://doi.org/10.1016/j.jfineco.2004.03.007>; Petajisto, “The Index Premium and Its Hidden Cost for Index Funds.”

²³ Some studies show that the returns of the same company listed in different markets diverges in the long-term. See for example Abe De Jong, Leonard Rosenthal, and Mathijs A. Van Dijk, “The Risk and Return of Arbitrage in Dual-Listed Companies,” *Review of Finance* 13, no. 3 (July 1, 2009): 495–520, <https://doi.org/10.1093/rof/rfn031>. There is also some debatable evidence that divestment has had an effect on the price of sin stocks; see Harrison Hong and Marcin Kacperczyk, “The Price of Sin: The Effects of Social Norms on Markets,” *Journal of Financial Economics* 93, no. 1 (July 1, 2009): 15–36, <https://doi.org/10.1016/j.jfineco.2008.09.001>.

²⁴ Ansar, Caldecott, and Tilbury, “Stranded Assets,” 30.



company's cash flows, the effect on the stock price will be either highly attenuated or completely nullified in the long-term.

Impact investing could have a more sustained effect on stock prices if there were a sufficient move in the market away from socially neutral investing and towards impact investing. It might be argued that if (say) harmful industries comprise 20% of global equities, then the tipping point would be when 80% of stocks are divested because at lower levels of divestment, socially neutral investors could always buy up the equity left by socially responsible investors.²⁵ However, this line of argument seems mistaken. First, the efficient capital market hypothesis states that current market prices reflect all available information; it does not say that each individual investor has the same view of the correct valuation. (In this way, the efficient market hypothesis works by a 'wisdom of crowds' effect.) This means that not all socially neutral investors would believe that divested stocks are incorrectly priced. So, it is more plausible to suppose that 40%–60% would think that they are incorrectly priced, meaning that the tipping point must be that 20%–40% of the market is socially responsible (assuming that 20% of industries are harmful).

Moreover, many of those who think the stock is incorrectly priced face restrictions on their portfolios and so may not be able to move in. Investors also would not borrow to invest in a divested stock because doing so would be risky and it is not plausible that investors would significantly change their risk preference.²⁶ Thus, as a rule of thumb, it seems reasonable that if harmful industries constitute 20% of global equities, around 20% of stocks would have to be divested to substantially affect the cost of capital for harmful industries.²⁷

In many cases, this tipping point seems a long way off. Of the \$40 trillion of assets under active management in the US, around \$9 trillion is now labelled as 'socially responsible investing'.²⁸ However, it is difficult to know what proportion of "socially responsible investing" is truly socially responsible. For example, Vanguard's socially responsible European Stock Fund includes British American Tobacco and Royal Dutch Shell.²⁹ This kind of misleading labelling seems fairly typical in the space.³⁰ The willingness to be strict seems limited. As Ansar *et al* note:

²⁵ The same applies to investing in socially beneficial industries, changing what needs to be changed.

²⁶ On this, see Andrei Shleifer and Robert W. Vishny, "The Limits of Arbitrage," *The Journal of Finance* 52, no. 1 (March 1, 1997): 35–55, <https://doi.org/10.1111/j.1540-6261.1997.tb03807.x>. We are grateful to Martijn Kaag for critical discussion of the points in this section.

²⁷ This echoes Ansar et al, see Ansar, Caldecott, and Tilbury, "Stranded Assets," 29.

²⁸ US SIF Forum for Sustainable and Responsible Investment, "Sustainable and Impact Investing in the United States Overview," 2016, <https://www.ussif.org/files/Infographics/Overview%20Infographic.pdf>.

²⁹ Jeff Kauflin, "Do-Good Investors: Watch Out For What These Funds Hold," *Forbes*, accessed October 22, 2018, <https://www.forbes.com/sites/jeffkauflin/2017/12/06/portfolio-placebos/>.

³⁰ Kauflin.



“Despite the huge interest in the media and a three-decade evolution only about 80 organisations and funds (out of a likely universe of over 1,000) have ever substantially divested from tobacco equity and even fewer from tobacco debt.”³¹

As of 2013, only \$5 billion has been divested from tobacco. The market cap of Philip Morris alone is \$140 billion.³² Other harmful industries are also large enough to make mounting an effective impact investing campaign seem difficult. For example, the world’s largest oil and gas companies had a market capitalisation of \$4 trillion at the end of 2012.³³ Overall, the impact investing tipping point seems a long way away at present. Moreover, given the market cap of the targeted companies, unless an investor controls a very large amount of capital, their own direct effect on stock prices is likely to be small in the first place.

The evidence on the direct effect of socially responsible investment on stock prices is ambiguous. Most previous divestment campaigns, with targets including tobacco, alcohol, weapons, gambling, and companies in apartheid-era South Africa, had modest or zero direct effect on the stock price of targeted companies.³⁴ On the other hand, there is also evidence that “sin stocks” – stocks of companies producing harmful products, such as tobacco and gambling – consistently outperform the market. Hong and Kacperczyk (2009) note that this could either be because socially responsible investment has had a long-term effect, or because of other features of these industries, such as that they face a greater risk of litigation.³⁵

To conclude, regarding the question of whether short-run price effects are negated after around 6 months, the evidence is limited and ambiguous, and the experts disagree. However, on the basis of the evidence we do have, it seems likely that in most cases probably more than half of the price effect is negated on this timescale. Moreover, since few individual investors will be able to have a noticeable effect on stock prices in the first place, our view is that, for the vast majority of socially responsible impact investors, their own actions will have at best modest direct effect on stock prices in large and liquid stock markets.

This being said, two points should be emphasised. Firstly, as we discuss in section 3, socially responsible investing plausibly has more substantial indirect effects. Secondly, as we discuss in section 4, this does **not** justify the refusal to engage in divestment campaigns solely for the sake of financial returns.

³¹ Ansar, Caldecott, and Tilbury, “Stranded Assets,” 12.

³² Yahoo Finance, “Summary for Philip Morris International Inc.,” accessed October 22, 2018, <https://finance.yahoo.com/quote/PM/>.

³³ Ansar, Caldecott, and Tilbury, “Stranded Assets,” 52.

³⁴ For a review, see Ansar, Caldecott, and Tilbury, 64. See also Brigitte Roth Tran, “Divest, Disregard, or Double Down?,” SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, April 1, 2017), sec. 4, <https://papers.ssrn.com/abstract=2952257>.

³⁵ Hong and Kacperczyk, “The Price of Sin.”



The counterfactual replaceability problem

The second argument concerns what we can call the *counterfactual replaceability problem*.³⁶ In an efficient market, if a company's stock offers market-rate returns, then socially neutral investors who have the right information would invest anyway if you did not do so. Thus, in these markets, the sole effect of socially responsible investment is to free up capital for the average investor. The effect of this could be positive *but is very different from the supposed benefits of impact investing*.

This is like the old adage about the economist and student who comes across a \$100 bill lying on the ground. As the student stops to pick it up, the economist says, "Don't bother—if it were really a \$100 bill, it wouldn't be there." In highly efficient markets, there are no \$100 bills lying on the floor: if an opportunity offers market-rate returns, socially neutral investors would take it anyway if you did not do so. Thus, non-concessionary investment is likely to have much smaller counterfactual impact than one would naively assume.

Changes in stock prices do not always strongly affect companies

The final reason that it is difficult to have substantial impact in public stock markets is much less important than the first and the second but is worth mentioning. Movements in the share price do not always affect the capital available to companies. In industries that generate a lot of cash and so do not need to raise capital, changes in share prices will not have much impact. In addition, public equity is generally traded in secondary markets, meaning that changes in stock price affect shareholders, rather than the usable capital the company has at hand.³⁷ However, if the company is planning to issue new stock, a change in the share price will affect their available capital. Moreover, shareholders will, of course, care about movements in the stock price, so if one were to occur, it would change the incentives of shareholders and the behaviour of companies. Entrepreneurs would also be less willing to set up harmful companies if they knew the value of their shares would be low. Moreover, for growing companies or those that periodically face busts, the share price does affect the capital available to the company.³⁸ Thus, this factor is less important than the others.

For all these reasons, in efficient markets, it is difficult for impact investors to affect the capital available to a company. Still, as we discuss in section 3, impact investing might be a way to stigmatise the affected firms, encouraging restrictive regulation and legislation.

Impact investing in inefficient markets

Founders Pledge members who are considering impact investing will probably be interested in VC or angel investing in socially beneficial tech startups. These markets have fewer investors

³⁶ Ansar, Caldecott, and Tilbury, "Stranded Assets."

³⁷ Brest and Born, "Unpacking the Impact in Impact Investing."

³⁸ Chris B. Murphy, "Why Do Companies Care About Their Stock Prices?," Investopedia, May 22, 2018, <https://www.investopedia.com/investing/why-do-companies-care-about-their-stock-prices/>.



and are subject to more imperfections and frictions, such as imperfect information. In inefficient markets, the factors discussed above may not hold, or at least not to the same extent.

Firstly, the dynamic at play in public stock markets, of socially neutral investors returning the stock price to its equilibrium, is not at play in private markets with fewer investors and imperfect information.

Second, the counterfactual replaceability problem is attenuated. In inefficient markets, investors with special expertise and knowledge may be able to find socially beneficial opportunities offering market-rate returns that would be missed by other investors. Nonetheless, **the counterfactual replaceability problem is not eliminated and remains a fundamental consideration in assessing additionality**. Impact investors need to consider their replaceability by both socially neutral investors and socially responsible investors. In most private markets, there will still often be a large number of socially neutral VCs trying to find great returns and exploiting market inefficiency is extremely difficult.

Impact investing without financial sacrifice?

With this background set, we can now explore the merits of non-concessionary vs. concessionary impact investing. In a January 2013 survey of impact investors, 65% of respondents indicated that they were seeking market-rate returns.³⁹ However, the arguments above imply that there is a trade-off between financial returns and social impact. If a company offers market-rate returns, then in efficient markets, your investment is highly unlikely to increase the capital available to the company. In inefficient markets, there may be more impactful opportunities with good financial returns but finding and exploiting this market inefficiency is nevertheless difficult. In general, the cost of capital for a company is equal to the investment returns for investors.⁴⁰ So, investors can only lower the cost of capital if they accept lower investment returns.

The amount of financial sacrifice investors should expect is very important when evaluating the overall merits of impact investing as a way to do good. The profits gained from a portfolio can be donated to high-impact charities. Thus, the greater the financial sacrifice, the greater the opportunity cost of impact investing. We discuss this in more detail in section 4.

There are a number of further arguments which suggest that all impact investors, including VCs working in private inefficient markets, should expect to make a financial sacrifice.

³⁹ Nick O'Donohoe, "When Can Impact Investing Create Real Impact?," Stanford Social Innovation Review, 2013, https://ssir.org/up_for_debate/impact_investing/nick_o'donohoe.

⁴⁰ InvestingAnswers, "Cost of Capital," accessed October 3, 2018, <https://investinganswers.com/financial-dictionary/stock-valuation/cost-capital-112>.



Firstly, there are strong theoretical grounds to believe that screening one's investment portfolio such that it only includes socially beneficial enterprises must decrease expected financial performance on average, as suggested by Modern Portfolio Theory.⁴¹ Only investing in socially beneficial companies shrinks the pool of available investments, thereby leading to less risk diversification and to missing out on high-value opportunities. The socially neutral investor seeking to maximise their personal return will always prefer as many options as possible. Moreover, as we mentioned above, sin stocks consistently outperform the market.⁴² If so, it follows that investors could improve the performance of their portfolio by investing in these socially harmful industries, and therefore that screening the portfolio for social impact must reduce performance. This is true in both public and private markets. In general, there is no reason to expect socially beneficial companies to have better performance than socially neutral or harmful ones.

Surprisingly, the empirical evidence on the financial performance of portfolios screened for impact or 'social responsibility' is ambiguous, with some studies finding that screening damages financial performance,⁴³ some meta-analyses finding no relationship between screening for social beneficial companies and financial performance,⁴⁴ while others even find a positive relationship.⁴⁵ What can explain the divergence between this data and the theory? Studies that investigate the financial performance of screened funds have important methodological shortcomings, which decrease the likelihood that they will identify a real effect:

1. They are correlational in nature and cannot shed light on whether there is a causal effect of screening portfolios on financial performance. For instance, it might be that screened portfolios do reasonably well because the fund managers who manage them are also more competent at investing.⁴⁶
2. As we discuss below, there can be substantial hidden costs associated with socially responsible investing that might be unaccounted for in financial performance measures: in public markets there are fees for building and managing screened portfolios, and in private markets, investors often pay substantial costs for acquiring the knowledge to

⁴¹ Harry Markowitz, "Portfolio Selection," *The Journal of Finance* 7, no. 1 (1952): 77–91.


⁴² Hong and Kacperczyk, "The Price of Sin"; Julie M. Salaber, "Sin Stock Returns Over the Business Cycle," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, April 24, 2009), <https://papers.ssrn.com/abstract=1443188>.

⁴³ Christopher Geczy, Robert Stambaugh, and David Levin, "Investing in Socially Responsible Mutual Funds," 2005.

⁴⁴ Christophe Revelli and Jean-Laurent Viviani, "Financial Performance of Socially Responsible Investing (SRI): What Have We Learned? A Meta-Analysis," *Business Ethics: A European Review* 24, no. 2 (April 1, 2015): 158–85, <https://doi.org/10.1111/beer.12076>.

⁴⁵ Gunnar Friede, Timo Busch, and Alexander Bassen, "ESG and Financial Performance: Aggregated Evidence from More than 2000 Empirical Studies," *Journal of Sustainable Finance & Investment* 5, no. 4 (October 2, 2015): 210–33, <https://doi.org/10.1080/20430795.2015.1118917>.

⁴⁶ Elias Erragragui and Thomas Lagoarde-Segot, "Solving the SRI Puzzle? A Note on the Mainstreaming of Ethical Investment," *Finance Research Letters* 18 (2016): 32–42.



vet their investments and by providing *pro bono* support to the companies they invest in.

3. Most socially responsible investment screening is based on a ‘best in class’ methodology which does not exclude whole industries, but rather only the ‘worst companies’ within an industry – this allows socially responsible investors indices to stay in sectors such as fast food and fossil energy.⁴⁷ As we saw above, much socially responsible investment is not that strict. When investors divest from harmful but clearly profitable industries, such as tobacco, there is clear evidence that they lose out financially.⁴⁸

Overall, our view is that screening portfolios will reduce expected financial performance. There is no theoretical explanation for the finding from some studies that screening one’s portfolio improves or does not harm financial performance. Studies that investigate this question are prone to pick up a lot of noise rather than signal and are subject to numerous confounders. We therefore think it more likely that the studies are mistaken than that the theoretical argument is wrong.

Second, putting to one side whether portfolios of impact investments will offer worse returns than investors could have got if they didn’t impact invest, we need to consider the costs borne by impact investors themselves. Impact investors face additional costs, which reduce financial performance relative to socially neutral investors. Since impact investors aim to have impact, it makes sense for them to try to track the impact of the enterprises they have supported. Impact evaluation is costly and imposes an additional burden that is not borne by socially neutral investors, who only need to track their financial returns. As well as providing capital, investors can provide non-monetary support in the form of expertise and networks to the companies it has invested in. However, any gains that might accrue from providing this support will be partially offset by the increased cost of providing that support: providing above-market non-monetary support creates above-market costs.

The foregoing argument suggests that the only form of impact investing that stands to do substantial good is *concessionary impact investing in private markets*. At the extreme end, a grant to a company would likely have counterfactual impact because the replaceability concern is eliminated. (It is worth noting that making grants to companies is not very effective because a rational company would immediately distribute the cash to shareholders. Moreover, equity investments give you ownership and control, whereas donations do not.) If you aim for market returns, then you will very likely only displace a socially neutral investor. If you accept below-

⁴⁷ Erragragui and Lagoarde-Segot.

⁴⁸ Madison Marriage, “Dumping Tobacco Cost Norwegian Oil Fund \$1.9bn,” *Financial Times*, April 17, 2016, <https://www.ft.com/content/4b24e8a4-0304-11e6-99cb-83242733f755>.



market returns, then the main risk is displacing another impact investor. We discuss this possibility in Principle 5.

Find companies on the brink of viability

As there is a trade-off between impact and financial returns, the best case for impact investment is likely to provide bridge financing to companies that are on the brink of financial viability. In these cases, your investment would make a genuine difference to the performance of the company.

This raises the question: why would socially beneficial companies not produce great financial returns? If a product is socially beneficial, why would consumers not pay a premium for it, creating a successful business that attracts socially neutral investors? The answer is that the product:


1. **Produces positive externalities** – As we saw above, one way for companies to have impact is to produce positive externalities in the form of reduced pollution, reduced meat consumption etc. Positive externalities are not rewarded in the marketplace, and so the financial value of companies producing them will not be commensurate to their social impact.
2. **Serves poor communities** – We saw in the discussion of Principle 1 that serving poor communities is a promising way to have impact. However, because poor communities have low willingness to pay, there are limited financial rewards for providing goods that benefit them.
3. **Is undervalued by consumers** – Some products are undervalued by consumers. Even though they would be highly beneficial, consumers are unwilling to pay for them. Products in this category include things like the Headspace mental health app.

It is instructive to consider companies that are socially beneficial but are largely *not* characterised by these three features. Some companies, such as Airbnb, Amazon, and Uber, produce a lot of social value for consumers in rich countries with high willingness to pay.⁴⁹ Thus, the financial returns of companies like these are reflective of the social value they produce. However, because the financial returns are so strong, socially neutral investors are likely to be attracted to them anyway. Conversely, companies that produce social benefits that are not fully rewarded in the marketplace will be less attractive to socially neutral investors.

Debt financing

We have thus far focused on the direct effects of equity investment on the capital available to different companies. It is also possible to help or hinder companies by providing or restricting debt. For large firms, debt is easily the largest source of finance. The market for debt is much

⁴⁹ For example, Cohen et al estimate that the UberX service generated \$6.8 billion in consumer surplus in the US in 2015. Peter Cohen et al., “Using Big Data to Estimate Consumer Surplus: The Case of Uber,” Working Paper (National Bureau of Economic Research, September 2016), <https://doi.org/10.3386/w22627>.



clumpier than the decentralised market for equity. As Ansar *et al*/note “five banks—J.P. Morgan, Bank of America Merrill Lynch, Citi, Wells Fargo, Mizuho—have a 40% market share of the global syndicated lending”.⁵⁰ Nonetheless, financial theory suggests that in deep debt markets, other lenders would substitute in for one of these lenders.


Financial depth is typically measured as the proportion of private credit to GDP of a country. Rich countries have deep financial markets, whereas developing countries do not. In emerging markets with shallow financial markets, there is less scope for substitution by other lenders. In these markets, the withdrawal of debt financing could have substantial direct implications for borrowers. This is mainly relevant for major institutional lenders rather than individual investors.

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

Principle 4 discussed how impact investors can have their counterfactual impact eaten up by socially neutral investors. They can also have their counterfactual impact eaten up by other impact investors. To use the analogy mentioned above, impact investors risk not only elbowing socially neutral investors out of the way, but socially responsible ones as well.

Investors calling themselves ‘socially responsible’ now comprise a substantial fraction of total investors (see Figure 2).

Figure 2. The size of Socially Responsible Investing (SRI) in the US in 2016



Source: US SIF Forum for Sustainable and Responsible Investment, [Impact Investing in the US Overview](#), (2016)

There has been a 14-fold increase in investors using the ‘socially responsible’ label in the US since 1995:

⁵⁰ The argument that follows is borrowed entirely from Ansar, Caldecott, and Tilbury, “Stranded Assets,” 32.



Figure 3. SRI and impact investing in the US 1995-2016



Source: US SIF Forum for Sustainable and Responsible Investment, [2016 Report Highlights](#) (2016)

If impact investors work in a crowded field, they risk counterfactual replacement by other impact investors. Due to the counterfactual replaceability effect, the counterfactual effect of the initial investment might not be providing additional capital for the company, but might rather be to free up capital for another impact investor. The size of the counterfactual benefits depends on what exactly the next impact investor is freed up to do. If the marginal impact of that other impact investor's money is as high as yours or if it is only slightly worse but the other impact investor would have to incur evaluation costs whilst you have already incurred them, then it is optimal to go ahead and fund the project. If few good opportunities are available to impact investors, then the benefits of your investment could be fairly small.

This shows that it is important to find problem areas to invest in that are neglected by other impact investors, relative to their scale. This is why investing in the cleantech market in the 2000s was a bad idea, from a social impact point of view – although the market was very large, it was also extremely crowded. Moreover, searching for promising projects is costly and due to diminishing marginal returns, it will be easier to find good opportunities in uncrowded areas.

Principle 6: Exploit an information or network advantage

In order to exploit imperfect markets, identify good opportunities, and provide valuable support to investees, investors need good information and networks relevant to the area they are trying to work in. Some impact investors may already have the in-house knowledge required, reducing the marginal cost of research. For example, the Bill & Melinda Gates Foundation already has significant expertise in global health and development, which it uses to identify promising for-profit projects via its billion-dollar Strategic Investment Fund.⁵¹ Building up from scratch the expertise necessary to reliably identify socially beneficial opportunities across a range of problem areas, such as climate change, agriculture, and global health, is likely to be costly.

⁵¹ <https://pri.gatesfoundation.org/portfolio/>



Moreover, specialising in a handful of areas is in tension with diversification, which is another reason to expect impact investing platforms to sacrifice financial returns. This may also be a reason to have a mixed portfolio of impact investments and socially neutral investments.

The comparative advantage of Founders Pledge members thinking about impact investing lies in the ability to identify promising startups and to provide expertise and networks for these startups. It is important to ensure that these abilities are better than what the average investor would have provided. Again, we need to consider the counterfactual. If a company would be identified and supported anyway by a similarly placed investor, then the counterfactual impact would at least be diminished.

Only working in areas in which you have a comparative information or network advantage may be in tension with risk diversification. There may therefore be a case for spreading your investments between socially neutral opportunities and socially beneficial ones. Having said that, the benefits of diversification are diminishing, and the risk would only rise significantly if the portfolio was highly concentrated in your niche. This means that many investors can take on some extra concentration without much downside.

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.⁵² 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**

⁵² For an introduction to different cause areas see, Effective Altruism, “Introduction to Effective Altruism.”



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.

Informed by Principles 1 and 2, we can draw up an initial list of companies that work in high-impact cause areas and produce positive externalities or serve poor consumers. The table below focuses only on three apparently high-impact cause areas, though other cause areas may also have promise:

Global poverty	Factory farming	Global catastrophic risk
Examples include; <ul style="list-style-type: none"> • Remittances • Banking services • Agricultural technology 	Examples include: <ul style="list-style-type: none"> • Animal product alternatives (e.g. The Good Food Institute) 	Examples include <ul style="list-style-type: none"> • Low carbon energy (e.g. Bill Gates Breakthrough Energy Ventures) • Improved disease diagnostics (to reduce the risk of pandemics)

This is by no means an exhaustive list. The challenge facing impact investors is to find companies that stand to make a real difference in any high value cause area. Even if you succeed in doing this, failure to observe the other principles will greatly reduce or eliminate any potential social impact. The ideal case would be a company working on a novel solution to a high impact problem that is on the brink of financial viability, but is ignored by socially neutral investors and impact investors.



3. How to evaluate impact investing

Impact investors face the dual challenge of measuring both financial returns and social impact. However, evaluating social impact is a difficult challenge and requires evaluating enterprise impact, **and** investment impact, **and** non-monetary impact. It is not enough to show that the business you invested in is socially beneficial, you must also show that your investment or non-monetary support made a difference to its performance.

The ultimate aim of socially motivated investors is to make *cost-effective* investments because the more cost-effective an investment is, the more good it does. The cost-effectiveness of an impact investment is given by its net benefits:

$$\text{Cost-effectiveness} = (\text{benefits of investment}) - (\text{opportunity cost of investment})$$


Cost-effectiveness is the equivalent of return on investment. Just as in business we should aim to get the best ROI, when we are trying to do good, we should aim for the best cost-effectiveness.

Enterprise impact, investment impact and non-monetary impact all only insofar as they affect cost-effectiveness. The importance of the counterfactual is the unifying theme when evaluating these three determinants of impact. Investors who only care about financial returns do not have to consider the counterfactual in the same way. If your aim is to maximise financial returns, then it does not matter whether someone else would have made the same investment if you had decided not to invest; what matters is what *you* got, not what happened to society. In this way, assessing social impact is a very different task to assessing financial performance.

3.1. Evaluating enterprise impact

Impact investing is a hybrid between investing and charitable do-gooding, but impact investors tend to come from the world of investing rather than charity. However, evaluating the social impact of a company poses a completely different challenge to evaluating the financial performance of a company, and goes beyond merely measuring social KPIs.

For context, consider the research time needed to assess the impact of charities. Establishing with reasonable certainty that a global health charity has impact takes months of research. In science and evidence-based development, evaluators survey the academic literature on an intervention carried out by a charity, consult with experts and assess charities' own internal research. Often, researchers rely on rigorous randomised controlled trials to assess whether a charity improved lives. For more indirect interventions, like think tank research and policy advocacy, this kind of rigour is not always possible, but impact evaluation is even more difficult. Evaluators need in-depth knowledge of the organisations involved, the wider ecosystem of actors, and the political and social constraints they operate under. Understanding a vast and complex area like climate change, biomedical research or mental



health, demands extensive engagement with the academic literature and discussion with experts.

Experience with charity impact evaluation shows that:

- It is very difficult to have social impact
- The path from action to impact is usually not straightforward
- Consequently, a serious and analytical approach to impact evaluation is essential

This is not to say that only the most rigorous evidence will do, it is rather that whatever you are evaluating, an impact evaluation based on a few days' research on the social KPIs of a business will probably be inadequate. Our impression is that most impact investors do not carry out rigorous impact evaluations.⁵³

In this section, we discuss five of the most important considerations to take into account when evaluating enterprise impact:

- Focus on outcome not outputs
- The historical marginal approach vs. the growth approach
- Crowdedness of the market
- Price and quality competition
- Costs and benefits to consumers

Outcomes not outputs

How easy is it to tell in advance whether a project will be socially beneficial? If you wanted to improve (say) education, would distributing textbooks and laptops or building libraries be a good approach? Intuitively, one would think the answer is 'yes', but the evidence suggests that when they have been tested these approaches have had no effect:

⁵³ See the discussion of the approach to impact evaluation in the field in Acumen, "Energy Impact Report," 2017, 17, <https://acumen.org/wp-content/uploads/2018/02/Acumen-Energy-Impact-Report.pdf>. Giving Evidence have carried out a brief review of some impact investing platforms and have found limited commitment to impact evaluation. Caroline Fiennes, "Perils of Ensuring an Investment Has Impact," Financial Times, September 24, 2018, <https://www.ft.com/content/a269bd5c-785e-11e8-af48-190d103e32a4>.




Figure 4. The cost-effectiveness of education interventions



Source: JPAL, [Improving learning by increasing motivation, targeting instruction, and addressing school governance](#) (2013)

The same is true in all other domains impact investors will encounter: it is rarely trivial to show that something genuinely produces social impact. Indeed, the evidence suggests that of social programmes that have been rigorously tested, most do not work.⁵⁴ (To find out for yourself whether you can identify which social programmes will work and which will fail, take [this test](#) by the impact research organisation 80,000 Hours.)

This means it is important to focus on outcomes rather than outputs. *Outcomes* are things we actually care about, such as welfare improvements, lives saved, increases in consumption, or tonnes of CO₂ averted.⁵⁵ *Outputs* measure the things we use to try to achieve those outcomes. Outputs can be measured in terms of number of solar panels sold, number of clean cookstoves distributed, number of learning aids sold, savings to consumers, and so on. Measuring outputs

⁵⁴ 80,000 Hours, "Is It Fair to Say That Most Social Programmes Don't Work?," accessed September 27, 2018, <https://80000hours.org/articles/effective-social-program/>.

⁵⁵ See 'How we think about charity' at www.founderspledge.com/research



is useful for measuring outcomes, but the path from output to outcome is often counterintuitive.⁵⁶

Many impact evaluation tools for impact investing, such as IRIS, focus on outputs not outcomes and so provide limited information on social impact.⁵⁷

Once you have clarified the outcome you are aiming to achieve, you need to decide on a metric relevant to that outcome. Examples of key metrics include:

- **Global health** – Disability Adjusted Life Years (DALYs) are the most widely accepted and used metric to measure health outcomes. A DALY can be thought of as one lost year of healthy life, and it is used to measure the “burden” of a disease. The worse the effect of a disease on one’s life, the higher its DALY weighting on a scale from 0 (perfect health) to 1 (dead). For example, mild distance vision impairment has a DALY weight of 0.003 and acute schizophrenia has a weight of 0.778. Evaluators of social programmes often rank interventions in terms of the cost per DALY averted.
- **Poverty** – Economic benefits are often calculated in dollar terms, but money is not intrinsically valuable, and its value depends on how much you already have. If you earn \$700 a year, then an additional \$1,000 would be life changing, whereas it would be trivial if you earn \$1 million a year. Money has *diminishing marginal utility*. One common way of modelling this is to assume that the welfare returns of money are logarithmic.⁵⁸ Doubling your income from any level produces the same welfare increase no matter where you start: moving from \$1,000 to \$2,000 is just as good as moving from \$20,000 to \$40,000.⁵⁹
- **Climate change** – The most useful metric to calculate the cost-effectiveness of work on climate change is *tonnes of CO₂-equivalent averted*. CO₂-equivalent (CO₂e) is a metric that measures the warming effect of different greenhouse gases over 100 years in terms of the functionally equivalent amount of CO₂. In our recent climate change report, we calculated the past cost-effectiveness of our recommended charities in terms of “\$ per tonne of CO₂e averted”.⁶⁰


⁵⁶ Note that the line between the two is sometimes blurry. CO₂ is a good case. We do not ultimately care about the amount of CO₂ in the atmosphere, but rather about the effects of additional CO₂ on human welfare via global warming. However, tonnes of CO₂ are a strong proxy for damage to human welfare and the connection between the two is relatively firm. The main risk of using output measures arises when the link between the output and the outcome is not well-established.

⁵⁷ Brest and Born, “Unpacking the Impact in Impact Investing.”

⁵⁸ Toby Ord, “The Value of Money Going to Different Groups,” Centre for Effective Altruism, May 2017, <https://www.centreforeffectivealtruism.org/blog/the-value-of-money-going-to-different-groups/>.

⁵⁹ See GiveWell’s cost-effectiveness analyses at <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

⁶⁰ We found that our recommended charities have *in the past* averted a tonne of CO₂e for less than \$5. See www.founderspledge.com/research



These metrics can be used to assess the impact of for-profit companies. Converting across heterogeneous metrics is a difficult though sometimes surmountable problem. Our research partner, GiveWell have estimated a conversion factor between health benefits and consumption benefits.⁶¹ Conversion in other cases, for example from tonnes of CO₂e into DALYs, is much more difficult because most of the health costs of climate change will be realised many decades hence, their magnitude is extremely uncertain, and evaluating them depends on difficult value questions about the moral importance of future people.

The historical extrapolation approach vs. the growth approach

There are two divergent approaches to assessing the enterprise impact of a company:⁶²

1. **The historical extrapolation approach:** What social impact minus cost did the company achieve over the last year? Can we expect this rate of return to hold up at the margin in the next year?
2. **The growth approach:** Does the project have a high growth rate, a large potential market, address an important problem, have an effective solution to this problem, and a great team?

The suitability of these two approaches depends on the type of business under consideration. For established and stable businesses like Coca Cola, you can assume that performance last year approximates the performance next year, and that the performance in the near future is a significant fraction of the total value. So, you can use proxies that look at recent history and try to project forward. Thus, for stable and established businesses, the historical extrapolation approach may be the best bet.

However, this approach is not suited to evaluating startups. To see why, consider the approach you should take if your sole aim is profit. Early on, few startups are profitable. In the first two years, Google did not make a profit, and until recently, Amazon had at most made a small profit. A socially neutral investor who took the historical approach would have erroneously concluded that these were not good investment opportunities. But this would neglect the long-term growth prospects of these two companies. Indeed, these companies' strategy of foregoing short-term profit makes sense – if you can grow by 100x, then it makes sense to invest all your spare money in growth rather than turn a profit in the short term.

Similarly, since the social benefits of a beneficial company will scale with its sales, it would be a mistake to use the historical extrapolation approach to assess the potential social impact of a

⁶¹ See their cost-effectiveness analyses at <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>

⁶² The argument here is indebted to Benjamin Todd, "Take the Growth Approach to Evaluating Startup Non-Profits, Not the Marginal Approach," 80,000 Hours, November 26, 2015, <https://80000hours.org/2015/11/take-the-growth-approach-to-evaluating-startup-non-profits-not-the-marginal-approach/>. For more on the growth approach, see Blake Masters and Peter Thiel, *Zero to One: Notes on Start Ups, or How to Build the Future* (London: Virgin Books, 2015).



startup. Doing this would fail to account for the future growth of the startup, which is where most of its social value comes from. Much of the value in a VC-style impact investing portfolio will come from a handful of 'big wins': startups that grow fast and come to dominate a market, producing large social benefits in the process. We discuss the importance of big wins in more detail in section 3.2.


You can use the following heuristics to test the social impact of for-profit businesses using the growth approach:

1. Does the company have a replicable approach to solving an important problem?
2. Do they have a method to do this at a much larger scale, such that if done, the value created would be much larger than costs?
3. If the company reached its entire addressable market, how good would it be?
4. Does the company have a competitive advantage over its rivals? Is it in a crowded market?
5. Is the company growing quickly, and is their growth rate plausibly exponential? (Bear in mind that assessing the growth rate isn't appropriate at the earliest stages of the project, before you have product-market fit).
6. Is the team smart and determined? Do they have a track record and the relevant skills and experience?

Crowdedness of market

As discussed above, the crowdedness of the market is a crucial determinant of a company's social impact. A naïve way to assess the social benefits of a business would be to calculate the social benefit of the product and then multiply by the number of products sold. For a solar power company, for example, one would look at how many solar panels were sold and multiply by the emissions benefit of one solar panel to calculate the total emissions benefit of the company. But as we have seen, the social benefits are actually those given by the naïve estimate minus what would have happened had the company not existed. The more crowded a market, the less counterfactual impact a company can expect to have. Of course, although enterprise impact is lower than one might naïvely assume, it can still often be substantial.

The counterfactual consideration constitutes a crucial difference between assessing the social impact of a business and assessing its financial value. When assessing the financial value of a business like Airbnb, we can completely ignore the alternative home sharing businesses that would later have arisen had Airbnb not existed. We cannot do this when assessing Airbnb's social impact. Imagine that Airbnb came up with its product only a month before a rival would have come up with something identical. Winning the race would have made Airbnb enormously financially successful, but the fact that it was so easily replaceable reduces its social impact. (Note that this hypothetical is meant to illustrate the concept of counterfactual impact; it is not an accurate description of what would have happened if Airbnb had not existed).



Price and quality effects

We saw above that a company's counterfactual impact can be lower than one might naïvely assume. It can also be higher than one might naïvely assume. Companies can have impact by instigating competition on price and quality. For example, the benefits of Uber's service extend beyond the benefits provided to Uber customers because Uber competes with other taxi firms, causing those firms to compete on price and quality. Therefore, the customers of other taxi firms also benefit. This effect is very difficult to estimate.

This is an advantage that for-profits have over non-profits. Due to the incentives in competitive markets, competition leads to benefits for consumers in the whole marketplace. Competition between non-profits, by contrast, does not create comparably strong market-wide incentives to improve the quality of the service because, as we discuss in section 4.2, the structure of incentives and information in the non-profit marketplace is not ideal.

Costs and benefits to consumers

Any sensible evaluation of a social enterprise will consider benefits to consumers, but it is also important to consider costs to consumers. A key difference between for-profits and non-profits is that non-profits usually do not charge for their services. This could be especially important for firms serving very poor communities where any purchase would be a sizeable fraction of household expenditure. In this case, charging a fee might have dramatic effects on usage.

3.2. Evaluating investment impact

As we have seen, investing in a business with high enterprise impact is no guarantee of real impact: an impact investor cannot produce social benefits unless they have investment impact or non-monetary impact. The key factors to consider when assessing investment impact are:

- Consider the counterfactual
- Think on the margin
- Consider all costs
- Focus on average, not typical, performance

Consider the counterfactual

We discussed the conditions in which investors can have counterfactual investment impact in section 2. In brief, the 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?
 - a. Have you invested in an efficient market?
 - b. 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?
 - a. Is the space crowded with other socially motivated investors?



b. 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.⁶³ It will always be difficult to quantify this probability. Nevertheless, without trying to quantify it, an evaluation would be incomplete.

Think on the margin

Impact investors should think about the long-run effect of their marginal donation, rather than about the total or average benefits produced by a company invested in. Suppose you invest \$1 million to acquire 30% of a company, and the company averts 50,000 tonnes of CO₂ from the time of your investment onward. It would be a mistake to calculate the benefits of your investment as a 50,000 tonne of CO₂ return on \$1 million invested. The relevant question is what difference your investment made to the long-term social impact of the company. If the company was already well capitalised, then the investment might not have had much effect on their output, but if the company would have collapsed without the investment, then the marginal effect of the donation would have been much larger.

Consider all costs

When assessing their own cost-effectiveness, it is important for impact investors to consider not only the costs of the investment itself, but also the other costs they bear. There are staff costs involved in finding good opportunities and in providing non-monetary support to companies. We should also consider the *opportunity cost* of staff time, which is what the staff could have been doing had they not worked for the impact investor. It might be that staff could otherwise be involved in highly socially beneficial activities or successful businesses. Calculating these costs is likely to be difficult but is again important.

Focus on average, not typical, performance

We should expect most investments in startups to fail; the performance of the typical investment will be poor. Does this mean that investing in startups is a bad idea? No, because we should focus on the average performance – the expected value of the portfolio – rather than the performance of the typical company in the portfolio. The average performance of the portfolio can be pulled up by a handful of very big wins. This is because returns in venture

⁶³ Paddy Carter, Nicholas van Sijpe, and Raphael Calel, “The Elusive Quest for Additionality” (Center for Global Development, September 2018), <https://www.cgdev.org/publication/elusive-quest-for-additionality>.




capital and angel investing are distributed according to a *power law*, rather than being *normally distributed*.⁶⁴

The height of American women is *normally distributed*. The average (mean) American woman is 5 foot 4 and so is the typical American woman, and 95% of American women are within two standard deviations of the mean – between 4 foot 10 and 5 foot 10. In normal distributions, the mean is the same as the median (the typical event), so if VC returns were normally distributed, the fate of the typical project would be a good guide to the success of the overall portfolio. Another notable thing about normal distributions is that extreme events are highly unlikely. For instance, we will never encounter a woman who is 23 standard deviations from the mean, or 11 feet tall.

Not all things follow a normal distribution; stock market prices, the size of cities, the magnitude of earthquakes, and the death toll of wars all seem to follow a *power law distribution*.⁶⁵ Figure 5 compares a normal distribution and a power law distribution. Power law distributions have fat right tails such that extreme events occur much more frequently. Returns in VC and angel investing seem to follow a power law.⁶⁶ If returns were normally distributed, we would almost never expect to see unicorns such as Uber, Klarna, Airbnb and Amazon – these are the equivalent of seeing an 11-foot-tall woman. Big wins such as these can pull up the average value of a portfolio even if most investments fail.

Figure 5. A normal distribution and a power law distribution



Source: Jerry Neumann, '[Power laws in venture](#)' (2015)

⁶⁴ Jerry Neumann, "Power Laws in Venture," *Reaction Wheel* (blog), June 25, 2015, <http://reactionwheel.net/2015/06/power-laws-in-venture.html>.

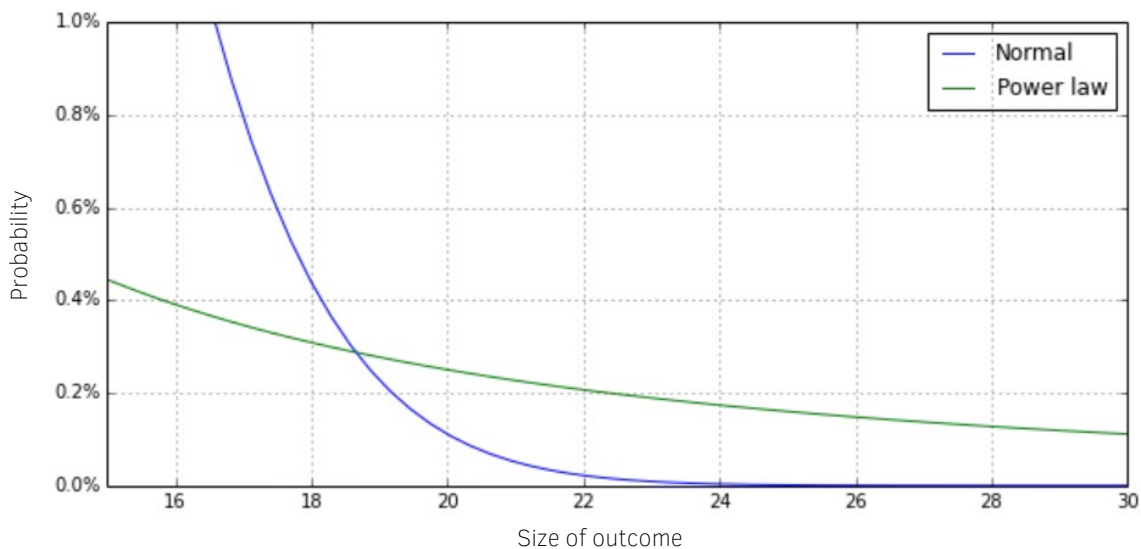
⁶⁵ Neumann.

⁶⁶ Neumann.



Figure 6 zooms in on the right-hand tail of the distribution, showing how a normal distribution drops off much faster than a power law:

Figure 6. The right tails of a normal distribution and a power law distribution



Source: Jerry Neumann, '[Power laws in venture](#)' (2015)

This is relevant to social impact of companies as well as their financial returns because the social impact of a socially beneficial product scales with the units sold.⁶⁷ Socially motivated investors should assess success by evaluating the average social impact of the whole portfolio rather than the typical investment.


Here, it might be good to discuss one potential confusion related to whether market rate returns are possible while having a social impact. Due to power law distribution of returns, any one company in a social impact investing portfolio can achieve market- or above-market- rate returns while at the same time generate social impact. However, the average of the whole portfolio will be at market rate in the long-term in expectation, if the portfolio is to have social impact. A good example here is the PRI initiative of the Gates Foundation which highlights some companies that have generated both social impact and above market rate returns, yet their overall portfolio is instructed to make 90 cents on the dollar - i.e. lose some money - in the long-term in order for it to have social impact.

3.3. Evaluating non-monetary impact

In section 1, we identified four main ways in which impact investors can have non-monetary impact:

1. Finding and promoting impact investment opportunities
2. Providing technical assistance and access to networks

⁶⁷ In many cases, we should expect the benefits to scale linearly, but in others they might scale super-linearly due to network effects.



3. Securing and protecting the enterprise's social mission
4. Gaining publicity for an advocacy campaign

The theme of the previous two sections – the importance of the counterfactual – is again central for assessing non-monetary impact.

Finding and promoting impact investment opportunities to other socially neutral or socially motivated investors

To promote opportunities to socially neutral investors, impact investors have to show that the opportunity has strong private returns. Socially motivated investors will be more interested in the combination of private and social returns. The degree to which you could promote investment opportunities to impact investors who would not have found them anyway depends on your knowledge advantage and on the crowdedness of the market.

Providing expertise and access to networks

Since we are focused on the counterfactual, the question we should ask here is not “did I provide useful expertise and access to networks?”, but “did I provide better expertise and access to networks than the company would have otherwise received?” Even if you provide useful technical assistance to a company, your non-monetary impact could be *negative* if another potential investor would have provided even better expertise and access to networks. Thus, investors need to consider whether what they offer is better than the market average.

Securing and protecting the enterprise's social mission

Socially valuable businesses can sometimes drift from their core values and mission. Owning equity in a company gives you some control over its future and so allows you to prevent this from happening.⁶⁸ For example, you could ensure that a pharmaceutical company stays focused on reducing the cost of vaccines as much as possible and does not get sidetracked by a less socially valuable mission. Socially responsible investors could also form a voting bloc to influence the direction of potentially harmful companies.

Gaining publicity for an advocacy campaign

We argued in section 2 that it is usually difficult for investors to affect the stock price of publicly traded companies. This means that divestment campaigns are unlikely to have much direct effect on their targets' stock price. However, they may have more impact by stigmatising the company, alienating customers and suppliers, and opening up the threat of future regulation.⁶⁹ Ansar *et al*/note that “in almost every divestment campaign we reviewed from adult services to Darfur, from tobacco to South Africa, divestment campaigns were successful

⁶⁸ See Paul Brest, “Investing for Impact with Program-Related Investments,” accessed August 6, 2018, https://ssir.org/articles/entry/investing_for_impact_with_program_related_investments.

⁶⁹ Ansar, Caldecott, and Tilbury, “Stranded Assets,” 14.



in lobbying for restrictive legislation affecting stigmatised firms".⁷⁰ Thus, the indirect effects of divestment campaigns are likely to be much greater than the direct effects.

This being said, we think that, given the financial opportunity costs of socially responsible investing, *for people aiming to have maximal social impact*, socially neutral investing to donate to effective charities will usually be more effective. We discuss this in more detail in section 4.3.

⁷⁰ Ansar, Caldecott, and Tilbury, 14.



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:

1. **Donating now** – Making grants to high-impact charities
2. **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-\$10.⁷¹ 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).⁷²


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

⁷¹ See www.founderspledge.com/research

⁷² See <https://www.givewell.org/how-we-work/our-criteria/cost-effectiveness/cost-effectiveness-models>



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 2015.⁷³ 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.⁷⁴ 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 2017.⁷⁵ 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.⁷⁶ They could also have followed the approach of 'mission hedging' and used the money to advocate for regulation of tobacco products.⁷⁷ The money that could have been donated would have matched that currently donated by Gates and Bloomberg, the leading philanthropists working in this space.⁷⁸ 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.

⁷³ Marriage, "Dumping Tobacco Cost Norwegian Oil Fund \$1.9bn."

⁷⁴ Ansar, Caldecott, and Tilbury, "Stranded Assets," 61.

⁷⁵ "WEI 2018," accessed October 22, 2018, <https://www.iea.org/wei2018/>.

⁷⁶ This is assuming that if the money was spent on the most cost-effective interventions, it could have saved a life for \$11,500, as per Dean T. Jamison et al., "Global Health 2035: A World Converging within a Generation," *The Lancet* 382, no. 9908 (2013): 1921.

⁷⁷ Roth Tran, "Divest, Disregard, or Double Down?"; Hauke Hillebrandt, "A Generalized Strategy of 'Mission Hedging': Investing in 'evil' to Do More Good," Effective Altruism Forum, February 2018, <https://forum.effectivealtruism.org/posts/iZp7TtZdFyW8eT5dA/a-generalized-strategy-of-mission-hedging-investing-in-evil>.

⁷⁸ Donald G. McNeil Jr, "Bill Gates and Michael Bloomberg to Fund Anti-Smoking Campaign," *The New York Times*, July 23, 2008, sec. Americas, <https://www.nytimes.com/2008/07/23/world/americas/23iht-tobacco.4.14730339.html>.



To be a little bit more precise and technical: multi-objective optimisation is harder than single-objective optimisation,⁷⁹ 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*.⁸⁰ ‘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.⁸¹

⁷⁹ “Multi-Objective Optimization,” *Wikipedia*, July 4, 2018, https://en.wikipedia.org/w/index.php?title=Multi-objective_optimization&oldid=848847142.

⁸⁰ The following is borrowed from Tyler Cowen, “Public Goods,” *Econlib*, accessed August 21, 2018, <https://www.econlib.org/library/Enc/PublicGoods.html>.

⁸¹ Edward L. Glaeser, Joseph Gyourko, and Raven Saks, “Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices,” *The Journal of Law and Economics* 48, no. 2 (October 1, 2005): 331–69, <https://doi.org/10.1086/429979>.



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.⁸² 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.⁸³ 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,⁸⁴ 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.⁸⁵ 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-

⁸² See https://www.givewell.org/international/technical/programs/insecticide-treated-nets#Free_vs_cost-recovering_distributions

⁸³ Fiennes, "Perils of Ensuring an Investment Has Impact."

⁸⁴ Giving USA, "Giving USA 2017: Total Charitable Donations Rise to New High of \$390.05 Billion," accessed September 26, 2018, <https://givingusa.org/giving-usa-2017-total-charitable-donations-rise-to-new-high-of-390-05-billion/>.

⁸⁵ <https://beta.observablehq.com/@henryaj/where-we-donate-and-why-it-matters>



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*.⁸⁶ 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.⁸⁷ 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

⁸⁶ For an overview, see Richard Steinberg, "Principal-Agent Theory and Nonprofit Accountability," in *Comparative Corporate Governance of Non-Profit Organizations*, ed. Klaus J. Hopt and Thomas von Hippel (Cambridge University Press, 2010).

⁸⁷ For discussion see for example Carl Shulman, "Annual 'splitting' of Funding Gaps Can Be Partial Funging When Gaps Carry over across Years," *Reflective Disequilibrium* (blog), August 2016, <http://reflectivedisequilibrium.blogspot.com/2016/08/annual-splitting-of-funding-gaps-can-be.html>.



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.⁸⁸ Therefore, careful philanthropists should be able to significantly reduce the replaceability concern.

4.3. Should you impact invest?

We have 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. As discussed above, the constraints set down by the six principles are stringent: they narrow down the possible space of effective impact investments significantly and make clear the trade-offs between returns and impact. However, in some cases, careful impact investors may be able to satisfy the six principles, and if so, impact investing could 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 in many cases 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 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 in many cases, 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.⁸⁹ In their 2017 [Energy Impact Report](#), Acumen state that they invested \$22.1 million in companies that averted 6.4 million tonnes of CO₂,⁹⁰ and they are on track to get all of their investment back.⁹¹ 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

⁸⁸ See www.founderspledge.com/methodology

⁸⁹ See <https://acumen.org/lean-data/>

⁹⁰ Acumen, "Energy Impact Report," 10.

⁹¹ Acumen, 13.



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.⁹² It claims that its overall \$22.1 million investment has crowded in \$104.5 million from other investors,⁹³ 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 (2017-27)	Total tonnes of CO ₂ averted in 2027
Acumen Energy Fund	\$22.1m	6.4 million tonnes	\$22.1m	6.4m	12.8m tonnes
Donate now	\$22.1m	22m	\$0	0	22m tonnes
Invest to give	\$22.1m	0	\$35m	35m	35m tonnes

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.

⁹² Acumen, 38–39.

⁹³ Acumen, 11.



- 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).



5. Recommended reading

There are several high-quality articles relevant to impact investing:

1. Paul Brest and Kelly Born, “Unpacking the Impact in Impact Investing,” Stanford Social Innovation Review, August 2013,
https://ssir.org/articles/entry/unpacking_the_impact_in_impact_investing
 - An excellent introduction to the key concepts of impact investing and how to have counterfactual impact in the field.
2. Atif Ansar, Ben Caldecott, and James Tilbury, [*Stranded assets and the fossil fuel divestment campaign: what does divestment mean for the valuation of fossil fuel assets?*](#), Smith School of Enterprise and the Environment, University of Oxford, October 2013
 - An excellent overview of the theory and evidence surrounding the direct and indirect effects of divestment campaigns.
3. Center for Global Development, [*The Elusive Quest for Additionality*](#), (2018)
 - Discusses the challenges associated with assessing additionality through investments in the private sector.
4. Jerry Neumann, “Power Laws in Venture,” *Reaction Wheel* (blog), June 25, 2015,
<http://reactionwheel.net/2015/06/power-laws-in-venture.html>
 - An introduction to the idea of power law distributions and how they apply to returns in VC.
5. Effective Altruism, “Introduction to Effective Altruism,” accessed August 24, 2018,
<https://www.effectivealtruism.org/articles/introduction-to-effective-altruism>
 - Explains the importance of choosing a high-impact cause and discusses some promising cause areas.

